



Robert H. Shumway
David S. Stoffer

Time Series Analysis and Its Applications

With R Examples

Third Edition

Robert H. Shumway
David S. Stoffer

Time Series Analysis and Its Applications

With R Examples

Third Edition

Blue Printing 12.2015

Prof. Robert H. Shumway
Department of Statistics
University of California
Davis, California
USA

Prof. David S. Stoffer
Department of Statistics
University of Pittsburgh
Pittsburgh, Pennsylvania
USA

ISSN 1431-875X
ISBN 978-1-4419-7864-6 e-ISBN 978-1-4419-7865-3
DOI 10.1007/978-1-4419-7865-3
Springer New York Dordrecht Heidelberg London

© Springer Science+Business Media, LLC 2011

Preface to the Blue Printing

The blue printing is a revision of the third edition that will be updated periodically; the printing date is on the title page. No errata will be kept; if you see a typo, send me (DSS) an email and I'll fix it at the next update. The blue printing follows the original yellow version but with internal links, various fixes, additions, and updates. Hopefully you downloaded this for free from <http://www.stat.pitt.edu/stoffer/tsa3/>. All the code is included here and probably won't be posted on the website.

The changes so far are:

- Chapter 1: Minor changes to some examples. Added definition of trend stationarity. Changed a few graphics. Added the concept of prewhitening when using CCF. Explained non-negative definite of the ACF in more detail (used to only be in a problem).
- Chapter 2: Rewrote some of the regression review. Changed smoothing examples from lame mortality data example to SOI and finding El Niño (some examples removed). Expanded on the lagged regression example and carried it on to Chapter 3.
- Chapter 3: Removed iid normal from definition of ARMA models. Added section on regression with ARMA errors and the corresponding problems (moved from Chapter 5). Changed/added some examples in the seasonal ARMA section. Added an example on why you should NOT use `auto.arima` from the `forecast` package. Carried lagged regression through to the end (lagged regression with autocorrelated errors).
- Chapter 4: Minor changes to some examples. Added star magnitude example and discussed modulated series. Moved some of the filtering section forward for easier access to information when needed. Removed reliance on `spec.pgram` (from `stats` package) to `mvspec` (from `astsa` package) so we don't have to spend pages explaining the quirks of `stats`. Removed section on wavelets because there are so many accessible texts available.

Chapter 5: Moved regression with ARMA errors to Chapter 3. Some minor changes to examples.

Chapter 6: Same basic layout, but some sections have been moved. The Bayesian section (now §6.11) is completely rewritten and is on linear Gaussian state space models only. The nonlinear stuff is removed because it's old and the newer stuff is in Douc, Moulines, and Stoffer (2014). Also added a smoothing splines section (now §6.5). Polished and fixed some other examples.

Chapter 7: No change, it's perfect.

Problems: Moved regression with ARMA errors problems to Chapter 3. Removed Chapter 4 wavelet problems. Tweaked, rewrote, improved, revised some other problems, but the overall ordering and coverage is roughly the same.

Appendices: Fixed some typos, but otherwise unchanged.

Internal links are [dark red](#), external links are [magenta](#), R code is [blue](#), comments are [green](#), and output is [purple](#).

Preface to the Third Edition

The goals of this book are to develop an appreciation for the richness and versatility of modern time series analysis as a tool for analyzing data, and still maintain a commitment to theoretical integrity, as exemplified by the seminal works of Brillinger (1975) and Hannan (1970) and the texts by Brockwell and Davis (1991) and Fuller (1995). The advent of inexpensive powerful computing has provided both real data and new software that can take one considerably beyond the fitting of simple time domain models, such as have been elegantly described in the landmark work of Box and Jenkins (1970). This book is designed to be useful as a text for courses in time series on several different levels and as a reference work for practitioners facing the analysis of time-correlated data in the physical, biological, and social sciences.

We have used earlier versions of the text at both the undergraduate and graduate levels over the past decade. However, there is now an “EZ Green Edition” on the website <http://www.stat.pitt.edu/stoffer/tsa3/> that is more accessible to students with only basic skills. Our experience is that an undergraduate course can be accessible to students with a background in regression analysis and may include §1.1–§1.6, §2.1–§2.3, the results and numerical parts of §3.1–§3.9, and briefly the results and numerical parts of §4.1–§4.6. At the advanced undergraduate or master’s level, where the students have some mathematical statistics background, more detailed coverage of the same sections, with the inclusion of §2.4 and extra topics from Chapter 5 or Chapter 6 can be used as a one-semester course. Often, the extra topics are chosen by the students according to their interests. Finally, a two-semester upper-level graduate course for mathematics, statistics, and engineering graduate students can be crafted by adding selected theoretical appendices. For the upper-level graduate course, we should mention that we are striving for a broader but less rigorous level of coverage than that which is attained by Brockwell and Davis (1991), the classic entry at this level.

The major difference between this third edition of the text and the second edition is that we provide R code for almost all of the numerical examples. An R package called `astsa` is provided for use with the text; see §R.2 for details. R code is provided simply to enhance the exposition by making the numerical examples reproducible.

We have tried, where possible, to keep the problem sets in order so that an instructor may have an easy time moving from the second edition to the third edition. However, some of the old problems have been revised and there are some new problems. Also, some of the data sets have been updated. We added one section in Chapter 5 on unit roots and enhanced some of the presentations throughout the text. The exposition on state-space modeling, ARMAX models, and (multivariate) regression with autocorrelated errors in Chapter 6 have been expanded. In this edition, we use standard R functions as much as possible, but we use our own scripts (included in [astsa](#)) when we feel it is necessary to avoid problems with a particular R function; these problems are discussed in detail on the website for the text under R Issues.

We thank John Kimmel, Executive Editor, Springer Statistics, for his guidance in the preparation and production of this edition of the text. We are grateful to Don Percival, University of Washington, for numerous suggestions that led to substantial improvement to the presentation in the second edition, and consequently in this edition. We thank Doug Wiens, University of Alberta, for help with some of the R code in Chapters 4 and 7, and for his many suggestions for improvement of the exposition. We are grateful for the continued help and advice of Pierre Duchesne, University of Montreal, and Alexander Aue, University of California, Davis. We also thank the many students and other readers who took the time to mention typographical errors and other corrections to the first and second editions. Finally, work on the this edition was supported by the National Science Foundation while one of us (D.S.S.) was working at the Foundation under the Intergovernmental Personnel Act.

Davis, CA
Pittsburgh, PA
September 2010

Robert H. Shumway
David S. Stoffer

Contents

Preface to the Blue Printing	vi
Preface to the Third Edition	viii
1 Characteristics of Time Series	1
1.1 Introduction	1
1.2 The Nature of Time Series Data	3
1.3 Time Series Statistical Models	10
1.4 Measures of Dependence: Autocorrelation and Cross-Correlation	16
1.5 Stationary Time Series	20
1.6 Estimation of Correlation	27
1.7 Vector-Valued and Multidimensional Series	33
Problems	38
2 Time Series Regression and Exploratory Data Analysis	46
2.1 Classical Regression in the Time Series Context	46
2.2 Exploratory Data Analysis	55
2.3 Smoothing in the Time Series Context	68
Problems	73
3 ARIMA Models	79
3.1 Introduction	79
3.2 Autoregressive Moving Average Models	79
3.3 Difference Equations	93
3.4 Autocorrelation and Partial Autocorrelation	98
3.5 Forecasting	105
3.6 Estimation	117
3.7 Integrated Models for Nonstationary Data	136
3.8 Building ARIMA Models	140
3.9 Regression with Autocorrelated Errors	148

3.10	Multiplicative Seasonal ARIMA Models	151
	Problems	161
4	Spectral Analysis and Filtering	170
4.1	Introduction	170
4.2	Cyclical Behavior and Periodicity	172
4.3	The Spectral Density	178
4.4	Periodogram and Discrete Fourier Transform	185
4.5	Nonparametric Spectral Estimation	193
4.6	Parametric Spectral Estimation	208
4.7	Multiple Series and Cross-Spectra	212
4.8	Linear Filters	217
4.9	Lagged Regression Models	223
4.10	Signal Extraction and Optimum Filtering	227
4.11	Spectral Analysis of Multidimensional Series	232
	Problems	235
5	Additional Time Domain Topics	245
5.1	Introduction	245
5.2	Long Memory ARMA and Fractional Differencing	245
5.3	Unit Root Testing	254
5.4	GARCH Models	257
5.5	Threshold Models	265
5.6	Lagged Regression: Transfer Function Modeling	269
5.7	Multivariate ARMAX Models	274
	Problems	287
6	State-Space Models	290
6.1	Introduction	290
6.2	Filtering, Smoothing, and Forecasting	295
6.3	Maximum Likelihood Estimation	305
6.4	Missing Data Modifications	314
6.5	Smoothing Splines and the Kalman Smoother	319
6.6	Structural Models: Signal Extraction and Forecasting	322
6.7	State-Space Models with Correlated Errors	326
6.7.1	ARMAX Models	327
6.7.2	Multivariate Regression with Autocorrelated Errors	328
6.8	Bootstrapping State-Space Models	331
6.9	Dynamic Linear Models with Switching	337
6.10	Stochastic Volatility	349
6.11	Bayesian Analysis of State Space Models	357
	Problems	367

7	Statistical Methods in the Frequency Domain	374
7.1	Introduction	374
7.2	Spectral Matrices and Likelihood Functions	378
7.3	Regression for Jointly Stationary Series	379
7.4	Regression with Deterministic Inputs	388
7.5	Random Coefficient Regression	397
7.6	Analysis of Designed Experiments	402
7.7	Discriminant and Cluster Analysis	416
7.8	Principal Components and Factor Analysis	433
7.9	The Spectral Envelope	449
	Problems	464
Appendix A	Large Sample Theory	470
A.1	Convergence Modes	470
A.2	Central Limit Theorems	477
A.3	The Mean and Autocorrelation Functions	481
Appendix B	Time Domain Theory	490
B.1	Hilbert Spaces and the Projection Theorem	490
B.2	Causal Conditions for ARMA Models	494
B.3	Large Sample Distribution of the AR(p) Conditional Least Squares Estimators	496
B.4	The Wold Decomposition	499
Appendix C	Spectral Domain Theory	501
C.1	Spectral Representation Theorem	501
C.2	Large Sample Distribution of the DFT and Smoothed Periodogram	505
C.3	The Complex Multivariate Normal Distribution	515
Appendix R	R Supplement	521
R.1	First Things First	521
R.2	astsa	521
R.3	Getting Started	522
R.3.1	Basic Statistics	525
R.3.2	Packages	526
R.3.3	Word	527
R.4	Time Series Primer	527
	References	535
	Index	547

Chapter 1

Characteristics of Time Series

1.1 Introduction

The analysis of experimental data that have been observed at different points in time leads to new and unique problems in statistical modeling and inference. The obvious correlation introduced by the sampling of adjacent points in time can severely restrict the applicability of the many conventional statistical methods traditionally dependent on the assumption that these adjacent observations are independent and identically distributed. The systematic approach by which one goes about answering the mathematical and statistical questions posed by these time correlations is commonly referred to as time series analysis.

The impact of time series analysis on scientific applications can be partially documented by producing an abbreviated listing of the diverse fields in which important time series problems may arise. For example, many familiar time series occur in the field of economics, where we are continually exposed to daily stock market quotations or monthly unemployment figures. Social scientists follow population series, such as birthrates or school enrollments. An epidemiologist might be interested in the number of influenza cases observed over some time period. In medicine, blood pressure measurements traced over time could be useful for evaluating drugs used in treating hypertension. Functional magnetic resonance imaging of brain-wave time series patterns might be used to study how the brain reacts to certain stimuli under various experimental conditions.

Many of the most intensive and sophisticated applications of time series methods have been to problems in the physical and environmental sciences. This fact accounts for the basic engineering flavor permeating the language of time series analysis. One of the earliest recorded series is the monthly sunspot numbers studied by Schuster (1906). More modern investigations may center on whether a warming is present in global temperature measurements or whether levels of pollution may influence daily mortality in Los Angeles. The modeling of speech series is an important problem related to the efficient transmission of voice recordings. Common features in a

time series characteristic known as the power spectrum are used to help computers recognize and translate speech. Geophysical time series such as those produced by yearly depositions of various kinds can provide long-range proxies for temperature and rainfall. Seismic recordings can aid in mapping fault lines or in distinguishing between earthquakes and nuclear explosions.

The above series are only examples of experimental databases that can be used to illustrate the process by which classical statistical methodology can be applied in the correlated time series framework. In our view, the first step in any time series investigation always involves careful scrutiny of the recorded data plotted over time. This scrutiny often suggests the method of analysis as well as statistics that will be of use in summarizing the information in the data. Before looking more closely at the particular statistical methods, it is appropriate to mention that two separate, but not necessarily mutually exclusive, approaches to time series analysis exist, commonly identified as the time domain approach and the frequency domain approach.

The time domain approach is generally motivated by the presumption that correlation between adjacent points in time is best explained in terms of a dependence of the current value on past values. The time domain approach focuses on modeling some future value of a time series as a parametric function of the current and past values. In this scenario, we begin with linear regressions of the present value of a time series on its own past values and on the past values of other series. This modeling leads one to use the results of the time domain approach as a forecasting tool and is particularly popular with economists for this reason.

One approach, advocated in the landmark work of Box and Jenkins (1970; see also Box et al., 1994), develops a systematic class of models called autoregressive integrated moving average (ARIMA) models to handle time-correlated modeling and forecasting. The approach includes a provision for treating more than one input series through multivariate ARIMA or through transfer function modeling. The defining feature of these models is that they are multiplicative models, meaning that the observed data are assumed to result from products of factors involving differential or difference equation operators responding to a white noise input.

A more recent approach to the same problem uses additive models more familiar to statisticians. In this approach, the observed data are assumed to result from sums of series, each with a specified time series structure; for example, in economics, assume a series is generated as the sum of trend, a seasonal effect, and error. The state-space model that results is then treated by making judicious use of the celebrated Kalman filters and smoothers, developed originally for estimation and control in space applications. Two relatively complete presentations from this point of view are in Harvey (1991) and Kitagawa and Gersch (1996). Time series regression is introduced in Chapter 2, and ARIMA and related time domain models are studied in Chapter 3, with the emphasis on classical, statistical, univariate linear regression. Special topics on time domain analysis are covered in Chapter 5; these topics include modern treatments of, for example, time series with long memory and GARCH models for the analysis of volatility. The state-space model, Kalman filtering and smoothing, and related topics are developed in Chapter 6.

Conversely, the frequency domain approach assumes the primary characteristics of interest in time series analyses relate to periodic or systematic sinusoidal variations found naturally in most data. These periodic variations are often caused by biological, physical, or environmental phenomena of interest. A series of periodic shocks may influence certain areas of the brain; wind may affect vibrations on an airplane wing; sea surface temperatures caused by El Niño oscillations may affect the number of fish in the ocean. The study of periodicity extends to economics and social sciences, where one may be interested in yearly periodicities in such series as monthly unemployment or monthly birth rates.

In spectral analysis, the partition of the various kinds of periodic variation in a time series is accomplished by evaluating separately the variance associated with each periodicity of interest. This variance profile over frequency is called the power spectrum. In our view, no schism divides time domain and frequency domain methodology, although cliques are often formed that advocate primarily one or the other of the approaches to analyzing data. In many cases, the two approaches may produce similar answers for long series, but the comparative performance over short samples is better done in the time domain. In some cases, the frequency domain formulation simply provides a convenient means for carrying out what is conceptually a time domain calculation. Hopefully, this book will demonstrate that the best path to analyzing many data sets is to use the two approaches in a complementary fashion. Expositions emphasizing primarily the frequency domain approach can be found in Bloomfield (1976, 2000), Priestley (1981), or Jenkins and Watts (1968). On a more advanced level, Hannan (1970), Brillinger (1981, 2001), Brockwell and Davis (1991), and Fuller (1996) are available as theoretical sources. Our coverage of the frequency domain is given in Chapters 4 and 7.

The objective of this book is to provide a unified and reasonably complete exposition of statistical methods used in time series analysis, giving serious consideration to both the time and frequency domain approaches. Because a myriad of possible methods for analyzing any particular experimental series can exist, we have integrated real data from a number of subject fields into the exposition and have suggested methods for analyzing these data.

1.2 The Nature of Time Series Data

Some of the problems and questions of interest to the prospective time series analyst can best be exposed by considering real experimental data taken from different subject areas. The following cases illustrate some of the common kinds of experimental time series data as well as some of the statistical questions that might be asked about such data.

Example 1.1 Johnson & Johnson Quarterly Earnings

Figure 1.1 shows quarterly earnings per share for the U.S. company Johnson & Johnson, furnished by Professor Paul Griffin (personal communication) of the Graduate School of Management, University of California, Davis. There are 84




Fig. 1.1. Johnson & Johnson quarterly earnings per share, 84 quarters, 1960-I to 1980-IV.

quarters (21 years) measured from the first quarter of 1960 to the last quarter of 1980. Modeling such series begins by observing the primary patterns in the time history. In this case, note the gradually increasing underlying trend and the rather regular variation superimposed on the trend that seems to repeat over quarters. Methods for analyzing data such as these are explored in [Chapter 2](#) (see [Problem 2.1](#)) using regression techniques and in [Chapter 6](#), §6.6, using structural equation modeling.

To plot the data using the R statistical package, type the following:¹

```
require(astsa)      # SEE THE FOOTNOTE
plot(jj, type="o", ylab="Quarterly Earnings per Share")
```

Example 1.2 Global Warming

Consider the global temperature series record shown in [Figure 1.2](#). The data are the global mean land–ocean temperature index from 1880 to 2009, with the base period 1951–1980. In particular, the data are deviations, measured in degrees centigrade, from the 1951–1980 average, and are an update of Hansen et al. (2006). We note an apparent upward trend in the series during the latter part of the twentieth century that has been used as an argument for the global warming hypothesis. Note also the leveling off at about 1935 and then another rather sharp upward trend at about 1970. The question of interest for global warming proponents and opponents is whether the overall trend is natural or whether it is caused by some human-induced interface. [Problem 2.8](#) examines 634 years of glacial sediment data that might be taken as a long-term temperature proxy. Such percentage changes in temperature do not seem to be unusual over a time period of 100 years. Again, the question of trend is of more interest than particular periodicities.

The R code for this example is similar to the code in [Example 1.1](#):

```
plot(gtemp, type="o", ylab="Global Temperature Deviations")
```

¹ Throughout the text, we assume that the R package for the book, [astsa](#), has been installed and loaded. See [§R.2](#) for further details.




Fig. 1.2. Yearly average global temperature deviations (1880–2009) in degrees centigrade.

Example 1.3 Speech Data

More involved questions develop in applications to the physical sciences. **Figure 1.3** shows a small .1 second (1000 point) sample of recorded speech for the phrase *aaa . . . hhh*, and we note the repetitive nature of the signal and the rather regular periodicities. One current problem of great interest is computer recognition of speech, which would require converting this particular signal into the recorded phrase *aaa . . . hhh*. Spectral analysis can be used in this context to produce a signature of this phrase that can be compared with signatures of various library syllables to look for a match. One can immediately notice the rather regular repetition of small wavelets. The separation between the packets is known as the pitch period and represents the response of the vocal tract filter to a periodic sequence of pulses stimulated by the opening and closing of the glottis. In R, you can reproduce **Figure 1.3** using `plot(speech)`.




Fig. 1.3. Speech recording of the syllable *aaa . . . hhh* sampled at 10,000 points per second with $n = 1020$ points.




Fig. 1.4. Returns of the NYSE. The data are daily value weighted market returns from February 2, 1984 to December 31, 1991 (2000 trading days). The crash of October 19, 1987 occurs at $t = 938$.

Example 1.4 New York Stock Exchange

As an example of financial time series data, [Figure 1.4](#) shows the daily returns (or percent change) of the New York Stock Exchange (NYSE) from February 2, 1984 to December 31, 1991. It is easy to spot the crash of October 19, 1987 in the figure. The data shown in [Figure 1.4](#) are typical of return data. The mean of the series appears to be stable with an average return of approximately zero, however, the volatility (or variability) of data changes over time. In fact, the data show volatility clustering; that is, highly volatile periods tend to be clustered together. A problem in the analysis of these type of financial data is to forecast the volatility of future returns. Models such as ARCH and GARCH models (Engle, 1982; Bollerslev, 1986) and stochastic volatility models (Harvey, Ruiz and Shephard, 1994) have been developed to handle these problems. We will discuss these models and the analysis of financial data in Chapters 5 and 6. The R code for this example is similar to the previous examples:

```
plot(nyse, ylab="NYSE Returns")
```

Example 1.5 El Niño and Fish Population

We may also be interested in analyzing several time series at once. [Figure 1.5](#) shows monthly values of an environmental series called the Southern Oscillation Index (SOI) and associated Recruitment (number of new fish) furnished by Dr. Roy Mendelssohn of the Pacific Environmental Fisheries Group (personal communication). Both series are for a period of 453 months ranging over the years 1950–1987. The SOI measures changes in air pressure, related to sea surface temperatures in the central Pacific Ocean. The central Pacific warms every three to seven years due to the El Niño effect, which has been blamed, in particular, for the 1997 floods




Fig. 1.5. Monthly SOI and Recruitment (estimated new fish), 1950-1987.

in the midwestern portions of the United States. Both series in [Figure 1.5](#) tend to exhibit repetitive behavior, with regularly repeating cycles that are easily visible. This periodic behavior is of interest because underlying processes of interest may be regular and the rate or frequency of oscillation characterizing the behavior of the underlying series would help to identify them. One can also remark that the cycles of the SOI are repeating at a faster rate than those of the Recruitment series. The Recruitment series also shows several kinds of oscillations, a faster frequency that seems to repeat about every 12 months and a slower frequency that seems to repeat about every 50 months. The study of the kinds of cycles and their strengths is the subject of Chapter 4. The two series also tend to be somewhat related; it is easy to imagine that somehow the fish population is dependent on the SOI. Perhaps even a lagged relation exists, with the SOI signaling changes in the fish population. This possibility suggests trying some version of regression analysis as a procedure for relating the two series. Transfer function modeling, as considered in Chapter 5, can be applied in this case to obtain a model relating Recruitment to its own past and the past values of the SOI.

The following R code will reproduce [Figure 1.5](#):

```
par(mfrow = c(2,1)) # set up the graphics
plot(soi, ylab="", xlab="", main="Southern Oscillation Index")
plot(rec, ylab="", xlab="", main="Recruitment")
```




Fig. 1.6. fMRI data from various locations in the cortex, thalamus, and cerebellum; $n = 128$ points, one observation taken every 2 seconds.

Example 1.6 fMRI Imaging

A fundamental problem in classical statistics occurs when we are given a collection of independent series or vectors of series, generated under varying experimental conditions or treatment configurations. Such a set of series is shown in [Figure 1.6](#), where we observe data collected from various locations in the brain via functional magnetic resonance imaging (fMRI). In this example, five subjects were given periodic brushing on the hand. The stimulus was applied for 32 seconds and then stopped for 32 seconds; thus, the signal period is 64 seconds. The sampling rate was one observation every 2 seconds for 256 seconds ($n = 128$). For this example, we averaged the results over two subjects (these were evoked responses, and all subjects were in phase). The series shown in [Figure 1.6](#) are consecutive measures of blood oxygenation-level dependent (BOLD) signal intensity, which measures areas of activation in the brain. Notice that the periodicities appear strongly in the motor cortex series and less strongly in the thalamus and cerebellum. The fact that one has series from different areas of the brain suggests testing whether the areas are responding differently to the brush stimulus. Analysis of variance techniques accomplish this in classical statistics, and we show in [Chapter 7](#) how these classical techniques extend to the time series case, leading to a spectral analysis of variance.

The following R commands were used to plot the data:

```
par(mfrow=c(2,1), mar=c(3,2,1,0)+.5, mgp=c(1.6,.6,0))
ts.plot(fmri1[,2:5], col=1:4, ylab="BOLD", xlab="", main="Cortex")
ts.plot(fmri1[,6:9], col=1:4, ylab="BOLD", xlab="", main="Thalamus &
  Cerebellum")
mtext("Time (1 pt = 2 sec)", side=1, line=2)
```




Fig. 1.7. Arrival phases from an earthquake (top) and explosion (bottom) at 40 points per second.

Example 1.7 Earthquakes and Explosions

As a final example, the series in [Figure 1.7](#) represent two phases or arrivals along the surface, denoted by P ($t = 1, \dots, 1024$) and S ($t = 1025, \dots, 2048$), at a seismic recording station. The recording instruments in Scandinavia are observing earthquakes and mining explosions with one of each shown in [Figure 1.7](#). The general problem of interest is in distinguishing or discriminating between waveforms generated by earthquakes and those generated by explosions. Features that may be important are the rough amplitude ratios of the first phase P to the second phase S, which tend to be smaller for earthquakes than for explosions. In the case of the two events in [Figure 1.7](#), the ratio of maximum amplitudes appears to be somewhat less than .5 for the earthquake and about 1 for the explosion. Otherwise, note a subtle difference exists in the periodic nature of the S phase for the earthquake. We can again think about spectral analysis of variance for testing the equality of the periodic components of earthquakes and explosions. We would also like to be able to classify future P and S components from events of unknown origin, leading to the time series discriminant analysis developed in [Chapter 7](#).

To plot the data as in this example, use the following commands in R:

```
par(mfrow=c(2,1))
plot(EQ5, main="Earthquake")
plot(EXP6, main="Explosion")
```

1.3 Time Series Statistical Models

The primary objective of time series analysis is to develop mathematical models that provide plausible descriptions for sample data, like that encountered in the previous section. In order to provide a statistical setting for describing the character of data that seemingly fluctuate in a random fashion over time, we assume a time series can be defined as a collection of random variables indexed according to the order they are obtained in time. For example, we may consider a time series as a sequence of random variables, $x_1, x_2, x_3, \dots$, where the random variable x_1 denotes the value taken by the series at the first time point, the variable x_2 denotes the value for the second time period, x_3 denotes the value for the third time period, and so on. In general, a collection of random variables, $\{x_t\}$, indexed by t is referred to as a stochastic process. In this text, t will typically be discrete and vary over the integers $t = 0, \pm 1, \pm 2, \dots$, or some subset of the integers. The observed values of a stochastic process are referred to as a realization of the stochastic process. Because it will be clear from the context of our discussions, we use the term time series whether we are referring generically to the process or to a particular realization and make no notational distinction between the two concepts.

It is conventional to display a sample time series graphically by plotting the values of the random variables on the vertical axis, or ordinate, with the time scale as the abscissa. It is usually convenient to connect the values at adjacent time periods to reconstruct visually some original hypothetical continuous time series that might have produced these values as a discrete sample. Many of the series discussed in the previous section, for example, could have been observed at any continuous point in time and are conceptually more properly treated as continuous time series. The approximation of these series by discrete time parameter series sampled at equally spaced points in time is simply an acknowledgment that sampled data will, for the most part, be discrete because of restrictions inherent in the method of collection. Furthermore, the analysis techniques are then feasible using computers, which are limited to digital computations. Theoretical developments also rest on the idea that a continuous parameter time series should be specified in terms of finite-dimensional distribution functions defined over a finite number of points in time. This is not to say that the selection of the sampling interval or rate is not an extremely important consideration. The appearance of data can be changed completely by adopting an insufficient sampling rate. We have all seen wagon wheels in movies appear to be turning backwards because of the insufficient number of frames sampled by the camera. This phenomenon leads to a distortion called aliasing (see §4.2).

The fundamental visual characteristic distinguishing the different series shown in [Example 1.1–Example 1.7](#) is their differing degrees of smoothness. One possible explanation for this smoothness is that it is being induced by the supposition that adjacent points in time are correlated, so the value of the series at time t , say, x_t , depends in some way on the past values $x_{t-1}, x_{t-2}, \dots$. This model expresses a fundamental way in which we might think about generating realistic-looking time series. To begin to develop an approach to using collections of random variables to model time series, consider [Example 1.8](#).




Fig. 1.8. Gaussian white noise series (top) and three-point moving average of the Gaussian white noise series (bottom).

Example 1.8 White Noise (3 flavors)

A simple kind of generated series might be a collection of uncorrelated random variables, w_t , with mean 0 and finite variance σ_w^2 . The time series generated from uncorrelated variables is used as a model for noise in engineering applications, where it is called *white noise*; we shall denote this process as $\mathbf{w}_t \sim \mathbf{wn}(\mathbf{0}, \sigma_w^2)$. The designation white originates from the analogy with white light and indicates that all possible periodic oscillations are present with equal strength.

We will sometimes require the noise to be independent and identically distributed (iid) random variables with mean 0 and variance σ_w^2 . We distinguish this by writing $\mathbf{w}_t \sim \mathbf{iid}(\mathbf{0}, \sigma_w^2)$ or by saying *white independent noise* or *iid noise*. A particularly useful white noise series is *Gaussian white noise*, wherein the w_t are independent normal random variables, with mean 0 and variance σ_w^2 ; or more succinctly, $\mathbf{w}_t \sim \mathbf{iid N}(\mathbf{0}, \sigma_w^2)$. Figure 1.8 shows in the upper panel a collection of 500 such random variables, with $\sigma_w^2 = 1$, plotted in the order in which they were drawn. The resulting series bears a slight resemblance to the explosion in Figure 1.7 but is not smooth enough to serve as a plausible model for any of the other experimental series. The plot tends to show visually a mixture of many different kinds of oscillations in the white noise series.

If the stochastic behavior of all time series could be explained in terms of the white noise model, classical statistical methods would suffice. Two ways of introducing serial correlation and more smoothness into time series models are given in (1.9) and (1.10).

Example 1.9 Moving Averages and Filtering

We might replace the white noise series w_t by a moving average that smooths the series. For example, consider replacing w_t in [Example 1.8](#) by an average of its current value and its immediate neighbors in the past and future. That is, let

$$v_t = \frac{1}{3}(w_{t-1} + w_t + w_{t+1}), \quad (1.1)$$

which leads to the series shown in the lower panel of [Figure 1.8](#). Inspecting the series shows a smoother version of the first series, reflecting the fact that the slower oscillations are more apparent and some of the faster oscillations are taken out. We begin to notice a similarity to the SOI in [Figure 1.5](#), or perhaps, to some of the fMRI series in [Figure 1.6](#).

To reproduce [Figure 1.8](#) in R use the following commands. A linear combination of values in a time series such as in (1.1) is referred to, generically, as a filtered series; hence the command `filter`.

```
w = rnorm(500,0,1) # 500 N(0,1) variates
v = filter(w, sides=2, filter=rep(1/3,3)) # moving average
par(mfrow=c(2,1))
plot.ts(w, main="white noise")
plot.ts(v, ylim=c(-3,3), main="moving average")
```

The speech series in [Figure 1.3](#) and the Recruitment series in [Figure 1.5](#), as well as some of the MRI series in [Figure 1.6](#), differ from the moving average series because one particular kind of oscillatory behavior seems to predominate, producing a sinusoidal type of behavior. A number of methods exist for generating series with this quasi-periodic behavior; we illustrate a popular one based on the autoregressive model considered in Chapter 3.

Example 1.10 Autoregressions

Suppose we consider the white noise series w_t of [Example 1.8](#) as input and calculate the output using the second-order equation

$$x_t = x_{t-1} - .9x_{t-2} + w_t \quad (1.2)$$

successively for $t = 1, 2, \dots, 500$. Equation (1.2) represents a regression or prediction of the current value x_t of a time series as a function of the past two values of the series, and, hence, the term autoregression is suggested for this model. A problem with startup values exists here because (1.2) also depends on the initial conditions x_0 and x_{-1} , but, for now, we assume that we are given these values and generate the succeeding values by substituting into (1.2). The resulting output series is shown in [Figure 1.9](#), and we note the periodic behavior of the series, which is similar to that displayed by the speech series in [Figure 1.3](#). The autoregressive model above and its generalizations can be used as an underlying model for many observed series and will be studied in detail in Chapter 3.

Similar to the previous example, the data are obtained by a filter of white noise. Thus we can simulate and plot data from the model (1.2) in R using the following commands (another way is to use `arima.sim`).




Fig. 1.9. Autoregressive series generated from model (1.2).

```
w = rnorm(550,0,1) # 50 extra to avoid startup problems
x = filter(w, filter=c(1,-.9), method="recursive")[-(1:50)]
plot.ts(x, main="autoregression")
```

Example 1.11 Random Walk with Drift

A model for analyzing trend such as seen in the global temperature data in Figure 1.2, is the random walk with drift model given by

$$x_t = \delta + x_{t-1} + w_t \quad (1.3)$$

for $t = 1, 2, \dots$, with initial condition $x_0 = 0$, and where w_t is white noise. The constant δ is called the drift, and when $\delta = 0$, (1.3) is called simply a random walk. The term random walk comes from the fact that, when $\delta = 0$, the value of the time series at time t is the value of the series at time $t - 1$ plus a completely random movement determined by w_t . Note that we may rewrite (1.3) as a cumulative sum of white noise variates. That is,

$$x_t = \delta t + \sum_{j=1}^t w_j \quad (1.4)$$

for $t = 1, 2, \dots$; either use induction, or plug (1.4) into (1.3) to verify this statement. Figure 1.10 shows 200 observations generated from the model with $\delta = 0$ and $.2$, and with $\sigma_w = 1$. For comparison, we also superimposed the straight line $.2t$ on the graph.

To reproduce Figure 1.10 in R use the following code (notice the use of multiple commands per line using a semicolon).

```
set.seed(154) # so you can reproduce the results
w = rnorm(200,0,1); x = cumsum(w) # two commands in one line
wd = w +.2; xd = cumsum(wd)
```




Fig. 1.10. Random walk, $\sigma_w = 1$, with drift $\delta = .2$ (upper jagged line), without drift, $\delta = 0$ (lower jagged blue line), and straight (dashed) lines with slope δ .

```
plot.ts(xd, ylim=c(-5,55), main="random walk", ylab='')
lines(x, col=4); abline(h=0, col=4, lty=2)
abline(a=0, b=.2, lty=2)
```

Example 1.12 Signal in Noise

Many realistic models for generating time series assume an underlying signal with some consistent periodic variation, contaminated by adding a random noise. For example, it is easy to detect the regular cycle fMRI series displayed on the top of [Figure 1.6](#). Consider the model

$$x_t = 2 \cos(2\pi \frac{t+15}{50}) + w_t \quad (1.5)$$

for $t = 1, 2, \dots, 500$, where the first term is regarded as the signal, shown in the upper panel of [Figure 1.11](#). We note that a sinusoidal waveform can be written as

$$A \cos(2\pi\omega t + \phi), \quad (1.6)$$

where A is the amplitude, ω is the frequency of oscillation, and ϕ is a phase shift. In (1.5), $A = 2$, $\omega = 1/50$ (one cycle every 50 time points), and $\phi = 2\pi 15/50 = .6\pi$.

An additive noise term was taken to be white noise with $\sigma_w = 1$ (middle panel) and $\sigma_w = 5$ (bottom panel), drawn from a normal distribution. Adding the two together obscures the signal, as shown in the lower panels of [Figure 1.11](#). Of course, the degree to which the signal is obscured depends on the amplitude of the signal and the size of σ_w . The ratio of the amplitude of the signal to σ_w (or some function of the ratio) is sometimes called the signal-to-noise ratio (SNR); the larger the SNR, the easier it is to detect the signal. Note that the signal is easily discernible in the middle panel of [Figure 1.11](#), whereas the signal is obscured in the bottom panel. Typically, we will not observe the signal but the signal obscured by noise.

To reproduce [Figure 1.11](#) in R, use the following commands:




Fig. 1.11. Cosine wave with period 50 points (top panel) compared with the cosine wave contaminated with additive white Gaussian noise, $\sigma_w = 1$ (middle panel) and $\sigma_w = 5$ (bottom panel); see (1.5).

```
cs = 2*cos(2*pi*1:500/50 + .6*pi); w = rnorm(500,0,1)
par(mfrow=c(3,1), mar=c(3,2,2,1), cex.main=1.5)
plot.ts(cs, main=expression(2*cos(2*pi*t/50+.6*pi)))
plot.ts(cs+w, main=expression(2*cos(2*pi*t/50+.6*pi) + N(0,1)))
plot.ts(cs+5*w, main=expression(2*cos(2*pi*t/50+.6*pi) + N(0,25)))
```

In Chapter 4, we will study the use of spectral analysis as a possible technique for detecting regular or periodic signals, such as the one described in [Example 1.12](#). In general, we would emphasize the importance of simple additive models such as given above in the form

$$x_t = s_t + v_t, \quad (1.7)$$

where s_t denotes some unknown signal and v_t denotes a time series that may be white or correlated over time. The problems of detecting a signal and then in estimating or extracting the waveform of s_t are of great interest in many areas of engineering and the physical and biological sciences. In economics, the underlying signal may be a trend or it may be a seasonal component of a series. Models such as (1.7), where the signal has an autoregressive structure, form the motivation for the state-space model of Chapter 6.

In the above examples, we have tried to motivate the use of various combinations of random variables emulating real time series data. Smoothness characteristics of observed time series were introduced by combining the random variables in various ways. Averaging independent random variables over adjacent time points, as in

Example 1.9, or looking at the output of difference equations that respond to white noise inputs, as in **Example 1.10**, are common ways of generating correlated data. In the next section, we introduce various theoretical measures used for describing how time series behave. As is usual in statistics, the complete description involves the multivariate distribution function of the jointly sampled values $x_1, x_2, \dots, x_n$, whereas more economical descriptions can be had in terms of the mean and autocorrelation functions. Because correlation is an essential feature of time series analysis, the most useful descriptive measures are those expressed in terms of covariance and correlation functions.

1.4 Measures of Dependence: Autocorrelation and Cross-Correlation

A complete description of a time series, observed as a collection of n random variables at arbitrary integer time points $t_1, t_2, \dots, t_n$, for any positive integer n , is provided by the joint distribution function, evaluated as the probability that the values of the series are jointly less than the n constants, $c_1, c_2, \dots, c_n$; i.e.,

$$F(c_1, c_2, \dots, c_n) = P(x_{t_1} \leq c_1, x_{t_2} \leq c_2, \dots, x_{t_n} \leq c_n). \quad (1.8)$$

Unfortunately, the multidimensional distribution function cannot usually be written easily unless the random variables are jointly normal, in which case the joint density has the well-known form displayed in (1.31).

Although the joint distribution function describes the data completely, it is an unwieldy tool for displaying and analyzing time series data. The distribution function (1.8) must be evaluated as a function of n arguments, so any plotting of the corresponding multivariate density functions is virtually impossible. The marginal distribution functions

$$F_t(x) = P\{x_t \leq x\}$$

or the corresponding marginal density functions

$$f_t(x) = \frac{\partial F_t(x)}{\partial x},$$

when they exist, are often informative for examining the marginal behavior of a series.² Another informative marginal descriptive measure is the mean function.

Definition 1.1 *The mean function is defined as*

$$\mu_{xt} = E(x_t) = \int_{-\infty}^{\infty} x f_t(x) dx, \quad (1.9)$$

provided it exists, where E denotes the usual expected value operator. When no confusion exists about which time series we are referring to, we will drop a subscript and write μ_{xt} as μ_t .

² If x_t is Gaussian with mean μ_t and variance σ_t^2 , abbreviated as $x_t \sim N(\mu_t, \sigma_t^2)$, the marginal density is given by $f_t(x) = \frac{1}{\sigma_t \sqrt{2\pi}} \exp \left\{ -\frac{1}{2\sigma_t^2} (x - \mu_t)^2 \right\}$.

Example 1.13 Mean Function of a Moving Average Series

If w_t denotes a white noise series, then $\mu_{wt} = E(w_t) = 0$ for all t . The top series in [Figure 1.8](#) reflects this, as the series clearly fluctuates around a mean value of zero. Smoothing the series as in [Example 1.9](#) does not change the mean because we can write

$$\mu_{vt} = E(v_t) = \frac{1}{3}[E(w_{t-1}) + E(w_t) + E(w_{t+1})] = 0.$$

Example 1.14 Mean Function of a Random Walk with Drift

Consider the random walk with drift model given in (1.4),

$$x_t = \delta t + \sum_{j=1}^t w_j, \quad t = 1, 2, \dots$$

Because $E(w_t) = 0$ for all t , and δ is a constant, we have

$$\mu_{xt} = E(x_t) = \delta t + \sum_{j=1}^t E(w_j) = \delta t$$

which is a straight line with slope δ . A realization of a random walk with drift can be compared to its mean function in [Figure 1.10](#).

Example 1.15 Mean Function of Signal Plus Noise

A great many practical applications depend on assuming the observed data have been generated by a fixed signal waveform superimposed on a zero-mean noise process, leading to an additive signal model of the form (1.5). It is clear, because the signal in (1.5) is a fixed function of time, we will have

$$\begin{aligned} \mu_{xt} &= E(x_t) = E[2 \cos(2\pi t/50 + .6\pi) + w_t] \\ &= 2 \cos(2\pi t/50 + .6\pi) + E(w_t) \\ &= 2 \cos(2\pi t/50 + .6\pi), \end{aligned}$$

and the mean function is just the cosine wave.

The lack of independence between two adjacent values x_s and x_t can be assessed numerically, as in classical statistics, using the notions of covariance and correlation. Assuming the variance of x_t is finite, we have the following definition.

Definition 1.2 *The autocovariance function is defined as the second moment product*

$$\gamma_x(s, t) = \text{cov}(x_s, x_t) = E[(x_s - \mu_s)(x_t - \mu_t)], \quad (1.10)$$

for all s and t . When no possible confusion exists about which time series we are referring to, we will drop the subscript and write $\gamma_x(s, t)$ as $\gamma(s, t)$. Note that $\gamma_x(s, t) = \gamma_x(t, s)$ for all time points s and t .

The autocovariance measures the *linear* dependence between two points on the same series observed at different times. Very smooth series exhibit autocovariance functions that stay large even when the t and s are far apart, whereas choppy series tend to have autocovariance functions that are nearly zero for large separations. Recall from classical statistics that if $\gamma_x(s, t) = 0$, x_s and x_t are not linearly related, but there still may be some dependence structure between them. If, however, x_s and x_t are bivariate normal, $\gamma_x(s, t) = 0$ ensures their independence. It is clear that, for $s = t$, the autocovariance reduces to the (assumed finite) variance, because

$$\gamma_x(t, t) = E[(x_t - \mu_t)^2] = \text{var}(x_t). \quad (1.11)$$

Example 1.16 Autocovariance of White Noise

The white noise series w_t has $E(w_t) = 0$ and

$$\gamma_w(s, t) = \text{cov}(w_s, w_t) = \begin{cases} \sigma_w^2 & s = t, \\ 0 & s \neq t. \end{cases} \quad (1.12)$$

A realization of white noise with $\sigma_w^2 = 1$ is shown in the top panel of [Figure 1.8](#).

Example 1.17 Autocovariance of a Moving Average

Consider applying a three-point moving average to the white noise series w_t of the previous example as in [Example 1.9](#). In this case,

$$\gamma_v(s, t) = \text{cov}(v_s, v_t) = \text{cov}\left\{\frac{1}{3}(w_{s-1} + w_s + w_{s+1}), \frac{1}{3}(w_{t-1} + w_t + w_{t+1})\right\}.$$

When $s = t$ we have³

$$\begin{aligned} \gamma_v(t, t) &= \frac{1}{9} \text{cov}\{(w_{t-1} + w_t + w_{t+1}), (w_{t-1} + w_t + w_{t+1})\} \\ &= \frac{1}{9} [\text{cov}(w_{t-1}, w_{t-1}) + \text{cov}(w_t, w_t) + \text{cov}(w_{t+1}, w_{t+1})] \\ &= \frac{3}{9} \sigma_w^2. \end{aligned}$$

When $s = t + 1$,

$$\begin{aligned} \gamma_v(t + 1, t) &= \frac{1}{9} \text{cov}\{(w_t + w_{t+1} + w_{t+2}), (w_{t-1} + w_t + w_{t+1})\} \\ &= \frac{1}{9} [\text{cov}(w_t, w_t) + \text{cov}(w_{t+1}, w_{t+1})] \\ &= \frac{2}{9} \sigma_w^2, \end{aligned}$$

using (1.12). Similar computations give $\gamma_v(t - 1, t) = 2\sigma_w^2/9$, $\gamma_v(t + 2, t) = \gamma_v(t - 2, t) = \sigma_w^2/9$, and 0 when $|t - s| > 2$. We summarize the values for all s and t as

$$\gamma_v(s, t) = \begin{cases} \frac{3}{9} \sigma_w^2 & s = t, \\ \frac{2}{9} \sigma_w^2 & |s - t| = 1, \\ \frac{1}{9} \sigma_w^2 & |s - t| = 2, \\ 0 & |s - t| > 2. \end{cases} \quad (1.13)$$

³ If the random variables $U = \sum_{j=1}^m a_j X_j$ and $V = \sum_{k=1}^r b_k Y_k$ are linear combinations of random variables $\{X_j\}$ and $\{Y_k\}$, respectively, then $\text{cov}(U, V) = \sum_{j=1}^m \sum_{k=1}^r a_j b_k \text{cov}(X_j, Y_k)$. Furthermore, $\text{var}(U) = \text{cov}(U, U)$.

Example 1.17 shows clearly that the smoothing operation introduces a covariance function that decreases as the separation between the two time points increases and disappears completely when the time points are separated by three or more time points. This particular autocovariance is interesting because it only depends on the time separation or lag and not on the absolute location of the points along the series. We shall see later that this dependence suggests a mathematical model for the concept of weak stationarity.

Example 1.18 Autocovariance of a Random Walk

For the random walk model, $x_t = \sum_{j=1}^t w_j$, we have

$$\gamma_x(s, t) = \text{cov}(x_s, x_t) = \text{cov}\left(\sum_{j=1}^s w_j, \sum_{k=1}^t w_k\right) = \min\{s, t\} \sigma_w^2,$$

because the w_t are uncorrelated random variables. Note that, as opposed to the previous examples, the autocovariance function of a random walk depends on the particular time values s and t , and not on the time separation or lag. Also, notice that the variance of the random walk, $\text{var}(x_t) = \gamma_x(t, t) = t \sigma_w^2$, increases without bound as time t increases. The effect of this variance increase can be seen in **Figure 1.10** where the processes start to move away from their mean functions δt (note that $\delta = 0$ and .2 in that example).

As in classical statistics, it is more convenient to deal with a measure of association between -1 and 1 , and this leads to the following definition.

Definition 1.3 *The autocorrelation function (ACF) is defined as*

$$\rho(s, t) = \frac{\gamma(s, t)}{\sqrt{\gamma(s, s)\gamma(t, t)}}. \quad (1.14)$$

The ACF measures the linear predictability of the series at time t , say x_t , using only the value x_s . We can show easily that $-1 \leq \rho(s, t) \leq 1$ using the Cauchy–Schwarz inequality.⁴ If we can predict x_t *perfectly* from x_s through a linear relationship, $x_t = \beta_0 + \beta_1 x_s$, then the correlation will be $+1$ when $\beta_1 > 0$, and -1 when $\beta_1 < 0$. Hence, we have a rough measure of the ability to forecast the series at time t from the value at time s .

Often, we would like to measure the predictability of another series y_t from the series x_s . Assuming both series have finite variances, we have the following definition.

Definition 1.4 *The cross-covariance function between two series, x_t and y_t , is*

$$\gamma_{xy}(s, t) = \text{cov}(x_s, y_t) = E[(x_s - \mu_{xs})(y_t - \mu_{yt})]. \quad (1.15)$$

There is also a scaled version of the cross-covariance function.

⁴ The Cauchy–Schwarz inequality implies $|\gamma(s, t)|^2 \leq \gamma(s, s)\gamma(t, t)$.

Definition 1.5 *The cross-correlation function (CCF) is given by*

$$\rho_{xy}(s, t) = \frac{\gamma_{xy}(s, t)}{\sqrt{\gamma_x(s, s)\gamma_y(t, t)}}. \quad (1.16)$$

We may easily extend the above ideas to the case of more than two series, say, $x_{t1}, x_{t2}, \dots, x_{tr}$; that is, multivariate time series with r components. For example, the extension of (1.10) in this case is

$$\gamma_{jk}(s, t) = E[(x_{sj} - \mu_{sj})(x_{tk} - \mu_{tk})] \quad j, k = 1, 2, \dots, r. \quad (1.17)$$

In the definitions above, the autocovariance and cross-covariance functions may change as one moves along the series because the values depend on both s and t , the locations of the points in time. In **Example 1.17**, the autocovariance function depends on the separation of x_s and x_t , say, $h = |s - t|$, and not on where the points are located in time. As long as the points are separated by h units, the location of the two points does not matter. This notion, called weak stationarity, when the mean is constant, is fundamental in allowing us to analyze sample time series data when only a single series is available.

1.5 Stationary Time Series

The preceding definitions of the mean and autocovariance functions are completely general. Although we have not made any special assumptions about the behavior of the time series, many of the preceding examples have hinted that a sort of regularity may exist over time in the behavior of a time series. We introduce the notion of regularity using a concept called stationarity.

Definition 1.6 *A strictly stationary time series is one for which the probabilistic behavior of every collection of values*

$$\{x_{t_1}, x_{t_2}, \dots, x_{t_k}\}$$

is identical to that of the time shifted set

$$\{x_{t_1+h}, x_{t_2+h}, \dots, x_{t_k+h}\}.$$

That is,

$$P\{x_{t_1} \leq c_1, \dots, x_{t_k} \leq c_k\} = P\{x_{t_1+h} \leq c_1, \dots, x_{t_k+h} \leq c_k\} \quad (1.18)$$

for all $k = 1, 2, \dots$, all time points $t_1, t_2, \dots, t_k$, all numbers $c_1, c_2, \dots, c_k$, and all time shifts $h = 0, \pm 1, \pm 2, \dots$.

If a time series is strictly stationary, then all of the multivariate distribution functions for subsets of variables must agree with their counterparts in the shifted set for all values of the shift parameter h . For example, when $k = 1$, (1.18) implies that

$$P\{x_s \leq c\} = P\{x_t \leq c\} \quad (1.19)$$

for any time points s and t . This statement implies, for example, that the probability that the value of a time series sampled hourly is negative at 1 AM is the same as at 10 AM. In addition, if the mean function, μ_t , of the series x_t exists, (1.19) implies that $\mu_s = \mu_t$ for all s and t , and hence μ_t must be constant. Note, for example, that a random walk process with drift is *not* strictly stationary because its mean function changes with time; see [Example 1.14](#).

When $k = 2$, we can write (1.18) as

$$P\{x_s \leq c_1, x_t \leq c_2\} = P\{x_{s+h} \leq c_1, x_{t+h} \leq c_2\} \quad (1.20)$$

for any time points s and t and shift h . Thus, if the variance function of the process exists, (1.20) implies that the autocovariance function of the series x_t satisfies

$$\gamma(s, t) = \gamma(s + h, t + h)$$

for all s and t and h . We may interpret this result by saying the autocovariance function of the process depends only on the time difference between s and t , and not on the actual times.

The version of stationarity in [Definition 1.6](#) is too strong for most applications. Moreover, it is difficult to assess strict stationarity from a single data set. Rather than imposing conditions on all possible distributions of a time series, we will use a milder version that imposes conditions only on the first two moments of the series. We now have the following definition.

Definition 1.7 *A **weakly stationary** time series, x_t , is a finite variance process such that*

- (i) *the mean value function, μ_t , defined in (1.9) is constant and does not depend on time t , and*
- (ii) *the autocovariance function, $\gamma(s, t)$, defined in (1.10) depends on s and t only through their difference $|s - t|$.*

*Henceforth, we will use the term **stationary** to mean weakly stationary; if a process is stationary in the strict sense, we will use the term **strictly stationary**.*

It should be clear from the discussion of strict stationarity following [Definition 1.6](#) that a strictly stationary, finite variance, time series is also stationary. The converse is not true unless there are further conditions. One important case where stationarity implies strict stationarity is if the time series is Gaussian [meaning all finite distributions, (1.18), of the series are Gaussian]. We will make this concept more precise at the end of this section.

Because the mean function, $E(x_t) = \mu_t$, of a stationary time series is independent of time t , we will write

$$\mu_t = \mu. \quad (1.21)$$

Also, because the autocovariance function, $\gamma(s, t)$, of a stationary time series, x_t , depends on s and t only through their difference $|s - t|$, we may simplify the notation. Let $s = t + h$, where h represents the time shift or lag. Then

$$\gamma(t + h, t) = \text{cov}(x_{t+h}, x_t) = \text{cov}(x_h, x_0) = \gamma(h, 0)$$

because the time difference between times $t + h$ and t is the same as the time difference between times h and 0. Thus, the autocovariance function of a stationary time series does not depend on the time argument t . Henceforth, for convenience, we will drop the second argument of $\gamma(h, 0)$.

Definition 1.8 *The autocovariance function of a stationary time series will be written as*

$$\gamma(h) = \text{cov}(x_{t+h}, x_t) = E[(x_{t+h} - \mu)(x_t - \mu)]. \quad (1.22)$$

Definition 1.9 *The autocorrelation function (ACF) of a stationary time series will be written using (1.14) as*

$$\rho(h) = \frac{\gamma(t + h, t)}{\sqrt{\gamma(t + h, t + h)\gamma(t, t)}} = \frac{\gamma(h)}{\gamma(0)}. \quad (1.23)$$

The Cauchy–Schwarz inequality shows again that $-1 \leq \rho(h) \leq 1$ for all h , enabling one to assess the relative importance of a given autocorrelation value by comparing with the extreme values -1 and 1 .

Example 1.19 Stationarity of White Noise

The mean and autocovariance functions of the white noise series discussed in [Example 1.8](#) and [Example 1.16](#) are easily evaluated as $\mu_{wt} = 0$ and

$$\gamma_w(h) = \text{cov}(w_{t+h}, w_t) = \begin{cases} \sigma_w^2 & h = 0, \\ 0 & h \neq 0. \end{cases}$$

Thus, white noise satisfies the conditions of [Definition 1.7](#) and is weakly stationary or stationary. If the white noise variates are also normally distributed or Gaussian, the series is also strictly stationary, as can be seen by evaluating (1.18) using the fact that the noise would also be iid. The autocorrelation function is given by $\rho_w(0) = 1$ and $\rho(h) = 0$ for $h \neq 0$.

Example 1.20 Trend Stationarity

If $x_t = \alpha + \beta t + y_t$, where y_t is (mean-zero) stationary, then the mean function is $\mu_{x,t} = E(x_t) = \alpha + \beta t$, which is not independent of time. Therefore, the process is not stationary. The autocovariance function, however, is independent of time, because $\gamma_x(h) = \text{cov}(x_{t+h}, x_t) = E[(x_{t+h} - \mu_{x,t+h})(x_t - \mu_{x,t})] = E(y_{t+h}y_t) = \gamma_y(h)$. Thus, the model may be considered as having stationary behavior around a linear trend; this behavior is sometimes called *trend stationarity*.




Fig. 1.12. Autocorrelation function of a three-point moving average.

Example 1.21 Stationarity of a Moving Average

The three-point moving average process of [Example 1.9](#) is stationary because, from [Example 1.13](#) and [Example 1.17](#), the mean and autocovariance functions $\mu_{v,t} = 0$, and

$$\gamma_v(h) = \begin{cases} \frac{3}{9}\sigma_w^2 & h = 0, \\ \frac{2}{9}\sigma_w^2 & h = \pm 1, \\ \frac{1}{9}\sigma_w^2 & h = \pm 2, \\ 0 & |h| > 2 \end{cases}$$

are independent of time t , satisfying the conditions of [Definition 1.7](#).

The autocorrelation function is given by

$$\rho_v(h) = \begin{cases} 1 & h = 0, \\ \frac{2}{3} & h = \pm 1, \\ \frac{1}{3} & h = \pm 2, \\ 0 & |h| > 2. \end{cases}$$

[Figure 1.12](#) shows a plot of the autocorrelations as a function of lag h . Note that the ACF is symmetric about lag zero.

The autocovariance function of a stationary process has several useful properties (also, see [Problem 1.25](#)). First, the value at $h = 0$, namely

$$\gamma(0) = E[(x_t - \mu)^2] \quad (1.24)$$

is the variance of the time series; note that the Cauchy–Schwarz inequality implies

$$|\gamma(h)| \leq \gamma(0).$$

A final useful property, noted in the previous example, is that the autocovariance function of a stationary series is symmetric around the origin; that is,

$$\gamma(h) = \gamma(-h) \quad (1.25)$$

for all h . This property follows because

$$\gamma((t+h) - t) = \text{cov}(x_{t+h}, x_t) = \text{cov}(x_t, x_{t+h}) = \gamma(t - (t+h)),$$

which shows how to use the notation as well as proving the result.

When several series are available, a notion of stationarity still applies with additional conditions.

Definition 1.10 Two time series, say, x_t and y_t , are said to be **jointly stationary** if they are each stationary, and the cross-covariance function

$$\gamma_{xy}(h) = \text{cov}(x_{t+h}, y_t) = E[(x_{t+h} - \mu_x)(y_t - \mu_y)] \quad (1.26)$$

is a function only of lag h .

Definition 1.11 The **cross-correlation function (CCF)** of jointly stationary time series x_t and y_t is defined as

$$\rho_{xy}(h) = \frac{\gamma_{xy}(h)}{\sqrt{\gamma_x(0)\gamma_y(0)}}. \quad (1.27)$$

Again, we have the result $-1 \leq \rho_{xy}(h) \leq 1$ which enables comparison with the extreme values -1 and 1 when looking at the relation between x_{t+h} and y_t . The cross-correlation function is not generally symmetric about zero, i.e., typically $\rho_{xy}(h) \neq \rho_{xy}(-h)$. This is an important concept; it should be clear that $\text{cov}(x_2, y_1)$ and $\text{cov}(x_1, y_2)$ need not be the same. It is the case, however, that

$$\rho_{xy}(h) = \rho_{yx}(-h), \quad (1.28)$$

which can be shown by manipulations similar to those used to show (1.25).

Example 1.22 Joint Stationarity

Consider the two series, x_t and y_t , formed from the sum and difference of two successive values of a white noise process, say,

$$x_t = w_t + w_{t-1}$$

and

$$y_t = w_t - w_{t-1},$$

where w_t are independent random variables with zero means and variance σ_w^2 . It is easy to show that $\gamma_x(0) = \gamma_y(0) = 2\sigma_w^2$ and $\gamma_x(1) = \gamma_x(-1) = \sigma_w^2$, $\gamma_y(1) = \gamma_y(-1) = -\sigma_w^2$. Also,

$$\gamma_{xy}(1) = \text{cov}(x_{t+1}, y_t) = \text{cov}(w_{t+1} + w_t, w_t - w_{t-1}) = \sigma_w^2$$

because only one term is nonzero. Similarly, $\gamma_{xy}(0) = 0$, $\gamma_{xy}(-1) = -\sigma_w^2$. We obtain, using (1.11),

$$\rho_{xy}(h) = \begin{cases} 0 & h = 0, \\ 1/2 & h = 1, \\ -1/2 & h = -1, \\ 0 & |h| \geq 2. \end{cases}$$

Clearly, the autocovariance and cross-covariance functions depend only on the lag separation, h , so the series are jointly stationary.




Fig. 1.13. Demonstration of the results of [Example 1.23](#) when $\ell = 5$.

Example 1.23 Prediction Using Cross-Correlation

As a simple example of cross-correlation, consider the problem of determining possible leading or lagging relations between two series x_t and y_t . If the model

$$y_t = Ax_{t-\ell} + w_t$$

holds, the series x_t is said to lead y_t for $\ell > 0$ and is said to lag y_t for $\ell < 0$. Hence, the analysis of leading and lagging relations might be important in predicting the value of y_t from x_t . Assuming that the noise w_t is uncorrelated with the x_t series, the cross-covariance function can be computed as

$$\begin{aligned}\gamma_{xy}(h) &= \text{cov}(x_{t+h}, y_t) = \text{cov}(x_{t+h}, Ax_{t-\ell} + w_t) \\ &= \text{cov}(x_{t+h}, Ax_{t-\ell}) = A\gamma_x(h + \ell).\end{aligned}$$

Since (Cauchy–Schwarz) the largest absolute value of $\gamma_x(h + \ell)$ is $\gamma_x(0)$, i.e., when $h = -\ell$, the cross-covariance function will look like the autocovariance of the input series x_t , and it will have a peak on the negative side if x_t leads y_t and a peak on the positive side if x_t lags y_t . Below is the R code of an example with $\ell = 5$ and $\hat{\gamma}_{xy}(h)$ is shown in [Figure 1.13](#).

```
set.seed(90210)
x = rnorm(100)
y = lag(x, -5) + rnorm(100)
ccf(x, y, ylab='CCovF', type='covariance')
```

The concept of weak stationarity forms the basis for much of the analysis performed with time series. The fundamental properties of the mean and autocovariance functions (1.21) and (1.22) are satisfied by many theoretical models that appear to generate plausible sample realizations. In [Example 1.9](#) and [Example 1.10](#), two series were generated that produced stationary looking realizations, and in [Example 1.21](#), we showed that the series in [Example 1.9](#) was, in fact, weakly stationary. Both examples are special cases of the so-called linear process.

Definition 1.12 A **linear process**, x_t , is defined to be a linear combination of white noise variates w_t , and is given by

$$x_t = \mu + \sum_{j=-\infty}^{\infty} \psi_j w_{t-j}, \quad \sum_{j=-\infty}^{\infty} |\psi_j| < \infty. \quad (1.29)$$

For the linear process (see [Problem 1.11](#)), we may show that the autocovariance function is given by

$$\gamma(h) = \sigma_w^2 \sum_{j=-\infty}^{\infty} \psi_{j+h} \psi_j \quad (1.30)$$

for $h \geq 0$; recall that $\gamma(-h) = \gamma(h)$. This method exhibits the autocovariance function of the process in terms of the lagged products of the coefficients. We only need $\sum_{j=-\infty}^{\infty} \psi_j^2 < \infty$ for the process to be stationary, but we will discuss this further in Chapter 5. Note that, for [Example 1.9](#), we have $\psi_0 = \psi_{-1} = \psi_1 = 1/3$ and the result in [Example 1.21](#) comes out immediately. The autoregressive series in [Example 1.10](#) can also be put in this form, as can the general autoregressive moving average processes considered in Chapter 3.

Notice that the linear process (1.29) is dependent on the future ($j < 0$), the present ($j = 0$), and the past ($j > 0$). For the purpose of forecasting, a future dependent model will be useless. Consequently, we will focus on processes that do not depend on the future. Such models are called *causal*, and a causal linear process has $\psi_j = 0$ for $j < 0$; we will discuss this further in Chapter 3.

Finally, as previously mentioned, an important case in which a weakly stationary series is also strictly stationary is the normal or Gaussian series.

Definition 1.13 A process, $\{x_t\}$, is said to be a **Gaussian process** if the n -dimensional vectors $\mathbf{x} = (x_{t_1}, x_{t_2}, \dots, x_{t_n})'$, for every collection of time points $t_1, t_2, \dots, t_n$, and every positive integer n , have a multivariate normal distribution.

Defining the $n \times 1$ mean vector $E(\mathbf{x}) \equiv \boldsymbol{\mu} = (\mu_{t_1}, \mu_{t_2}, \dots, \mu_{t_n})'$ and the $n \times n$ covariance matrix as $\text{var}(\mathbf{x}) \equiv \Gamma = \{\gamma(t_i, t_j); i, j = 1, \dots, n\}$, which is assumed to be positive definite, the multivariate normal density function can be written as

$$f(\mathbf{x}) = (2\pi)^{-n/2} |\Gamma|^{-1/2} \exp \left\{ -\frac{1}{2} (\mathbf{x} - \boldsymbol{\mu})' \Gamma^{-1} (\mathbf{x} - \boldsymbol{\mu}) \right\}, \quad (1.31)$$

where $|\cdot|$ denotes the determinant.

We list some important items regarding linear and Gaussian processes.

- If a Gaussian time series, $\{x_t\}$, is weakly stationary, then $\mu_t = \mu$ and $\gamma(t_i, t_j) = \gamma(|t_i - t_j|)$, so that the vector $\boldsymbol{\mu}$ and the matrix Γ are independent of time. These facts imply that all the finite distributions, (1.31), of the series $\{x_t\}$ depend only on time lag and not on the actual times, and hence the series must be strictly stationary.
- A result called the *Wold Decomposition* ([Theorem B.5](#)) states that a stationary non-deterministic time series is a causal linear process (but with $\sum \psi_j^2 < \infty$). A linear process need not be Gaussian, but if a time series is Gaussian, then it is a causal linear process with $w_t \sim \text{iid } N(0, \sigma_w^2)$. Hence, stationary Gaussian processes form the basis of modeling many time series.
- It is not enough for the marginal distributions to be Gaussian for the process to be Gaussian. It is easy to construct a situation where X and Y are normal, but (X, Y) is not bivariate normal; e.g., let X and Z be independent normals and let $Y = Z$ if $XZ > 0$ and $Y = -Z$ if $XZ \leq 0$.

1.6 Estimation of Correlation

Although the theoretical autocorrelation and cross-correlation functions are useful for describing the properties of certain hypothesized models, most of the analyses must be performed using sampled data. This limitation means the sampled points $x_1, x_2, \dots, x_n$ only are available for estimating the mean, autocovariance, and autocorrelation functions. From the point of view of classical statistics, this poses a problem because we will typically not have iid copies of x_t that are available for estimating the covariance and correlation functions. In the usual situation with only one realization, however, the assumption of stationarity becomes critical. Somehow, we must use averages over this single realization to estimate the population means and covariance functions.

Accordingly, if a time series is stationary, the mean function (1.21) $\mu_t = \mu$ is constant so that we can estimate it by the sample mean,

$$\bar{x} = \frac{1}{n} \sum_{t=1}^n x_t. \quad (1.32)$$

The standard error of the estimate is the square root of $\text{var}(\bar{x})$, which can be computed using first principles (recall footnote 3), and is given by

$$\begin{aligned} \text{var}(\bar{x}) &= \text{var}\left(\frac{1}{n} \sum_{t=1}^n x_t\right) = \frac{1}{n^2} \text{cov}\left(\sum_{t=1}^n x_t, \sum_{s=1}^n x_s\right) \\ &= \frac{1}{n^2} \left(n\gamma_x(0) + (n-1)\gamma_x(1) + (n-2)\gamma_x(2) + \dots + \gamma_x(n-1) \right. \\ &\quad \left. + (n-1)\gamma_x(-1) + (n-2)\gamma_x(-2) + \dots + \gamma_x(1-n) \right) \\ &= \frac{1}{n} \sum_{h=-n}^n \left(1 - \frac{|h|}{n}\right) \gamma_x(h). \end{aligned} \quad (1.33)$$

If the process is white noise, (1.33) reduces to the familiar σ_x^2/n recalling that $\gamma_x(0) = \sigma_x^2$. Note that, in the case of dependence, the standard error of $\bar{x}$ may be smaller or larger than the white noise case depending on the nature of the correlation structure (see Problem 1.19)

The theoretical autocovariance function, (1.22), is estimated by the sample autocovariance function defined as follows.

Definition 1.14 *The sample autocovariance function is defined as*

$$\hat{\gamma}(h) = n^{-1} \sum_{t=1}^{n-h} (x_{t+h} - \bar{x})(x_t - \bar{x}), \quad (1.34)$$

with $\hat{\gamma}(-h) = \hat{\gamma}(h)$ for $h = 0, 1, \dots, n-1$.

The sum in (1.34) runs over a restricted range because x_{t+h} is not available for $t + h > n$. The estimator in (1.34) is preferred to the one that would be obtained by dividing by $n - h$ because (1.34) is a non-negative definite function.

The autocovariance function, $\gamma(h)$, of a stationary process is non-negative definite (see Problem 1.25) ensuring that variances of linear combinations of the variates x_t will never be negative. That is, for any $n \geq 1$, and constants $a_1, \dots, a_n$,

$$0 \leq \text{var}(a_1 x_1 + \dots + a_n x_n) = \sum_{j=1}^n \sum_{k=1}^n a_j a_k \gamma(j - k),$$

using footnote 3. And, because a variance is never negative, the estimate of that variance

$$\widehat{\text{var}}(a_1 x_1 + \dots + a_n x_n) = \sum_{j=1}^n \sum_{k=1}^n a_j a_k \widehat{\gamma}(j - k),$$

should also be non-negative. The estimator in (1.34) guarantees this result, but no such guarantee exists if we divide by $n - h$. Note that neither dividing by n nor $n - h$ in (1.34) yields an unbiased estimator of $\gamma(h)$.

Definition 1.15 *The sample autocorrelation function is defined, analogously to (1.23), as*

$$\widehat{\rho}(h) = \frac{\widehat{\gamma}(h)}{\widehat{\gamma}(0)}. \quad (1.35)$$

The sample autocorrelation function has a sampling distribution that allows us to assess whether the data comes from a completely random or white series or whether correlations are statistically significant at some lags.

Property 1.1 Large-Sample Distribution of the ACF

Under general conditions,⁵ if x_t is white noise, then for n large, the sample ACF, $\widehat{\rho}_x(h)$, for $h = 1, 2, \dots, H$, where H is fixed but arbitrary, is approximately normally distributed with zero mean and standard deviation given by

$$\sigma_{\widehat{\rho}_x(h)} = \frac{1}{\sqrt{n}}. \quad (1.36)$$

Based on the previous result, we obtain a rough method of assessing whether peaks in $\widehat{\rho}(h)$ are significant by determining whether the observed peak is outside the interval $\pm 2/\sqrt{n}$ (or plus/minus two standard errors); for a white noise sequence, approximately 95% of the sample ACFs should be within these limits. The applications of this property develop because many statistical modeling procedures depend on reducing a time series to a white noise series using various kinds of transformations. After such a procedure is applied, the plotted ACFs of the residuals should then lie roughly within the limits given above.

⁵ The general conditions are that x_t is iid with finite fourth moment. A sufficient condition for this to hold is that x_t is white Gaussian noise. Precise details are given in Theorem A.7 in Appendix A.

Example 1.24 A Simulated Time Series

To compare the sample ACF for various sample sizes to the theoretical ACF, consider a contrived set of data generated by tossing a fair coin, letting $x_t = 1$ when a head is obtained and $x_t = -1$ when a tail is obtained. Then, construct y_t as

$$y_t = 5 + x_t - .7x_{t-1}. \quad (1.37)$$

To simulate data, we consider two cases, one with a small sample size ($n = 10$) and another with a moderate sample size ($n = 100$).

```
set.seed(101010)
x1 = 2*rbinom(11, 1, .5) - 1    # simulated sequence of coin tosses
x2 = 2*rbinom(101, 1, .5) - 1
y1 = 5 + filter(x1, sides=1, filter=c(1,-.7))[-1]
y2 = 5 + filter(x2, sides=1, filter=c(1,-.7))[-1]
plot.ts(y1, type='s')          # not shown
plot.ts(y2, type='s')          # not shown
acf(y1, lag.max=4, plot=FALSE) # 1/√10 = .32
  Autocorrelations of series 'y1', by lag
      0      1      2      3      4
  1.000 -0.688  0.425 -0.306 -0.007
acf(y2, lag.max=4, plot=FALSE) # 1/√100 = .1
  Autocorrelations of series 'y2', by lag
      0      1      2      3      4
  1.000 -0.480 -0.002 -0.004  0.000
# Note that the sample ACF at lag zero is always 1 (Why?).
```

The theoretical ACF can be obtained from the model (1.37) using the fact that the mean of x_t is zero and the variance of x_t is one. It can be shown that

$$\rho_y(1) = \frac{-.7}{1 + .7^2} = -.47$$

and $\rho_y(h) = 0$ for $|h| > 1$ (Problem 1.24). It is interesting to compare the theoretical ACF with sample ACFs for the realization where $n = 10$ and the other realization where $n = 100$; note the increased variability in the smaller size sample.

Example 1.25 ACF of a Speech Signal

Computing the sample ACF as in the previous example can be thought of as matching the time series h units in the future, say, x_{t+h} against itself, x_t . Figure 1.14 shows the ACF of the speech series of Figure 1.3. The original series appears to contain a sequence of repeating short signals. The ACF confirms this behavior, showing repeating peaks spaced at about 106-109 points. Autocorrelation functions of the short signals appear, spaced at the intervals mentioned above. The distance between the repeating signals is known as the pitch period and is a fundamental parameter of interest in systems that encode and decipher speech. Because the series is sampled at 10,000 points per second, the pitch period appears to be between .0106 and .0109 seconds. To compute the sample ACF in R, use

```
acf(speech, 250)
```




Fig. 1.14. ACF of the speech series.

Definition 1.16 The estimators for the cross-covariance function, $\gamma_{xy}(h)$, as given in (1.26) and the cross-correlation, $\rho_{xy}(h)$, in (1.11) are given, respectively, by the **sample cross-covariance function**

$$\hat{\gamma}_{xy}(h) = n^{-1} \sum_{t=1}^{n-h} (x_{t+h} - \bar{x})(y_t - \bar{y}), \quad (1.38)$$

where $\hat{\gamma}_{xy}(-h) = \hat{\gamma}_{yx}(h)$ determines the function for negative lags, and the **sample cross-correlation function**

$$\hat{\rho}_{xy}(h) = \frac{\hat{\gamma}_{xy}(h)}{\sqrt{\hat{\gamma}_x(0)\hat{\gamma}_y(0)}}. \quad (1.39)$$

The sample cross-correlation function can be examined graphically as a function of lag h to search for leading or lagging relations in the data using the property mentioned in **Example 1.23** for the theoretical cross-covariance function. Because $-1 \leq \hat{\rho}_{xy}(h) \leq 1$, the practical importance of peaks can be assessed by comparing their magnitudes with their theoretical maximum values. Furthermore, for x_t and y_t independent linear processes of the form (1.29), we have the following property.

Property 1.2 Large-Sample Distribution of Cross-Correlation

The large sample distribution of $\hat{\rho}_{xy}(h)$ is normal with mean zero and

$$\sigma_{\hat{\rho}_{xy}} = \frac{1}{\sqrt{n}} \quad (1.40)$$

if at least one of the processes is independent white noise (see **Theorem A.8**).




Fig. 1.15. Sample ACFs of the SOI series (top) and of the Recruitment series (middle), and the sample CCF of the two series (bottom); negative lags indicate SOI leads Recruitment. The lag axes are in terms of seasons (12 months).

Example 1.26 SOI and Recruitment Correlation Analysis

The autocorrelation and cross-correlation functions are also useful for analyzing the joint behavior of two stationary series whose behavior may be related in some unspecified way. In [Example 1.5](#) (see [Figure 1.5](#)), we have considered simultaneous monthly readings of the SOI and the number of new fish (Recruitment) computed from a model. [Figure 1.15](#) shows the autocorrelation and cross-correlation functions (ACFs and CCF) for these two series. Both of the ACFs exhibit periodicities corresponding to the correlation between values separated by 12 units. Observations 12 months or one year apart are strongly positively correlated, as are observations at multiples such as 24, 36, 48, . . . Observations separated by six months are negatively correlated, showing that positive excursions tend to be associated with negative excursions six months removed.

The sample CCF in [Figure 1.15](#), however, shows some departure from the cyclic component of each series and there is an obvious peak at $h = -6$. This result implies that SOI measured at time $t - 6$ months is associated with the Recruitment series at time t . We could say the SOI leads the Recruitment series by six months. The sign of the ACF is negative, leading to the conclusion that the two series move in different directions; that is, increases in SOI lead to decreases in Recruitment and vice versa. We will discover in Chapter 2 that there is a relationship between the series, but the relationship is nonlinear. The dashed lines shown on the plots indicate $\pm 2/\sqrt{453}$ [see (1.40)], but since neither series is noise, these lines do not apply.




Fig. 1.16. Display for Example 1.27

To reproduce Figure 1.15 in R, use the following commands:

```
par(mfrow=c(3,1))
acf(soi, 48, main="Southern Oscillation Index")
acf(rec, 48, main="Recruitment")
ccf(soi, rec, 48, main="SOI vs Recruitment", ylab="CCF")
```

Example 1.27 Prewhitening and Cross Correlation Analysis

Although we do not have all the tools necessary yet, it is worthwhile to discuss the idea of prewhitening a series prior to a cross-correlation analysis. The basic idea is simple; in order to use Property 1.2, at least one of the series must be white noise. If this is not the case, there is no simple way to tell if a cross-correlation estimate is significantly different from zero. Hence, in Example 1.26, we were only guessing at the linear dependence relationship between SOI and Recruitment.

For example, in Figure 1.16 we generated two series, x_t and y_t , for $t = 1, \dots, 120$ independently as

$$x_t = 2 \cos(2\pi t \frac{1}{12}) + w_{t1} \quad \text{and} \quad y_t = 2 \cos(2\pi [t + 5] \frac{1}{12}) + w_{t2}$$

where $\{w_{t1}, w_{t2}; t = 1, \dots, 120\}$ are all independent standard normals. The series are made to resemble SOI and Recruitment. The generated data are shown in the

top row of the figure. The middle row of **Figure 1.16** shows the sample ACF of each series, each of which exhibits the cyclic nature of each series. The bottom row (left) of **Figure 1.16** shows the sample CCF between x_t and y_t , which appears to show cross-correlation even though the series are independent. The bottom row (right) also displays the sample CCF between x_t and the prewhitened y_t , which shows that the two sequences are uncorrelated. By prewhitening y_t , we mean that the signal has been removed from the data by running a regression of y_t on $\cos(2\pi t)$ and $\sin(2\pi t)$ [see **Example 2.9**] and then putting $\tilde{y}_t = y_t - \hat{y}_t$, where $\hat{y}_t$ are the predicted values from the regression.

The following code will reproduce **Figure 1.16**.

```
set.seed(1492)
num=120; t=1:num
X = ts(2*cos(2*pi*t/12) + rnorm(num), freq=12)
Y = ts(2*cos(2*pi*(t+5)/12) + rnorm(num), freq=12)
Yw = resid( lm(Y~ cos(2*pi*t/12) + sin(2*pi*t/12), na.action=NULL) )
par(mfrow=c(3,2), mgp=c(1.6,.6,0), mar=c(3,3,1,1) )
plot(X)
plot(Y)
acf(X,48, ylab='ACF(X)')
acf(Y,48, ylab='ACF(Y)')
ccf(X,Y,24, ylab='CCF(X,Y)')
ccf(X,Yw,24, ylab='CCF(X,Yw)', ylim=c(-.6,.6))
```

1.7 Vector-Valued and Multidimensional Series

We frequently encounter situations in which the relationships between a number of jointly measured time series are of interest. For example, in the previous sections, we considered discovering the relationships between the SOI and Recruitment series. Hence, it will be useful to consider the notion of a vector time series $\mathbf{x}_t = (x_{t1}, x_{t2}, \dots, x_{tp})'$, which contains as its components p univariate time series. We denote the $p \times 1$ column vector of the observed series as $\mathbf{x}_t$. The row vector $\mathbf{x}_t'$ is its transpose. For the stationary case, the $p \times 1$ mean vector

$$\boldsymbol{\mu} = E(\mathbf{x}_t) \quad (1.41)$$

of the form $\boldsymbol{\mu} = (\mu_{t1}, \mu_{t2}, \dots, \mu_{tp})'$ and the $p \times p$ autocovariance matrix

$$\Gamma(h) = E[(\mathbf{x}_{t+h} - \boldsymbol{\mu})(\mathbf{x}_t - \boldsymbol{\mu})'] \quad (1.42)$$

can be defined, where the elements of the matrix $\Gamma(h)$ are the cross-covariance functions

$$\gamma_{ij}(h) = E[(x_{t+h,i} - \mu_i)(x_{tj} - \mu_j)] \quad (1.43)$$

for $i, j = 1, \dots, p$. Because $\gamma_{ij}(h) = \gamma_{ji}(-h)$, it follows that

$$\Gamma(-h) = \Gamma'(h). \quad (1.44)$$

Now, the sample autocovariance matrix of the vector series $\mathbf{x}_t$ is the $p \times p$ matrix of sample sample cross-covariances, defined as




Fig. 1.17. Two-dimensional time series of temperature measurements taken on a rectangular field (64×36 with 17-foot spacing). Data are from Bazza et al. (1988).

$$\widehat{\Gamma}(h) = n^{-1} \sum_{t=1}^{n-h} (\mathbf{x}_{t+h} - \bar{\mathbf{x}})(\mathbf{x}_t - \bar{\mathbf{x}})', \quad (1.45)$$

where

$$\bar{\mathbf{x}} = n^{-1} \sum_{t=1}^n \mathbf{x}_t \quad (1.46)$$

denotes the $p \times 1$ sample mean vector. The symmetry property of the theoretical autocovariance (1.44) extends to the sample autocovariance (1.45), which is defined for negative values by taking

$$\widehat{\Gamma}(-h) = \widehat{\Gamma}(h)'. \quad (1.47)$$

In many applied problems, an observed series may be indexed by more than time alone. For example, the position in space of an experimental unit might be described by two coordinates, say, s_1 and s_2 . We may proceed in these cases by defining a multidimensional process $\mathbf{x}_{\mathbf{s}}$ as a function of the $r \times 1$ vector $\mathbf{s} = (s_1, s_2, \dots, s_r)'$, where s_i denotes the coordinate of the i th index.

Example 1.28 Soil Surface Temperatures

As an example, the two-dimensional ($r = 2$) temperature series x_{s_1, s_2} in [Figure 1.17](#) is indexed by a row number s_1 and a column number s_2 that represent positions on a 64×36 spatial grid set out on an agricultural field. The value of the temperature measured at row s_1 and column s_2 , is denoted by $x_{\mathbf{s}} = x_{s_1, s_2}$. We can note from the two-dimensional plot that a distinct change occurs in the character of the two-dimensional surface starting at about row 40, where the oscillations along the row axis become fairly stable and periodic. For example, averaging over the 36 columns, we may compute an average value for each s_1 as in [Figure 1.18](#). It is clear that the noise present in the first part of the two-dimensional series is nicely averaged out, and we see a clear and consistent temperature signal.

To generate [Figure 1.17](#) and [Figure 1.18](#) in R, use the following commands:

```
persp(1:64, 1:36, soiltemp, phi=30, theta=30, scale=FALSE, expand=4,
      ticktype="detailed", xlab="rows", ylab="cols", zlab="temperature")
plot.ts(rowMeans(soiltemp), xlab="row", ylab="Average Temperature")
```

The autocovariance function of a stationary multidimensional process, $x_{\mathbf{s}}$, can be defined as a function of the multidimensional lag vector, say, $\mathbf{h} = (h_1, h_2, \dots, h_r)'$, as

$$\gamma(\mathbf{h}) = E[(x_{\mathbf{s}+\mathbf{h}} - \mu)(x_{\mathbf{s}} - \mu)], \quad (1.48)$$

where

$$\mu = E(x_{\mathbf{s}}) \quad (1.49)$$

does not depend on the spatial coordinate $\mathbf{s}$. For the two dimensional temperature process, (1.48) becomes

$$\gamma(h_1, h_2) = E[(x_{s_1+h_1, s_2+h_2} - \mu)(x_{s_1, s_2} - \mu)], \quad (1.50)$$

which is a function of lag, both in the row (h_1) and column (h_2) directions.

The multidimensional sample autocovariance function is defined as

$$\hat{\gamma}(\mathbf{h}) = (S_1 S_2 \cdots S_r)^{-1} \sum_{s_1} \sum_{s_2} \cdots \sum_{s_r} (x_{\mathbf{s}+\mathbf{h}} - \bar{x})(x_{\mathbf{s}} - \bar{x}), \quad (1.51)$$

where $\mathbf{s} = (s_1, s_2, \dots, s_r)'$ and the range of summation for each argument is $1 \leq s_i \leq S_i - h_i$, for $i = 1, \dots, r$. The mean is computed over the r -dimensional array, that is,

$$\bar{x} = (S_1 S_2 \cdots S_r)^{-1} \sum_{s_1} \sum_{s_2} \cdots \sum_{s_r} x_{s_1, s_2, \dots, s_r}, \quad (1.52)$$

where the arguments s_i are summed over $1 \leq s_i \leq S_i$. The multidimensional sample autocorrelation function follows, as usual, by taking the scaled ratio

$$\hat{\rho}(\mathbf{h}) = \frac{\hat{\gamma}(\mathbf{h})}{\hat{\gamma}(\mathbf{0})}. \quad (1.53)$$




Fig. 1.18. Row averages of the two-dimensional soil temperature profile. $\bar{x}_{s_1} = \sum_{s_2} x_{s_1, s_2} / 36$.

Example 1.29 Sample ACF of the Soil Temperature Series

The autocorrelation function of the two-dimensional (2d) temperature process can be written in the form

$$\hat{\rho}(h_1, h_2) = \frac{\hat{\gamma}(h_1, h_2)}{\hat{\gamma}(0, 0)},$$

where

$$\hat{\gamma}(h_1, h_2) = (S_1 S_2)^{-1} \sum_{s_1} \sum_{s_2} (x_{s_1+h_1, s_2+h_2} - \bar{x})(x_{s_1, s_2} - \bar{x})$$

Figure 1.19 shows the autocorrelation function for the temperature data, and we note the systematic periodic variation that appears along the rows. The autocovariance over columns seems to be strongest for $h_1 = 0$, implying columns may form replicates of some underlying process that has a periodicity over the rows. This idea can be investigated by examining the mean series over columns as shown in Figure 1.18.

The easiest way (that we know of) to calculate a 2d ACF in R is by using the fast Fourier transform (FFT) as shown below. Unfortunately, the material needed to understand this approach is given in Chapter 4, §4.4. The 2d autocovariance function is obtained in two steps and is contained in `cs` below; $\hat{\gamma}(0, 0)$ is the (1,1) element so that $\hat{\rho}(h_1, h_2)$ is obtained by dividing each element by that value. The 2d ACF is contained in `rs` below, and the rest of the code is simply to arrange the results to yield a nice display.

```
fs = Mod(fft(soiltemp-mean(soiltemp)))^2/(64*36)
cs = Re(fft(fs, inverse=TRUE)/sqrt(64*36)) # ACovF
rs = cs/cs[1,1] # ACF
rs2 = cbind(rs[1:41,21:2], rs[1:41,1:21])
```




Fig. 1.19. Two-dimensional autocorrelation function for the soil temperature data.

```
rs3 = rbind(rs2[41:2,], rs2)
par(mar = c(1,2.5,0,0)+.1)
persp(-40:40, -20:20, rs3, phi=30, theta=30, expand=30, scale="FALSE",
      ticktype="detailed", xlab="row lags", ylab="column lags", zlab="ACF")
```

The sampling requirements for multidimensional processes are rather severe because values must be available over some uniform grid in order to compute the ACF. In some areas of application, such as in soil science, we may prefer to sample a limited number of rows or *transects* and hope these are essentially replicates of the basic underlying phenomenon of interest. One-dimensional methods can then be applied. When observations are irregular in time space, modifications to the estimators need to be made. Systematic approaches to the problems introduced by irregularly spaced observations have been developed by Journel and Huijbregts (1978) or Cressie (1993). We shall not pursue such methods in detail here, but it is worth noting that the introduction of the variogram

$$2V_x(\mathbf{h}) = \text{var}\{x_{\mathbf{s}+\mathbf{h}} - x_{\mathbf{s}}\} \quad (1.54)$$

and its sample estimator

$$2\widehat{V}_x(\mathbf{h}) = \frac{1}{N(\mathbf{h})} \sum_{\mathbf{s}} (x_{\mathbf{s}+\mathbf{h}} - x_{\mathbf{s}})^2 \quad (1.55)$$

play key roles, where $N(\mathbf{h})$ denotes both the number of points located within $\mathbf{h}$, and the sum runs over the points in the neighborhood. Clearly, substantial indexing difficulties will develop from estimators of the kind, and often it will be difficult to find non-negative definite estimators for the covariance function. **Problem 1.27** investigates the relation between the variogram and the autocovariance function in the stationary case.

Problems

Section 1.2

1.1 To compare the earthquake and explosion signals, plot the data displayed in **Figure 1.7** on the same graph using different colors or different line types and comment on the results. (The R code in **Example 1.11** may be of help on how to add lines to existing plots.)

1.2 Consider a signal-plus-noise model of the general form $x_t = s_t + w_t$, where w_t is Gaussian white noise with $\sigma_w^2 = 1$. Simulate and plot $n = 200$ observations from each of the following two models.

(a) $x_t = s_t + w_t$, for $t = 1, \dots, 200$, where

$$s_t = \begin{cases} 0, & t = 1, \dots, 100 \\ 10 \exp\left\{-\frac{(t-100)}{20}\right\} \cos(2\pi t/4), & t = 101, \dots, 200. \end{cases}$$

Hint:

```
s = c(rep(0,100), 10*exp(-(1:100)/20)*cos(2*pi*1:100/4))
x = ts(s + rnorm(200, 0, 1))
plot(x)
```

(b) $x_t = s_t + w_t$, for $t = 1, \dots, 200$, where

$$s_t = \begin{cases} 0, & t = 1, \dots, 100 \\ 10 \exp\left\{-\frac{(t-100)}{200}\right\} \cos(2\pi t/4), & t = 101, \dots, 200. \end{cases}$$

(c) Compare the general appearance of the series (a) and (b) with the earthquake series and the explosion series shown in **Figure 1.7**. In addition, plot (or sketch) and compare the signal modulators (a) $\exp\{-t/20\}$ and (b) $\exp\{-t/200\}$, for $t = 1, 2, \dots, 100$.

Section 1.3

1.3 (a) Generate $n = 100$ observations from the autoregression

$$x_t = -.9x_{t-2} + w_t$$

with $\sigma_w = 1$, using the method described in [Example 1.10](#). Next, apply the moving average filter

$$v_t = (x_t + x_{t-1} + x_{t-2} + x_{t-3})/4$$

to x_t , the data you generated. Now plot x_t as a line and superimpose v_t as a dashed line. Comment on the behavior of x_t and how applying the moving average filter changes that behavior. [Hints: Use `v = filter(x, rep(1/4, 4), sides = 1)` for the filter and note that the R code in [Example 1.11](#) may be of help on how to add lines to existing plots.]

(b) Repeat (a) but with

$$x_t = \cos(2\pi t/4).$$

(c) Repeat (b) but with added $N(0, 1)$ noise,

$$x_t = \cos(2\pi t/4) + w_t.$$

(d) Compare and contrast (a)–(c); i.e., how does the moving average change each series.

Section 1.4

1.4 Show that the autocovariance function can be written as

$$\gamma(s, t) = E[(x_s - \mu_s)(x_t - \mu_t)] = E(x_s x_t) - \mu_s \mu_t,$$

where $E[x_t] = \mu_t$.

1.5 For the two series, x_t , in [Problem 1.2](#) (a) and (b):

- (a) Compute and plot the mean functions $\mu_x(t)$, for $t = 1, \dots, 200$.
- (b) Calculate the autocovariance functions, $\gamma_x(s, t)$, for $s, t = 1, \dots, 200$.

Section 1.5

1.6 Consider the time series

$$x_t = \beta_1 + \beta_2 t + w_t,$$

where β_1 and β_2 are known constants and w_t is a white noise process with variance σ_w^2 .

- (a) Determine whether x_t is stationary.
- (b) Show that the process $y_t = x_t - x_{t-1}$ is stationary.
- (c) Show that the mean of the moving average

$$v_t = \frac{1}{2q+1} \sum_{j=-q}^q x_{t-j}$$

is $\beta_1 + \beta_2 t$, and give a simplified expression for the autocovariance function.

1.7 For a moving average process of the form

$$x_t = w_{t-1} + 2w_t + w_{t+1},$$

where w_t are independent with zero means and variance σ_w^2 , determine the autocovariance and autocorrelation functions as a function of lag $h = s - t$ and plot the ACF as a function of h .

1.8 Consider the random walk with drift model

$$x_t = \delta + x_{t-1} + w_t,$$

for $t = 1, 2, \dots$, with $x_0 = 0$, where w_t is white noise with variance σ_w^2 .

- (a) Show that the model can be written as $x_t = \delta t + \sum_{k=1}^t w_k$.
- (b) Find the mean function and the autocovariance function of x_t .
- (c) Argue that x_t is not stationary.
- (d) Show $\rho_x(t-1, t) = \sqrt{\frac{t-1}{t}} \rightarrow 1$ as $t \rightarrow \infty$. What is the implication of this result?
- (e) Suggest a transformation to make the series stationary, and prove that the transformed series is stationary. (Hint: See **Problem 1.6b**.)

1.9 A time series with a periodic component can be constructed from

$$x_t = U_1 \sin(2\pi\omega_0 t) + U_2 \cos(2\pi\omega_0 t),$$

where U_1 and U_2 are independent random variables with zero means and $E(U_1^2) = E(U_2^2) = \sigma^2$. The constant ω_0 determines the period or time it takes the process to make one complete cycle. Show that this series is weakly stationary with autocovariance function

$$\gamma(h) = \sigma^2 \cos(2\pi\omega_0 h).$$

1.10 Suppose we would like to predict a single stationary series x_t with zero mean and autocorrelation function $\gamma(h)$ at some time in the future, say, $t + \ell$, for $\ell > 0$.

- (a) If we predict using only x_t and some scale multiplier A , show that the mean-square prediction error

$$MSE(A) = E[(x_{t+\ell} - Ax_t)^2]$$

is minimized by the value

$$A = \rho(\ell).$$

(b) Show that the minimum mean-square prediction error is

$$MSE(A) = \gamma(0)[1 - \rho^2(\ell)].$$

(c) Show that if $x_{t+\ell} = Ax_t$, then $\rho(\ell) = 1$ if $A > 0$, and $\rho(\ell) = -1$ if $A < 0$.

1.11 Consider the linear process defined in (1.29).

(a) Verify that the autocovariance function of the process is given by (1.30). Use the result to verify your answer to **Problem 1.7**. *Hint:* For $h \geq 0$, $\text{cov}(x_{t+h}, x_t) = \text{cov}(\sum_k \psi_k w_{t+h-k}, \sum_j \psi_j w_{t-j})$. For each $j \in \mathbb{Z}$, the only “survivor” will be when $k = h + j$.

(b) Show that x_t exists as a limit in mean square (see **Appendix A**).

1.12 For two weakly stationary series x_t and y_t , verify (1.28).

1.13 Consider the two series

$$x_t = w_t$$

$$y_t = w_t - \theta w_{t-1} + u_t,$$

where w_t and u_t are independent white noise series with variances σ_w^2 and σ_u^2 , respectively, and θ is an unspecified constant.

(a) Express the ACF, $\rho_y(h)$, for $h = 0, \pm 1, \pm 2, \dots$ of the series y_t as a function of σ_w^2 , σ_u^2 , and θ .

(b) Determine the CCF, $\rho_{xy}(h)$ relating x_t and y_t .

(c) Show that x_t and y_t are jointly stationary.

1.14 Let x_t be a stationary normal process with mean μ_x and autocovariance function $\gamma(h)$. Define the nonlinear time series

$$y_t = \exp\{x_t\}.$$

(a) Express the mean function $E(y_t)$ in terms of μ_x and $\gamma(0)$. The moment generating function of a normal random variable x with mean μ and variance σ^2 is

$$M_x(\lambda) = E[\exp\{\lambda x\}] = \exp\left\{\mu\lambda + \frac{1}{2}\sigma^2\lambda^2\right\}.$$

(b) Determine the autocovariance function of y_t . The sum of the two normal random variables $x_{t+h} + x_t$ is still a normal random variable.

1.15 Let w_t , for $t = 0, \pm 1, \pm 2, \dots$ be a normal white noise process, and consider the series

$$x_t = w_t w_{t-1}.$$

Determine the mean and autocovariance function of x_t , and state whether it is stationary.

1.16 Consider the series

$$x_t = \sin(2\pi Ut),$$

$t = 1, 2, \dots$, where U has a uniform distribution on the interval $(0, 1)$.

- (a) Prove x_t is weakly stationary.
- (b) Prove x_t is not strictly stationary.

1.17 Suppose we have the linear process x_t generated by

$$x_t = w_t - \theta w_{t-1},$$

$t = 0, 1, 2, \dots$, where $\{w_t\}$ is independent and identically distributed with characteristic function $\phi_w(\cdot)$, and θ is a fixed constant. [Replace “characteristic function” with “moment generating function” if instructed to do so.]

- (a) Express the joint characteristic function of $x_1, x_2, \dots, x_n$, say,

$$\phi_{x_1, x_2, \dots, x_n}(\lambda_1, \lambda_2, \dots, \lambda_n),$$

in terms of $\phi_w(\cdot)$.

- (b) Deduce from (a) that x_t is strictly stationary.

1.18 Suppose that x_t is a linear process of the form (1.29). Prove

$$\sum_{h=-\infty}^{\infty} |\gamma(h)| < \infty.$$

Section 1.6

1.19 Suppose $x_t = \mu + w_t + \theta w_{t-1}$, where $w_t \sim wn(0, \sigma_w^2)$.

- (a) Show that mean function is $E(x_t) = \mu$.
- (b) Show that the autocovariance function of x_t is given by $\gamma_x(0) = \sigma_w^2(1 + \theta^2)$, $\gamma_x(\pm 1) = \sigma_w^2\theta$, and $\gamma_x(h) = 0$ otherwise.
- (c) Show that x_t is stationary for all values of $\theta \in \mathbb{R}$.
- (d) Use (1.33) to calculate $\text{var}(\bar{x})$ for estimating μ when (i) $\theta = 1$, (ii) $\theta = 0$, and (iii) $\theta = -1$.
- (e) In time series, the sample size n is typically large, so that $\frac{(n-1)}{n} \approx 1$. With this as a consideration, comment on the results of part (d); in particular, how does the accuracy in the estimate of the mean μ change for the three different cases?

- 1.20** (a) Simulate a series of $n = 500$ Gaussian white noise observations as in **Example 1.8** and compute the sample ACF, $\hat{\rho}(h)$, to lag 20. Compare the sample ACF you obtain to the actual ACF, $\rho(h)$. [Recall **Example 1.19**.]
- (b) Repeat part (a) using only $n = 50$. How does changing n affect the results?

- 1.21** (a) Simulate a series of $n = 500$ moving average observations as in **Example 1.9** and compute the sample ACF, $\hat{\rho}(h)$, to lag 20. Compare the sample ACF you obtain to the actual ACF, $\rho(h)$. [Recall **Example 1.21**.]

(b) Repeat part (a) using only $n = 50$. How does changing n affect the results?

1.22 Although the model in **Problem 1.2(a)** is not stationary (Why?), the sample ACF can be informative. For the data you generated in that problem, calculate and plot the sample ACF, and then comment.

1.23 Simulate a series of $n = 500$ observations from the signal-plus-noise model presented in **Example 1.12** with $\sigma_w^2 = 1$. Compute the sample ACF to lag 100 of the data you generated and comment.

1.24 For the time series y_t described in **Example 1.24**, verify the stated result that $\rho_y(1) = -.47$ and $\rho_y(h) = 0$ for $h > 1$.

1.25 A real-valued function $g(t)$, defined on the integers, is non-negative definite if and only if

$$\sum_{i=1}^n \sum_{j=1}^n a_i g(t_i - t_j) a_j \geq 0$$

for all positive integers n and for all vectors $\mathbf{a} = (a_1, a_2, \dots, a_n)'$ and $\mathbf{t} = (t_1, t_2, \dots, t_n)'$. For the matrix $G = \{g(t_i - t_j); i, j = 1, 2, \dots, n\}$, this implies that $\mathbf{a}'G\mathbf{a} \geq 0$ for all vectors $\mathbf{a}$. It is called positive definite if we can replace ' $\geq$ ' with ' $>$ ' for all $\mathbf{a} \neq \mathbf{0}$, the zero vector.

- (a) Prove that $\gamma(h)$, the autocovariance function of a stationary process, is a non-negative definite function.
- (b) Verify that the sample autocovariance $\hat{\gamma}(h)$ is a non-negative definite function.

Section 1.7

1.26 Consider a collection of time series $x_{1t}, x_{2t}, \dots, x_{Nt}$ that are observing some common signal μ_t observed in noise processes $e_{1t}, e_{2t}, \dots, e_{Nt}$, with a model for the j -th observed series given by

$$x_{jt} = \mu_t + e_{jt}.$$

Suppose the noise series have zero means and are uncorrelated for different j . The common autocovariance functions of all series are given by $\gamma_e(s, t)$. Define the sample mean

$$\bar{x}_t = \frac{1}{N} \sum_{j=1}^N x_{jt}.$$

- (a) Show that $E[\bar{x}_t] = \mu_t$.
- (b) Show that $E[(\bar{x}_t - \mu)^2] = N^{-1}\gamma_e(t, t)$.
- (c) How can we use the results in estimating the common signal?

1.27 A concept used in geostatistics, see Journel and Huijbregts (1978) or Cressie (1993), is that of the variogram, defined for a spatial process $x_{\mathbf{s}}$, $\mathbf{s} = (s_1, s_2)$, for $s_1, s_2 = 0, \pm 1, \pm 2, \dots$, as

$$V_x(\mathbf{h}) = \frac{1}{2} E[(x_{\mathbf{s}+\mathbf{h}} - x_{\mathbf{s}})^2],$$

where $\mathbf{h} = (h_1, h_2)$, for $h_1, h_2 = 0, \pm 1, \pm 2, \dots$. Show that, for a stationary process, the variogram and autocovariance functions can be related through

$$V_x(\mathbf{h}) = \gamma(\mathbf{0}) - \gamma(\mathbf{h}),$$

where $\gamma(\mathbf{h})$ is the usual lag $\mathbf{h}$ covariance function and $\mathbf{0} = (0, 0)$. Note the easy extension to any spatial dimension.

The following problems require the material given in [Appendix A](#)

1.28 Suppose $x_t = \beta_0 + \beta_1 t$, where β_0 and β_1 are constants. Prove as $n \rightarrow \infty$, $\hat{\rho}_x(h) \rightarrow 1$ for fixed h , where $\hat{\rho}_x(h)$ is the ACF ([1.35](#)).

1.29 (a) Suppose x_t is a weakly stationary time series with mean zero and with absolutely summable autocovariance function, $\gamma(h)$, such that

$$\sum_{h=-\infty}^{\infty} \gamma(h) = 0.$$

Prove that $\sqrt{n} \bar{x} \xrightarrow{P} 0$, where $\bar{x}$ is the sample mean ([1.32](#)).

(b) Give an example of a process that satisfies the conditions of part (a). What is special about this process?

1.30 Let x_t be a linear process of the form ([A.43](#))–([A.44](#)). If we define

$$\tilde{\gamma}(h) = n^{-1} \sum_{t=1}^n (x_{t+h} - \mu_x)(x_t - \mu_x),$$

show that

$$n^{1/2}(\tilde{\gamma}(h) - \hat{\gamma}(h)) = o_p(1).$$

Hint: The Markov Inequality

$$P\{|x| \geq \epsilon\} < \frac{E|x|}{\epsilon}$$

can be helpful for the cross-product terms.

1.31 For a linear process of the form

$$x_t = \sum_{j=0}^{\infty} \phi^j w_{t-j},$$

where $\{w_t\}$ satisfies the conditions of Theorem A.7 and $|\phi| < 1$, show that

$$\sqrt{n} \frac{(\widehat{\rho}_x(1) - \rho_x(1))}{\sqrt{1 - \rho_x^2(1)}} \xrightarrow{d} N(0, 1),$$

and construct a 95% confidence interval for ϕ when $\widehat{\rho}_x(1) = .64$ and $n = 100$.

1.32 Let $\{x_t; t = 0, \pm 1, \pm 2, \dots\}$ be iid $(0, \sigma^2)$.

(a) For $h \geq 1$ and $k \geq 1$, show that $x_t x_{t+h}$ and $x_s x_{s+k}$ are uncorrelated for all $s \neq t$.

(b) For fixed $h \geq 1$, show that the $h \times 1$ vector

$$\sigma^{-2} n^{-1/2} \sum_{t=1}^n (x_t x_{t+1}, \dots, x_t x_{t+h})' \xrightarrow{d} (z_1, \dots, z_h)'$$

where $z_1, \dots, z_h$ are iid $N(0, 1)$ random variables. [Hint: Use the Cramér-Wold device.]

(c) Show, for each $h \geq 1$,

$$n^{-1/2} \left[\sum_{t=1}^n x_t x_{t+h} - \sum_{t=1}^{n-h} (x_t - \bar{x})(x_{t+h} - \bar{x}) \right] \xrightarrow{p} 0 \quad \text{as } n \rightarrow \infty$$

where $\bar{x} = n^{-1} \sum_{t=1}^n x_t$.

(d) Noting that $n^{-1} \sum_{t=1}^n x_t^2 \xrightarrow{p} \sigma^2$ by the WLLN, conclude that

$$n^{1/2} [\widehat{\rho}(1), \dots, \widehat{\rho}(h)]' \xrightarrow{d} (z_1, \dots, z_h)'$$

where $\widehat{\rho}(h)$ is the sample ACF of the data $x_1, \dots, x_n$.

Chapter 2

Time Series Regression and Exploratory Data Analysis

In this chapter we introduce classical multiple linear regression in a time series context, model selection, exploratory data analysis for preprocessing nonstationary time series, (for example trend removal), variance stabilization, and nonparametric smoothing of time series.

2.1 Classical Regression in the Time Series Context

We begin our discussion of linear regression in the time series context by assuming some output or dependent time series, say, x_t , for $t = 1, \dots, n$, is being influenced by a collection of possible inputs or independent series, say, $z_{t1}, z_{t2}, \dots, z_{tq}$, where we first regard the inputs as fixed and known. This assumption, necessary for applying conventional linear regression, will be relaxed later on. We express this relation through the linear regression model

$$x_t = \beta_0 + \beta_1 z_{t1} + \beta_2 z_{t2} + \dots + \beta_q z_{tq} + w_t, \quad (2.1)$$

where $\beta_0, \beta_1, \dots, \beta_q$ are unknown fixed regression coefficients, and $\{w_t\}$ is a random error or noise process consisting of independent and identically distributed (iid) normal variables with mean zero and variance σ_w^2 ; we will relax the iid assumption later. A more general setting within which to embed mean square estimation and linear regression is given in [Appendix B](#), where we introduce Hilbert spaces and the Projection Theorem.

Example 2.1 Estimating a Linear Trend

Consider the global temperature data, say x_t , shown in [Figure 1.2](#) and [Figure 2.1](#). As discussed in [Example 1.2](#), there is an apparent upward trend in the series that has been used to argue the global warming hypothesis. We might use simple linear regression to estimate that trend by fitting the model

$$x_t = \beta_0 + \beta_1 z_t + w_t, \quad z_t = 1880, 1857, \dots, 2009.$$

This is in the form of the regression model (2.1) with $q = 1$. Note that we are making the assumption that the errors, w_t , are an iid normal sequence, which may not be true; the problem of autocorrelated errors is discussed in detail in Chapter 3.

In ordinary least squares (OLS), we minimize the error sum of squares

$$Q = \sum_{t=1}^n w_t^2 = \sum_{t=1}^n (x_t - [\beta_0 + \beta_1 z_t])^2$$

with respect to β_i for $i = 0, 1$. In this case we can use simple calculus to evaluate $\partial Q / \partial \beta_i = 0$ for $i = 0, 1$, to obtain two equations to solve for the β s. The OLS estimates of the coefficients are explicit and given by

$$\hat{\beta}_1 = \frac{\sum_{t=1}^n (x_t - \bar{x})(z_t - \bar{z})}{\sum_{t=1}^n (z_t - \bar{z})^2} \quad \text{and} \quad \hat{\beta}_0 = \bar{x} - \hat{\beta}_1 \bar{z},$$

where $\bar{x} = \sum_t x_t / n$ and $\bar{z}_t = \sum_t z_t / n$ are the respective sample means. Note that we could have used, for example, $z_t = 1, \dots, 130$, without affecting the interpretation of the slope coefficient, β_1 ; only the intercept, β_0 , would change.

Using R, we obtained the estimated coefficients $\hat{\beta}_0 = -11.2$, and $\hat{\beta}_1 = .006$ (with a standard error of .0003) yielding a highly significant estimated increase of .6 degrees centigrade per 100 years. Finally, Figure 2.1 shows the global temperature data, say x_t , with the estimated trend, say $\hat{x}_t = -11.2 + .006z_t$, superimposed. *It is apparent that the estimated trend line obtained via simple linear regression does not quite capture the trend of the data and better models will be needed.* To perform this analysis in R, use the following commands:

```
summary(fit <- lm(gtemp~time(gtemp))) # regress gtemp on time - view results
plot(gtemp, type="o", ylab="Global Temperature Deviation")
abline(fit) # add the fitted regression line to the plot
```

The multiple linear regression model described by (2.1) can be conveniently written in a more general notation by defining the column vectors $\mathbf{z}_t = (1, z_{t1}, z_{t2}, \dots, z_{tq})'$ and $\boldsymbol{\beta} = (\beta_0, \beta_1, \dots, \beta_q)'$, where $'$ denotes transpose, so (2.1) can be written in the alternate form

$$x_t = \beta_0 + \beta_1 z_{t1} + \dots + \beta_q z_{tq} + w_t = \boldsymbol{\beta}' \mathbf{z}_t + w_t. \quad (2.2)$$

where $w_t \sim \text{iid } N(0, \sigma_w^2)$. As in the previous example, OLS estimation finds the coefficient vector $\boldsymbol{\beta}$ that minimizes the error sum of squares

$$Q = \sum_{t=1}^n w_t^2 = \sum_{t=1}^n (x_t - \boldsymbol{\beta}' \mathbf{z}_t)^2, \quad (2.3)$$

with respect to $\beta_0, \beta_1, \dots, \beta_q$. This minimization can be accomplished by differentiating (2.3) with respect to the vector $\boldsymbol{\beta}$ or by using the properties of projections. Either way, the solution must satisfy $\sum_{t=1}^n (x_t - \boldsymbol{\beta}' \mathbf{z}_t) \mathbf{z}_t' = 0$. This procedure gives the *normal equations*




Fig. 2.1. Global temperature deviations shown in Figure 1.2 with fitted linear trend line.

$$\left(\sum_{t=1}^n \mathbf{z}_t \mathbf{z}_t' \right) \hat{\boldsymbol{\beta}} = \sum_{t=1}^n \mathbf{z}_t x_t. \quad (2.4)$$

If $\sum_{t=1}^n \mathbf{z}_t \mathbf{z}_t'$ is non-singular, the least squares estimate of $\boldsymbol{\beta}$ is

$$\hat{\boldsymbol{\beta}} = \left(\sum_{t=1}^n \mathbf{z}_t \mathbf{z}_t' \right)^{-1} \sum_{t=1}^n \mathbf{z}_t x_t.$$

The minimized error sum of squares (2.3), denoted SSE , can be written as

$$SSE = \sum_{t=1}^n (x_t - \hat{\boldsymbol{\beta}}' \mathbf{z}_t)^2. \quad (2.5)$$

The ordinary least squares estimators are unbiased, i.e., $E(\hat{\boldsymbol{\beta}}) = \boldsymbol{\beta}$, and have the smallest variance within the class of linear unbiased estimators.

If the errors w_t are normally distributed, $\hat{\boldsymbol{\beta}}$ is also the maximum likelihood estimator for $\boldsymbol{\beta}$ and is normally distributed with

$$\text{cov}(\hat{\boldsymbol{\beta}}) = \sigma_w^2 \mathbf{C}, \quad (2.6)$$

where

$$\mathbf{C} = \left(\sum_{t=1}^n \mathbf{z}_t \mathbf{z}_t' \right)^{-1} \quad (2.7)$$

is a convenient notation. An unbiased estimator for the variance σ_w^2 is

$$s_w^2 = MSE = \frac{SSE}{n - (q + 1)}, \quad (2.8)$$

Table 2.1. Analysis of Variance for Regression

Source	df	Sum of Squares	Mean Square	F
$z_{t,r+1:q}$	$q - r$	$SSR = SSE_r - SSE$	$MSR = SSR/(q - r)$	$F = \frac{MSR}{MSE}$
Error	$n - (q + 1)$	SSE	$MSE = SSE/(n - q - 1)$	

where MSE denotes the *mean squared error*. Under the normal assumption,

$$t = \frac{(\hat{\beta}_i - \beta_i)}{s_w \sqrt{c_{ii}}} \quad (2.9)$$

has the t-distribution with $n - (q + 1)$ degrees of freedom; c_{ii} denotes the i -th diagonal element of C , as defined in (2.7). This result is often used for individual tests of the null hypothesis $H_0: \beta_i = 0$ for $i = 1, \dots, q$.

Various competing models are often of interest to isolate or select the best subset of independent variables. Suppose a proposed model specifies that only a subset $r < q$ independent variables, say, $z_{t,1:r} = \{z_{t1}, z_{t2}, \dots, z_{tr}\}$ is influencing the dependent variable x_t . The reduced model is

$$x_t = \beta_0 + \beta_1 z_{t1} + \dots + \beta_r z_{tr} + w_t \quad (2.10)$$

where $\beta_1, \beta_2, \dots, \beta_r$ are a subset of coefficients of the original q variables.

The null hypothesis in this case is $H_0: \beta_{r+1} = \dots = \beta_q = 0$. We can test the reduced model (2.10) against the full model (2.2) by comparing the error sums of squares under the two models using the F -statistic

$$F = \frac{(SSE_r - SSE)/(q - r)}{SSE/(n - q - 1)} = \frac{MSR}{MSE}, \quad (2.11)$$

where SSE_r is the error sum of squares under the reduced model (2.10). Note that $SSE_r \geq SSE$ because the full model has more parameters. If $H_0: \beta_{r+1} = \dots = \beta_q = 0$ is true, then $SSE_r \approx SSE$ because the estimates of those β s will be close to 0. Hence, we do not believe H_0 if $SSR = SSE_r - SSE$ is big. Under the null hypothesis, (2.11) has a central F -distribution with $q - r$ and $n - q - 1$ degrees of freedom when (2.10) is the correct model.

These results are often summarized in an Analysis of Variance (ANOVA) table as given in Table 2.1 for this particular case. The difference in the numerator is often called the regression sum of squares (SSR). The null hypothesis is rejected at level α if $F > F_{n-q-1}^{q-r}(\alpha)$, the $1 - \alpha$ percentile of the F distribution with $q - r$ numerator and $n - q - 1$ denominator degrees of freedom.

A special case of interest is the null hypothesis $H_0: \beta_1 = \dots = \beta_q = 0$. In this case $r = 0$, and the model in (2.10) becomes

$$x_t = \beta_0 + w_t.$$

We may measure the proportion of variation accounted for by all the variables using

$$R^2 = \frac{SSE_0 - SSE}{SSE_0}, \quad (2.12)$$

where the residual sum of squares under the reduced model is

$$SSE_0 = \sum_{t=1}^n (x_t - \bar{x})^2. \quad (2.13)$$

In this case SSE_0 is the sum of squared deviations from the mean $\bar{x}$ and is otherwise known as the adjusted total sum of squares. The measure R^2 is called the *squared multiple correlation* between x_t and the variables $z_{t1}, z_{t2}, \dots, z_{tq}$.

The techniques discussed in the previous paragraph can be used to test various models against one another using the F test given in (2.11). These tests have been used in the past in a stepwise manner, where variables are added or deleted when the values from the F -test either exceed or fail to exceed some predetermined levels. The procedure, called stepwise multiple regression, is useful in arriving at a set of useful variables. An alternative is to focus on a procedure for model selection that does not proceed sequentially, but simply evaluates each model on its own merits. Suppose we consider a normal regression model with k coefficients and denote the maximum likelihood estimator for the variance as

$$\hat{\sigma}_k^2 = \frac{SSE(k)}{n}, \quad (2.14)$$

where $SSE(k)$ denotes the residual sum of squares under the model with k regression coefficients. Then, Akaike (1969, 1973, 1974) suggested measuring the goodness of fit for this particular model by balancing the error of the fit against the number of parameters in the model; we define the following.¹

Definition 2.1 Akaike's Information Criterion (AIC)

$$AIC = \log \hat{\sigma}_k^2 + \frac{n + 2k}{n}, \quad (2.15)$$

where $\hat{\sigma}_k^2$ is given by (2.14) and k is the number of parameters in the model.

The value of k yielding the minimum AIC specifies the best model. The idea is roughly that minimizing $\hat{\sigma}_k^2$ would be a reasonable objective, except that it decreases monotonically as k increases. Therefore, we ought to penalize the error variance by a term proportional to the number of parameters. The choice for the penalty term given by (2.15) is not the only one, and a considerable literature is available advocating different penalty terms. A corrected form, suggested by Sugiura (1978), and expanded by Hurvich and Tsai (1989), can be based on small-sample distributional results for the linear regression model (details are provided in Problem 2.4 and Problem 2.5). The corrected form is defined as follows.

¹ Formally, AIC is defined as $-2 \log L_k + 2k$ where L_k is the maximized likelihood and k is the number of parameters in the model. For the normal regression problem, AIC can be reduced to the form given by (2.15). AIC is an estimate of the Kullback-Leibler discrepancy between a true model and a candidate model; see Problem 2.4 and Problem 2.5 for further details.

Definition 2.2 AIC, Bias Corrected (AICc)

$$\text{AICc} = \log \hat{\sigma}_k^2 + \frac{n+k}{n-k-2}, \quad (2.16)$$

where $\hat{\sigma}_k^2$ is given by (2.14), k is the number of parameters in the model, and n is the sample size.

We may also derive a correction term based on Bayesian arguments, as in Schwarz (1978), which leads to the following.

Definition 2.3 Bayesian Information Criterion (BIC)

$$\text{BIC} = \log \hat{\sigma}_k^2 + \frac{k \log n}{n}, \quad (2.17)$$

using the same notation as in *Definition 2.2*.

BIC is also called the Schwarz Information Criterion (SIC); see also Rissanen (1978) for an approach yielding the same statistic based on a minimum description length argument. Various simulation studies have tended to verify that BIC does well at getting the correct order in large samples, whereas AICc tends to be superior in smaller samples where the relative number of parameters is large; see McQuarrie and Tsai (1998) for detailed comparisons. In fitting regression models, two measures that have been used in the past are adjusted R-squared, which is essentially s_w^2 , and Mallows C_p , Mallows (1973), which we do not consider in this context.

Example 2.2 Pollution, Temperature and Mortality

The data shown in *Figure 2.2* are extracted series from a study by Shumway et al. (1988) of the possible effects of temperature and pollution on weekly mortality in Los Angeles County. Note the strong seasonal components in all of the series, corresponding to winter-summer variations and the downward trend in the cardiovascular mortality over the 10-year period.

A scatterplot matrix, shown in *Figure 2.3*, indicates a possible linear relation between mortality and the pollutant particulates and a possible relation to temperature. Note the curvilinear shape of the temperature mortality curve, indicating that higher temperatures as well as lower temperatures are associated with increases in cardiovascular mortality.

Based on the scatterplot matrix, we entertain, tentatively, four models where M_t denotes cardiovascular mortality, T_t denotes temperature and P_t denotes the particulate levels. They are

$$M_t = \beta_1 + \beta_2 t + w_t \quad (2.18)$$

$$M_t = \beta_1 + \beta_2 t + \beta_3 (T_t - T.) + w_t \quad (2.19)$$

$$M_t = \beta_1 + \beta_2 t + \beta_3 (T_t - T.) + \beta_4 (T_t - T.)^2 + w_t \quad (2.20)$$

$$M_t = \beta_1 + \beta_2 t + \beta_3 (T_t - T.) + \beta_4 (T_t - T.)^2 + \beta_5 P_t + w_t \quad (2.21)$$




Fig. 2.2. Average weekly cardiovascular mortality (top), temperature (middle) and particulate pollution (bottom) in Los Angeles County. There are 508 six-day smoothed averages obtained by filtering daily values over the 10 year period 1970-1979.

Table 2.2. Summary Statistics for Mortality Models

Model	k	SSE	df	MSE	R^2	AIC	BIC
(2.18)	2	40,020	506	79.0	.21	5.38	5.40
(2.19)	3	31,413	505	62.2	.38	5.14	5.17
(2.20)	4	27,985	504	55.5	.45	5.03	5.07
(2.21)	5	20,508	503	40.8	.60	4.72	4.77

where we adjust temperature for its mean, $T = 74.6$, to avoid scaling problems. It is clear that (2.18) is a trend only model, (2.19) is linear temperature, (2.20) is curvilinear temperature and (2.21) is curvilinear temperature and pollution. We summarize some of the statistics given for this particular case in 2.2. The values of R^2 were computed by noting that $SSE_1 = 50,687$ using (2.13).

We note that each model does substantially better than the one before it and that the model including temperature, temperature squared, and particulates does the




Fig. 2.3. Scatterplot matrix showing relations between mortality, temperature, and pollution.

best, accounting for some 60% of the variability and with the best value for AIC and BIC (because of the large sample size, AIC and AICc are nearly the same). Note that one can compare any two models using the residual sums of squares and (2.11). Hence, a model with only trend could be compared to the full model using $q = 5$, $r = 2$, $n = 508$, so

$$F_{3,503} = \frac{(40,020 - 20,508)/3}{20,508/503} = 160,$$

which exceeds $F_{3,503}(.001) = 5.51$. We obtain the best prediction model,

$$\begin{aligned} \widehat{M}_t = & 81.59 - .027_{(.002)}t - .473_{(.032)}(T_t - 74.6) \\ & + .023_{(.003)}(T_t - 74.6)^2 + .255_{(.019)}P_t, \end{aligned}$$

for mortality, where the standard errors, computed from (2.6)–(2.8), are given in parentheses. As expected, a negative trend is present in time as well as a negative coefficient for adjusted temperature. The quadratic effect of temperature can clearly be seen in the scatterplots of Figure 2.3. Pollution weights positively and can be

interpreted as the incremental contribution to daily deaths per unit of particulate pollution. It would still be essential to check the residuals $\hat{w}_t = M_t - \hat{M}_t$ for autocorrelation (of which there is a substantial amount), but we defer this question to to §3.9 when we discuss regression with correlated errors.

Below is the R code to plot the series, display the scatterplot matrix, fit the final regression model (2.21), and compute the corresponding values of AIC, AICc and BIC.² Finally, the use of `na.action` in `lm()` is to retain the time series attributes for the residuals and fitted values.

```
par(mfrow=c(3,1)) # plot the data
plot(cmort, main="Cardiovascular Mortality", xlab="", ylab="")
plot(temp, main="Temperature", xlab="", ylab="")
plot(part, main="Particulates", xlab="", ylab="")
dev.new() # open a new graphic device for the scatterplot matrix
pairs(cbind(Mortality=cmort, Temperature=temp, Particulates=part))
temp = temp - mean(temp) # center temperature
temp2 = temp^2
trend = time(cmort) # time
fit = lm(cmort ~ trend + temp + temp2 + part, na.action=NULL)
summary(fit) # regression results
summary(aov(fit)) # ANOVA table (compare to next line)
summary(aov(lm(cmort ~ cbind(trend, temp, temp2, part)))) # Table 2.1
num = length(cmort) # sample size
AIC(fit)/num - log(2*pi) # AIC
BIC(fit)/num - log(2*pi) # BIC
(AICc = log(sum(resid(fit)^2)/num) + (num+5)/(num-5-2)) # AICc
```

As previously mentioned, it is possible to include lagged variables in time series regression models and we will continue to discuss this type of problem throughout the text. This concept is explored further in **Problem 2.2** and **Problem 2.11**. The following is a simple example of lagged regression.

Example 2.3 Regression With Lagged Variables

In **Example 1.26**, we discovered that the Southern Oscillation Index (SOI) measured at time $t - 6$ months is associated with the Recruitment series at time t , indicating that the SOI leads the Recruitment series by six months. Although there is evidence that the relationship is not linear (this is discussed further in **Example 2.7** and **Example 2.8**), consider the following regression,

$$R_t = \beta_1 + \beta_2 S_{t-6} + w_t, \quad (2.22)$$

where R_t denotes Recruitment for month t and S_{t-6} denotes SOI six months prior. Assuming the w_t sequence is white, the fitted model is

$$\hat{R}_t = 65.79 - 44.28_{(2.78)} S_{t-6} \quad (2.23)$$

with $\hat{\sigma}_w = 22.5$ on 445 degrees of freedom. This result indicates the strong predictive ability of SOI for Recruitment six months in advance. Of course, it is still essential to check the the model assumptions, but again we defer this until later.

² The easiest way to extract AIC and BIC from an `lm()` run in R is to use the command `AIC()` or `BIC()`. Our definitions differ from R by terms that do not change from model to model. In the example, we show how to obtain (2.15) and (2.17) from the R output. It is more difficult to obtain AICc.

Performing lagged regression in R is a little difficult because the series must be aligned prior to running the regression. The easiest way to do this is to create a data frame that we call `fish` using `ts.intersect`, which aligns the lagged series.

```
fish = ts.intersect(rec, soiL6=lag(soi,-6), dframe=TRUE)
summary(fit1 <- lm(rec~soiL6, data=fish, na.action=NULL))
```

The headache of aligning the lagged series can be avoided by using the R package `dynlm`, which must be downloaded and installed.

```
require(dynlm)
summary(fit2 <- dynlm(rec~ L(soi,6)))
```

In the `dynlm` example, `fit2` is similar to the `fit1` object, but the time series attributes are retained without any additional commands.

2.2 Exploratory Data Analysis

In general, it is necessary for time series data to be stationary, so averaging lagged products over time, as in the previous section, will be a sensible thing to do. With time series data, it is the dependence between the values of the series that is important to measure; we must, at least, be able to estimate autocorrelations with precision. It would be difficult to measure that dependence if the dependence structure is not regular or is changing at every time point. Hence, to achieve any meaningful statistical analysis of time series data, it will be crucial that, if nothing else, the mean and the autocovariance functions satisfy the conditions of stationarity (for at least some reasonable stretch of time) stated in [Definition 1.7](#). Often, this is not the case, and we will mention some methods in this section for playing down the effects of nonstationarity so the stationary properties of the series may be studied.

A number of our examples came from clearly nonstationary series. The Johnson & Johnson series in [Figure 1.1](#) has a mean that increases exponentially over time, and the increase in the magnitude of the fluctuations around this trend causes changes in the covariance function; the variance of the process, for example, clearly increases as one progresses over the length of the series. Also, the global temperature series shown in [Figure 1.2](#) contains some evidence of a trend over time; human-induced global warming advocates seize on this as empirical evidence to advance their hypothesis that temperatures are increasing.

Perhaps the easiest form of nonstationarity to work with is the trend stationary model wherein the process has stationary behavior around a trend. We may write this type of model as

$$x_t = \mu_t + y_t \quad (2.24)$$

where x_t are the observations, μ_t denotes the trend, and y_t is a stationary process. Quite often, strong trend, μ_t , will obscure the behavior of the stationary process, y_t , as we shall see in numerous examples. Hence, there is some advantage to removing the trend as a first step in an exploratory analysis of such time series. The steps involved are to obtain a reasonable estimate of the trend component, say $\hat{\mu}_t$, and then work with the residuals

$$\hat{y}_t = x_t - \hat{\mu}_t. \quad (2.25)$$

Example 2.4 Detrending Global Temperature

Here we suppose the model is of the form of (2.24),

$$x_t = \mu_t + y_t,$$

where, as we suggested in the analysis of the global temperature data presented in Example 2.1, a straight line might be a reasonable model for the trend, i.e.,

$$\mu_t = \beta_1 + \beta_2 t.$$

In that example, we estimated the trend using ordinary least squares³ and found

$$\hat{\mu}_t = -11.2 + .006 t.$$

Figure 2.1 shows the data with the estimated trend line superimposed. To obtain the detrended series we simply subtract $\hat{\mu}_t$ from the observations, x_t , to obtain the detrended series

$$\hat{y}_t = x_t + 11.2 - .006 t.$$

The top graph of Figure 2.4 shows the detrended series. Figure 2.5 shows the ACF of the original data (top panel) as well as the ACF of the detrended data (middle panel).

To detrend in the series in R, use the following commands. We also show how to difference and plot the differenced data; we discuss differencing after this example. In addition, we show how to generate the sample ACFs displayed in Figure 2.5.

```
fit = lm(gtemp~time(gtemp), na.action=NULL) # regress gtemp on time
par(mfrow=c(2,1))
plot(resid(fit), type="o", main="detrended")
plot(diff(gtemp), type="o", main="first difference")
par(mfrow=c(3,1)) # plot ACFs
acf(gtemp, 48, main="gtemp")
acf(resid(fit), 48, main="detrended")
acf(diff(gtemp), 48, main="first difference")
```

In Example 1.11 and the corresponding Figure 1.10 we saw that a random walk might also be a good model for trend. That is, rather than modeling trend as fixed (as in Example 2.4), we might model trend as a stochastic component using the random walk with drift model,

$$\mu_t = \delta + \mu_{t-1} + w_t, \quad (2.26)$$

where w_t is white noise and is independent of y_t . If the appropriate model is (2.24), then differencing the data, x_t , yields a stationary process; that is,

$$\begin{aligned} x_t - x_{t-1} &= (\mu_t + y_t) - (\mu_{t-1} + y_{t-1}) \\ &= \delta + w_t + y_t - y_{t-1}. \end{aligned} \quad (2.27)$$

³ Because the error term, y_t , is not assumed to be iid, the reader may feel that weighted least squares is called for in this case. The problem is, we do not know the behavior of y_t and that is precisely what we are trying to assess at this stage. A notable result by Grenander and Rosenblatt (1957, Ch 7), however, is that under mild conditions on y_t , for polynomial regression or periodic regression, asymptotically, ordinary least squares is equivalent to weighted least squares.




Fig. 2.4. Detrended (top) and differenced (bottom) global temperature series. The original data are shown in [Figure 1.2](#) and [Figure 2.1](#).

It is easy to show $z_t = y_t - y_{t-1}$ is stationary using [footnote 3](#) of Chapter 1. That is, because y_t is stationary,

$$\begin{aligned}\gamma_z(h) &= \text{cov}(z_{t+h}, z_t) = \text{cov}(y_{t+h} - y_{t+h-1}, y_t - y_{t-1}) \\ &= 2\gamma_y(h) - \gamma_y(h+1) - \gamma_y(h-1)\end{aligned}$$

is independent of time; we leave it as an exercise ([Problem 2.7](#)) to show that $x_t - x_{t-1}$ in [\(2.27\)](#) is stationary.

One advantage of differencing over detrending to remove trend is that no parameters are estimated in the differencing operation. One disadvantage, however, is that differencing does not yield an estimate of the stationary process y_t as can be seen in [\(2.27\)](#). If an estimate of y_t is essential, then detrending may be more appropriate. If the goal is to coerce the data to stationarity, then differencing may be more appropriate. Differencing is also a viable tool if the trend is fixed, as in [Example 2.4](#). That is, e.g., if $\mu_t = \beta_1 + \beta_2 t$ in the model [\(2.24\)](#), differencing the data produces stationarity (see [Problem 2.6](#)):




Fig. 2.5. Sample ACFs of the global temperature (top), and of the detrended (middle) and the differenced (bottom) series.

$$x_t - x_{t-1} = (\mu_t + y_t) - (\mu_{t-1} + y_{t-1}) = \beta_2 + y_t - y_{t-1}.$$

Because differencing plays a central role in time series analysis, it receives its own notation. The first difference is denoted as

$$\nabla x_t = x_t - x_{t-1}. \quad (2.28)$$

As we have seen, the first difference eliminates a linear trend. A second difference, that is, the difference of (2.28), can eliminate a quadratic trend, and so on. In order to define higher differences, we need a variation in notation that we will use often in our discussion of ARIMA models in Chapter 3.

Definition 2.4 We define the **backshift operator** by

$$Bx_t = x_{t-1}$$

and extend it to powers $B^2x_t = B(Bx_t) = Bx_{t-1} = x_{t-2}$, and so on. Thus,

$$B^k x_t = x_{t-k}. \quad (2.29)$$

The idea of an inverse operator can also be given if we require $B^{-1}B = 1$, so that

$$x_t = B^{-1}Bx_t = B^{-1}x_{t-1}.$$

That is, B^{-1} is the *forward-shift operator*. In addition, it is clear that we may rewrite (2.28) as

$$\nabla x_t = (1 - B)x_t, \quad (2.30)$$

and we may extend the notion further. For example, the second difference becomes

$$\nabla^2 x_t = (1 - B)^2 x_t = (1 - 2B + B^2)x_t = x_t - 2x_{t-1} + x_{t-2} \quad (2.31)$$

by the linearity of the operator. To check, just take the difference of the first difference $\nabla(\nabla x_t) = \nabla(x_t - x_{t-1}) = (x_t - x_{t-1}) - (x_{t-1} - x_{t-2})$.

Definition 2.5 Differences of order d are defined as

$$\nabla^d = (1 - B)^d, \quad (2.32)$$

where we may expand the operator $(1 - B)^d$ algebraically to evaluate for higher integer values of d . When $d = 1$, we drop it from the notation.

The first difference (2.28) is an example of a linear filter applied to eliminate a trend. Other filters, formed by averaging values near x_t , can produce adjusted series that eliminate other kinds of unwanted fluctuations, as in Chapter 3. The differencing technique is an important component of the ARIMA model of Box and Jenkins (1970) (see also Box et al., 1994), to be discussed in Chapter 3.

Example 2.5 Differencing Global Temperature

The first difference of the global temperature series, also shown in Figure 2.4, produces different results than removing trend by detrending via regression. For example, the differenced series does not contain the long middle cycle we observe in the detrended series. The ACF of this series is also shown in Figure 2.5. In this case it appears that the differenced process shows minimal autocorrelation, which may imply the global temperature series is nearly a random walk with drift. It is interesting to note that if the series is a random walk with drift, the mean of the differenced series, which is an estimate of the drift, is about .0066 (but with a large standard error):

```
mean(diff(gtemp))    # = 0.00659 (drift)
sd(diff(gtemp))/sqrt(length(diff(gtemp))) # = 0.00966 (SE)
```

An alternative to differencing is a less-severe operation that still assumes stationarity of the underlying time series. This alternative, called fractional differencing, extends the notion of the difference operator (2.32) to fractional powers $-.5 < d < .5$, which still define stationary processes. Granger and Joyeux (1980) and Hosking (1981) introduced long memory time series, which corresponds to the case when $0 < d < .5$. This model is often used for environmental time series arising in hydrology. We will discuss long memory processes in more detail in §5.2.




Fig. 2.6. Glacial varve thicknesses (top) from Massachusetts for $n = 634$ years compared with log transformed thicknesses (bottom).

Often, obvious aberrations are present that can contribute nonstationary as well as nonlinear behavior in observed time series. In such cases, transformations may be useful to equalize the variability over the length of a single series. A particularly useful transformation is

$$y_t = \log x_t, \quad (2.33)$$

which tends to suppress larger fluctuations that occur over portions of the series where the underlying values are larger. Other possibilities are power transformations in the Box–Cox family of the form

$$y_t = \begin{cases} (x_t^\lambda - 1)/\lambda & \lambda \neq 0, \\ \log x_t & \lambda = 0. \end{cases} \quad (2.34)$$

Methods for choosing the power λ are available (see Johnson and Wichern, 1992, §4.7) but we do not pursue them here. Often, transformations are also used to improve the approximation to normality or to improve linearity in predicting the value of one series from another.

Example 2.6 Paleoclimatic Glacial Varves

Melting glaciers deposit yearly layers of sand and silt during the spring melting seasons, which can be reconstructed yearly over a period ranging from the time

deglaciation began in New England (about 12,600 years ago) to the time it ended (about 6,000 years ago). Such sedimentary deposits, called varves, can be used as proxies for paleoclimatic parameters, such as temperature, because, in a warm year, more sand and silt are deposited from the receding glacier. **Figure 2.6** shows the thicknesses of the yearly varves collected from one location in Massachusetts for 634 years, beginning 11,834 years ago. For further information, see Shumway and Verosub (1992). Because the variation in thicknesses increases in proportion to the amount deposited, a logarithmic transformation could remove the nonstationarity observable in the variance as a function of time. **Figure 2.6** shows the original and transformed varves, and it is clear that this improvement has occurred. We may also plot the histogram of the original and transformed data, as in **Problem 2.8**, to argue that the approximation to normality is improved. The ordinary first differences (2.30) are also computed in **Problem 2.8**, and we note that the first differences have a significant negative correlation at lag $h = 1$. Later, in Chapter 5, we will show that perhaps the varve series has long memory and will propose using fractional differencing.

Figure 2.6 was generated in R as follows:

```
par(mfrow=c(2,1))
plot(varve, main="varve", ylab="")
plot(log(varve), main="log(varve)", ylab="")
```

Next, we consider another preliminary data processing technique that is used for the purpose of visualizing the relations between series at different lags, namely, scatterplot matrices. In the definition of the ACF, we are essentially interested in relations between x_t and x_{t-h} ; the autocorrelation function tells us whether a substantial linear relation exists between the series and its own lagged values. The ACF gives a profile of the linear correlation at all possible lags and shows which values of h lead to the best predictability. The restriction of this idea to linear predictability, however, may mask a possible nonlinear relation between current values, x_t , and past values, x_{t-h} . This idea extends to two series where one may be interested in examining scatterplots of y_t versus x_{t-h} .

Example 2.7 Scatterplot Matrices, SOI and Recruitment

To check for nonlinear relations of this form, it is convenient to display a lagged scatterplot matrix, as in **Figure 2.7**, that displays values of the SOI, S_t , on the vertical axis plotted against S_{t-h} on the horizontal axis. The sample autocorrelations are displayed in the upper right-hand corner and superimposed on the scatterplots are locally weighted scatterplot smoothing (lowess) lines that can be used to help discover any nonlinearities. We discuss smoothing in the next section, but for now, think of lowess as a robust method for fitting nonlinear regression.

In **Figure 2.7**, we notice that the lowess fits are approximately linear, so that the sample autocorrelations are meaningful. Also, we see strong positive linear relations at lags $h = 1, 2, 11, 12$, that is, between S_t and $S_{t-1}, S_{t-2}, S_{t-11}, S_{t-12}$, and a negative linear relation at lags $h = 6, 7$. These results match up well with peaks noticed in the ACF in **Figure 1.15**.




Fig. 2.7. Scatterplot matrix relating current SOI values, S_t , to past SOI values, S_{t-h} , at lags $h = 1, 2, \dots, 12$. The values in the upper right corner are the sample autocorrelations and the lines are a lowess fit.

Similarly, we might want to look at values of one series, say Recruitment, denoted R_t plotted against another series at various lags, say the SOI, S_{t-h} , to look for possible nonlinear relations between the two series. Because, for example, we might wish to predict the Recruitment series, R_t , from current or past values of the SOI series, S_{t-h} , for $h = 0, 1, 2, \dots$ it would be worthwhile to examine the scatterplot matrix. Figure 2.8 shows the lagged scatterplot of the Recruitment series R_t on the vertical axis plotted against the SOI index S_{t-h} on the horizontal axis. In addition, the figure exhibits the sample cross-correlations as well as lowess fits.

Figure 2.8 shows a fairly strong nonlinear relationship between Recruitment, R_t , and the SOI series at S_{t-5} , S_{t-6} , S_{t-7} , S_{t-8} , indicating the SOI series tends to lead the Recruitment series and the coefficients are negative, implying that increases in the SOI lead to decreases in the Recruitment. The nonlinearity observed in the scatterplots (with the help of the superimposed lowess fits) indicate that the behavior




Fig. 2.8. Scatterplot matrix of the Recruitment series, R_t , on the vertical axis plotted against the SOI series, S_{t-h} , on the horizontal axis at lags $h = 0, 1, \dots, 8$. The values in the upper right corner are the sample cross-correlations and the lines are a lowess fit.

between Recruitment and the SOI is different for positive values of SOI than for negative values of SOI.

Simple scatterplot matrices for one series can be obtained in R using the `lag.plot` command. Figure 2.7 and Figure 2.8 may be reproduced using the following scripts provided with `astsa` (see Appendix R for details):

```
lag1.plot(soi, 12)      # Figure 2.7
lag2.plot(soi, rec, 8) # Figure 2.8
```

Example 2.8 Regression with Lagged Variables (cont)

In Example 2.3 we regressed Recruitment on lagged SOI,

$$R_t = \beta_1 + \beta_2 S_{t-6} + w_t.$$

However, in Example 2.7, we saw that the relationship is nonlinear and different when SOI is positive or negative. In this case, we may consider adding a dummy




Fig. 2.9. Display for [Example 2.8](#): Plot of Recruitment (R_t) vs SOI lagged 6 months (S_{t-6}) with the fitted values of the regression as points (+) and a lowess fit (—).

variable to account for this change. In particular, we fit the model

$$R_t = \beta_1 + \beta_2 S_{t-6} + \beta_3 D_{t-6} + \beta_4 D_{t-6} S_{t-6} + w_t,$$

where D_t is a dummy variable that is 0 if $S_t < 0$ and 1 otherwise. The result of the fit is given in the R code below. [Figure 2.9](#) shows R_t vs S_{t-6} with the fitted values of the regression and a lowess fit superimposed. The piecewise regression fit is similar to the lowess fit, but we note that the residuals are not white noise (see the code below). This is followed up in [Example 3.49](#).

```
dummy = ifelse(soi<0, 0, 1)
fish = ts.intersect(rec, soil6=lag(soi,-6), dL6=lag(dummy,-6), dframe=TRUE)
summary(fit <- lm(rec~ soil6*dL6, data=fish, na.action=NULL))
Coefficients:
              Estimate Std. Error t.value    p.value
(Intercept)   74.479      2.865   25.998 < 2e-16
soil6         -15.358      7.401   -2.075  0.0386
dL6            -1.139      3.711   -0.307  0.7590
soil6:dL6     -51.244      9.523   -5.381  1.2e-07
---
Residual standard error: 21.84 on 443 degrees of freedom
Multiple R-squared:  0.4024
F-statistic: 99.43 on 3 and 443 DF,  p-value: < 2.2e-16
attach(fish)
plot(soil6, rec)
lines(lowess(soil6, rec), col=4, lwd=2)
points(soil6, fitted(fit), pch='+', col=2)
plot(resid(fit)) # not shown ...
acf(resid(fit))  # ... but obviously not noise
```

As a final exploratory tool, we discuss assessing periodic behavior in time series data using regression analysis and the periodogram; this material may be thought of as an introduction to spectral analysis, which we discuss in detail in Chapter 4. In [Example 1.12](#), we briefly discussed the problem of identifying cyclic or periodic signals in time series. A number of the time series we have seen so far exhibit periodic behavior. For example, the data from the pollution study example shown in [Figure 2.2](#) exhibit strong yearly cycles. Also, the Johnson & Johnson data shown in [Figure 1.1](#)




Fig. 2.10. Data generated by (2.35) [top] and the fitted line, (2.37), superimposed on the data [bottom].

make one cycle every year (four quarters) on top of an increasing trend and the speech data in Figure 1.2 is highly repetitive. The monthly SOI and Recruitment series in Figure 1.6 show strong yearly cycles, but hidden in the series are clues to the El Niño cycle.

Example 2.9 Using Regression to Discover a Signal in Noise

In Example 1.12, we generated $n = 500$ observations from the model

$$x_t = A \cos(2\pi\omega t + \phi) + w_t, \quad (2.35)$$

where $\omega = 1/50$, $A = 2$, $\phi = .6\pi$, and $\sigma_w = 5$; the data are shown on the bottom panel of Figure 1.11 on page 15. At this point we assume the frequency of oscillation $\omega = 1/50$ is known, but A and ϕ are unknown parameters. In this case the parameters appear in (2.35) in a nonlinear way, so we use a trigonometric identity⁴ and write

$$A \cos(2\pi\omega t + \phi) = \beta_1 \cos(2\pi\omega t) + \beta_2 \sin(2\pi\omega t),$$

where $\beta_1 = A \cos(\phi)$ and $\beta_2 = -A \sin(\phi)$. Now the model (2.35) can be written in the usual linear regression form given by (no intercept term is needed here)

$$x_t = \beta_1 \cos(2\pi t/50) + \beta_2 \sin(2\pi t/50) + w_t. \quad (2.36)$$

Using linear regression on the generated data, the fitted model is

$$\hat{x}_t = -.71_{(.30)} \cos(2\pi t/50) - 2.55_{(.30)} \sin(2\pi t/50) \quad (2.37)$$

⁴ $\cos(\alpha \pm \beta) = \cos(\alpha) \cos(\beta) \mp \sin(\alpha) \sin(\beta)$.

with $\hat{\sigma}_w = 4.68$, where the values in parentheses are the standard errors. We note the actual values of the coefficients for this example are $\beta_1 = 2 \cos(.6\pi) = -.62$ and $\beta_2 = -2 \sin(.6\pi) = -1.90$. Because the parameter estimates are significant and close to the actual values, it is clear that we are able to detect the signal in the noise using regression, even though the signal appears to be obscured by the noise in the bottom panel of [Figure 1.11](#). [Figure 2.10](#) shows data generated by (2.35) with the fitted line, (2.37), superimposed.

To reproduce the analysis and [Figure 2.10](#) in R, use the following:

```
set.seed(1000) # so you can reproduce these results
x = 2*cos(2*pi*1:500/50 + .6*pi) + rnorm(500,0,5)
z1 = cos(2*pi*1:500/50)
z2 = sin(2*pi*1:500/50)
summary(fit <- lm(x~0+z1+z2)) # zero to exclude the intercept
par(mfrow=c(2,1), mar=c(3,3,1,1), mgp=c(1.6,.6,0))
plot.ts(x)
plot.ts(x, col=8, ylab=expression(hat(x)))
lines(fitted(fit), col=2)
```

We will discuss this and related approaches in more detail in Chapter 4.

Example 2.10 The Periodogram – Discover a Signal in Noise

The analysis in [Example 2.9](#) may seem like cheating because we assumed we knew the value of the frequency parameter ω . If we do not know ω , we could try to fit the model (2.35) using nonlinear regression with ω as a parameter. Another method is to try various values of ω in a systematic way. Using the regression results of §2.2, we can show the estimated regression coefficients in [Example 2.9](#) take on the special form given by

$$\hat{\beta}_1 = \frac{\sum_{t=1}^n x_t \cos(2\pi t/50)}{\sum_{t=1}^n \cos^2(2\pi t/50)} = \frac{2}{n} \sum_{t=1}^n x_t \cos(2\pi t/50); \quad (2.38)$$

$$\hat{\beta}_2 = \frac{\sum_{t=1}^n x_t \sin(2\pi t/50)}{\sum_{t=1}^n \sin^2(2\pi t/50)} = \frac{2}{n} \sum_{t=1}^n x_t \sin(2\pi t/50). \quad (2.39)$$

This suggests looking at all possible regression parameter estimates,⁵ say

$$\hat{\beta}_1(j/n) = \frac{2}{n} \sum_{t=1}^n x_t \cos(2\pi t j/n); \quad (2.40)$$

$$\hat{\beta}_2(j/n) = \frac{2}{n} \sum_{t=1}^n x_t \sin(2\pi t j/n), \quad (2.41)$$

where, $n = 500$ and $j = 1, \dots, \frac{n}{2} - 1$, and inspecting the results for large values. For the endpoints, $j = 0$ and $j = n/2$, we have $\hat{\beta}_1(0) = n^{-1} \sum_{t=1}^n x_t$ and $\hat{\beta}_1(\frac{1}{2}) = n^{-1} \sum_{t=1}^n (-1)^t x_t$, and $\hat{\beta}_2(0) = \hat{\beta}_2(\frac{1}{2}) = 0$.

⁵ In the notation of §2.2, the estimates are of the form $\sum_{t=1}^n x_t z_t / \sum_{t=1}^n z_t^2$ where $z_t = \cos(2\pi t j/n)$ or $z_t = \sin(2\pi t j/n)$. In this setup, unless $j = 0$ or $j = n/2$ if n is even, $\sum_{t=1}^n z_t^2 = n/2$; see [Problem 2.10](#).




Fig. 2.11. The scaled periodogram, (2.42), of the 500 observations generated by (2.35); the data are displayed in Figure 1.11 and Figure 2.10.

For this particular example, the values calculated in (2.38) and (2.39) are $\hat{\beta}_1(10/500)$ and $\hat{\beta}_2(10/500)$. By doing this, we have regressed a series, x_t , of length n using n regression parameters, so that we will have a perfect fit. The point, however, is that if the data contain any cyclic behavior we are likely to catch it by performing these saturated regressions.

Next, note that the regression coefficients $\hat{\beta}_1(j/n)$ and $\hat{\beta}_2(j/n)$, for each j , are essentially measuring the correlation of the data with a sinusoid oscillating at j cycles in n time points.⁶ Hence, an appropriate measure of the presence of a frequency of oscillation of j cycles in n time points in the data would be

$$P(j/n) = \hat{\beta}_1^2(j/n) + \hat{\beta}_2^2(j/n), \quad (2.42)$$

which is basically a measure of squared correlation. The quantity (2.42) is sometimes called the periodogram, but we will call $P(j/n)$ the scaled periodogram and we will investigate its properties in Chapter 4. Figure 2.11 shows the scaled periodogram for the data generated by (2.35), and it easily discovers the periodic component with frequency $\omega = .02 = 10/500$ even though it is difficult to visually notice that component in Figure 1.11 due to the noise.

Finally, we mention that it is not necessary to run a large regression

$$x_t = \sum_{j=0}^{n/2} \beta_1(j/n) \cos(2\pi t j/n) + \beta_2(j/n) \sin(2\pi t j/n) \quad (2.43)$$

to obtain the values of $\beta_1(j/n)$ and $\beta_2(j/n)$ [with $\beta_2(0) = \beta_2(1/2) = 0$] because they can be computed quickly if n (assumed even here) is a highly composite integer. There is no error in (2.43) because there are n observations and n parameters; the regression fit will be perfect. The discrete Fourier transform (DFT) is a complex-valued weighted average of the data given by

⁶ Sample correlations are of the form $\sum_t x_t z_t / (\sum_t x_t^2 \sum_t z_t^2)^{1/2}$.

$$\begin{aligned}
d(j/n) &= n^{-1/2} \sum_{t=1}^n x_t \exp(-2\pi i t j/n) \\
&= n^{-1/2} \left(\sum_{t=1}^n x_t \cos(2\pi t j/n) - i \sum_{t=1}^n x_t \sin(2\pi t j/n) \right)
\end{aligned} \tag{2.44}$$

where the frequencies j/n are called the Fourier or fundamental frequencies. Because of a large number of redundancies in the calculation, (2.44) may be computed quickly using the fast Fourier transform (FFT)⁷, which is available in many computing packages such as Matlab®, S-PLUS® and R. Note that⁸

$$|d(j/n)|^2 = \frac{1}{n} \left(\sum_{t=1}^n x_t \cos(2\pi t j/n) \right)^2 + \frac{1}{n} \left(\sum_{t=1}^n x_t \sin(2\pi t j/n) \right)^2 \tag{2.45}$$

and it is this quantity that is called the periodogram; we will write

$$I(j/n) = |d(j/n)|^2.$$

We may calculate the scaled periodogram, (2.42), using the periodogram as

$$P(j/n) = \frac{4}{n} I(j/n). \tag{2.46}$$

We will discuss this approach in more detail and provide examples with data in Chapter 4.

Figure 2.11 can be created in R using the following commands (and the data already generated in `x`):

```

I = abs(fft(x))^2/500 # the periodogram
P = (4/500)*I[1:250] # the scaled periodogram
f = 0:249/500 # frequencies
plot(f, P, type="l", xlab="Frequency", ylab="Scaled Periodogram")

```

2.3 Smoothing in the Time Series Context

In §1.3, we introduced the concept of filtering or smoothing a time series, and in Example 1.9, we discussed using a moving average to smooth white noise. This method is useful in discovering certain traits in a time series, such as long-term trend and seasonal components. In particular, if x_t represents the observations, then

$$m_t = \sum_{j=-k}^k a_j x_{t-j}, \tag{2.47}$$

where $a_j = a_{-j} \geq 0$ and $\sum_{j=-k}^k a_j = 1$ is a symmetric moving average of the data.

⁷ Different packages scale the FFT differently; consult the documentation. R calculates (2.44) without scaling by $n^{-1/2}$.

⁸ If $z = a - ib$ is complex, then $|z|^2 = z\bar{z} = (a - ib)(a + ib) = a^2 + b^2$.




Fig. 2.12. Moving average smoother of SOI. The insert shows the shape of the moving average (“boxcar”) kernel [not drawn to scale] described in (2.49).

Example 2.11 Moving Average Smoother

For example, Figure 2.12 shows the monthly SOI series discussed in Example 1.5 smoothed using (2.47) with weights $a_0 = a_{\pm 1} = \dots = a_{\pm 5} = 1/12$, and $a_{\pm 6} = 1/24$; $k = 6$. This particular method removes (filters out) the obvious annual temperature cycle and helps emphasize the El Niño cycle.

To reproduce Figure 2.12 in R:

```
wgts = c(.5, rep(1,11), .5)/12
soif = filter(soi, sides=2, filter=wgts)
plot(soi)
lines(soif, lwd=2, col=4)
```

Although the moving average smoother does a good job in highlighting the El Niño effect, it might be considered too choppy. We can obtain a smoother fit using the normal distribution for the weights, instead of boxcar-type weights of (2.47).

Example 2.12 Kernel Smoothing

Kernel smoothing is a moving average smoother that uses a weight function, or kernel, to average the observations. Figure 2.13 shows kernel smoothing of the mortality series, where m_t is now

$$m_t = \sum_{i=1}^n w_i(t) x_i, \quad (2.48)$$

where

$$w_i(t) = K\left(\frac{t-i}{b}\right) / \sum_{j=1}^n K\left(\frac{t-j}{b}\right) \quad (2.49)$$

are the weights and $K(\cdot)$ is a kernel function. This estimator, which was originally explored by Parzen (1962) and Rosenblatt (1956b), is often called the Nadaraya–Watson estimator (Watson, 1966). In this example, and typically, the normal kernel, $K(z) = \frac{1}{\sqrt{2\pi}} \exp(-z^2/2)$, is used.

To implement this in R, use the `ksmooth` function where a bandwidth can be chosen. The wider the bandwidth, b , the smoother the result. From the R `ksmooth`




Fig. 2.13. Kernel smoother of SOI. The insert shows the shape of the normal kernel [not drawn to scale].

help file: The kernels are scaled so that their quartiles (viewed as probability densities) are at $\pm 0.25 \cdot \text{bandwidth}$. For the standard normal distribution, the quartiles are ± 0.674 . In our case, we are smoothing over time, which is of the form $t/12$ for the SOI time series. In **Figure 2.13**, we used the value of $b = 1$ to correspond to approximately smoothing over one year.

Figure 2.13 can be reproduced in R as follows.

```
plot(soi)
lines(ksmooth(time(soi), soi, "normal", bandwidth=1), lwd=2, col=4)
```

Example 2.13 Lowess

Another approach to smoothing a time plot is nearest neighbor regression. The technique is based on k -nearest neighbors regression, wherein one uses only the data $\{x_{t-k/2}, \dots, x_t, \dots, x_{t+k/2}\}$ to predict x_t via regression, and then sets $m_t = \hat{x}_t$.

Lowess is a method of smoothing that is rather complex, but the basic idea is close to nearest neighbor regression. **Figure 2.14** shows smoothing of mortality using the R function `lowess` (see Cleveland, 1979). First, a certain proportion of nearest neighbors to x_t are included in a weighting scheme; values closer to x_t in time get more weight. Then, a robust weighted regression is used to predict x_t and obtain the smoothed values m_t . The larger the fraction of nearest neighbors included, the smoother the fit will be. In **Figure 2.14**, one smoother uses 5% of the data to obtain an estimate of the El Niño cycle of the data.

In addition, a (negative) trend in SOI would indicate the long-term warming of the Pacific Ocean. To investigate this, we used lowess with the default smoother span of $f=2/3$ of the data. **Figure 2.14** can be reproduced in R as follows.

```
plot(soi)
lines(lowess(soi, f=.05), lwd=2, col=4) # El Nino cycle
lines(lowess(soi), lty=2, lwd=2, col=2) # trend (with default span)
```





Fig. 2.14. Locally weighted scatterplot smoothers (**lowess**) of the SOI series.

Example 2.14 Smoothing Splines

An obvious way to smooth data would be to fit a polynomial regression in terms of time. For example, a cubic polynomial would have $x_t = m_t + w_t$ where

$$m_t = \beta_0 + \beta_1 t + \beta_2 t^2 + \beta_3 t^3.$$

We could then fit m_t via ordinary least squares.

An extension of polynomial regression is to first divide time $t = 1, \dots, n$, into k intervals, $[t_0 = 1, t_1]$, $[t_1 + 1, t_2]$, $\dots$, $[t_{k-1} + 1, t_k = n]$; the values $t_0, t_1, \dots, t_k$ are called *knots*. Then, in each interval, one fits a polynomial regression, typically the order is 3, and this is called cubic splines.

A related method is smoothing splines, which minimizes a compromise between the fit and the degree of smoothness given by

$$\sum_{t=1}^n [x_t - m_t]^2 + \lambda \int (m_t'')^2 dt, \quad (2.50)$$

where m_t is a cubic spline with a knot at each t and primes denote differentiation. The degree of smoothness is controlled by $\lambda > 0$.

Think of taking a long drive where m_t is the position of your car at time t . In this case, m_t'' is acceleration/deceleration, and $\int (m_t'')^2$ is a measure of the total amount of acceleration and deceleration on your trip. A smooth drive would be one where a constant velocity, v , is maintained (i.e., $m_t' = v$). A choppy ride would be when the driver is constantly accelerating and decelerating, such as beginning drivers tend to do.

If $\lambda = 0$, we don't care how choppy the ride is, and this leads to $m_t = x_t$, which is the data and not smooth. If $\lambda = \infty$, we insist on no acceleration or deceleration (i.e., $m_t'' = 0$); in this case, our drive must be at constant velocity, $m_t = c + vt$, and consequently very smooth. Thus, λ is seen as a tradeoff between linear regression (completely smooth) and the data itself (no smoothness). The larger the value of λ , the smoother the fit.




Fig. 2.15. Smoothing splines fit to the SOI series.

In R, the smoothing parameter is called `spar` and it is monotonically related to λ ; type `?smooth.spline` to view the help file for details. Figure 2.15 shows smoothing spline fits on the SOI series using `spar=.5` to emphasize the El Niño cycle, and `spar=1` to emphasize the trend. The figure can be reproduced in R as follows.

```
plot(soi)
lines(smooth.spline(time(soi), soi, spar=.5), lwd=2, col=4)
lines(smooth.spline(time(soi), soi, spar=1), lty=2, lwd=2, col=2)
```

Example 2.15 Smoothing One Series as a Function of Another

In addition to smoothing time plots, smoothing techniques can be applied to smoothing a time series as a function of another time series. We have already seen this idea used in Example 2.7 when we used lowess to visualize the nonlinear relationship between Recruitment and SOI at various lags. In this example, we smooth the scatterplot of two contemporaneously measured time series, mortality as a function of temperature. In Example 2.2, we discovered a nonlinear relationship between mortality and temperature. Continuing along these lines, Figure 2.16 shows scatterplots of mortality, M_t , and temperature, T_t , along with M_t smoothed as a function of T_t using lowess. Note that mortality increases at extreme temperatures, but in an asymmetric way; mortality is higher at colder temperatures than at hotter temperatures. The minimum mortality rate seems to occur at approximately 80° F.

Figure 2.16 can be reproduced in R as follows using the defaults.

```
plot(tempr, cmort, xlab="Temperature", ylab="Mortality")
lines(lowess(tempr, cmort))
```




Fig. 2.16. Smooth of mortality as a function of temperature using lowess.

Problems

Section 2.2

2.1 Structural Model For the Johnson & Johnson data, say y_t , shown in Figure 1.1, let $x_t = \log(y_t)$. In this problem, we are going to fit a special type of structural model, $x_t = T_t + S_t + N_t$ where T_t is a trend component, S_t is a seasonal component, and N_t is noise. In our case, time t is in quarters (1960.00, 1960.25, ...) so one unit of time is a year.

(a) Fit the regression model

$$x_t = \underbrace{\beta t}_{\text{trend}} + \underbrace{\alpha_1 Q_1(t) + \alpha_2 Q_2(t) + \alpha_3 Q_3(t) + \alpha_4 Q_4(t)}_{\text{seasonal}} + \underbrace{w_t}_{\text{noise}}$$

where $Q_i(t) = 1$ if time t corresponds to quarter $i = 1, 2, 3, 4$, and zero otherwise. The $Q_i(t)$'s are called indicator variables. We will assume for now that w_t is a Gaussian white noise sequence. *Hint:* Detailed code is given in Appendix R on page 531.

- (b) If the model is correct, what is the estimated average annual increase in the logged earnings per share?
- (c) If the model is correct, does the average logged earnings rate increase or decrease from the third quarter to the fourth quarter? And, by what percentage does it increase or decrease?
- (d) What happens if you include an intercept term in the model in (a)? Explain why there was a problem.
- (e) Graph the data, x_t , and superimpose the fitted values, say $\hat{x}_t$, on the graph. Examine the residuals, $x_t - \hat{x}_t$, and state your conclusions. Does it appear that the model fits the data well (do the residuals look white)?

2.2 For the mortality data examined in Example 2.2:

- (a) Add another component to the regression in (2.21) that accounts for the particulate count four weeks prior; that is, add P_{t-4} to the regression in (2.21). State your conclusion.
- (b) Draw a scatterplot matrix of M_t, T_t, P_t and P_{t-4} and then calculate the pairwise correlations between the series. Compare the relationship between M_t and P_t versus M_t and P_{t-4} .

2.3 Read the entire problem before you start because you can do parts (a) and (b) at the same time.

- (a) Generate *four* series that are random walk with drift, (1.4), of length $n = 100$ with $\delta = .01$ and $\sigma_w = 1$. Call the data x_t for $t = 1, \dots, 100$. Fit the regression $x_t = \beta t + w_t$ using least squares. Plot the data, the true mean function (i.e., $\mu_t = .01 t$) and the fitted line, $\hat{x}_t = \hat{\beta} t$, on the same graph.
- (b) Generate *four* series of length $n = 100$ that are linear trend plus noise, say $y_t = .01 t + w_t$, where t and w_t are as in part (a). Fit the regression $y_t = \beta t + w_t$ using least squares. Plot the data, the true mean function (i.e., $\mu_t = .01 t$) and the fitted line, $\hat{y}_t = \hat{\beta} t$, on the same graph.
- (c) Comment (what did you learn from this assignment).

The following R code may be useful for doing (a) and (b) at the same time.

```
par(mfrow=c(4,2), mar=c(2.5,2.5,0,0)+.5, mgp=c(1.6,.6,0)) # set up
for (i in 1:4){
  x = ts(cumsum(rnorm(100,.01,1)))          # data
  y = ts(.01*1:100 + rnorm(100))
  regx = lm(x~0+time(x), na.action=NULL)    # regressions
  regy = lm(y~0+time(y), na.action=NULL)
  plot(x)                                    # plots
  abline(a=0, b=.01, col=2, lty=2)          # true mean
  abline(regx, col=4)                       # fitted line
  plot(y)
  abline(a=0, b=.01, col=2, lty=2)
  abline(regy, col=4)
}
```

2.4 Kullback-Leibler Information Given the random vector $\mathbf{y}$, we define the information for discriminating between two densities in the same family, indexed by a parameter $\boldsymbol{\theta}$, say $f(\mathbf{y}; \boldsymbol{\theta}_1)$ and $f(\mathbf{y}; \boldsymbol{\theta}_2)$, as

$$I(\boldsymbol{\theta}_1; \boldsymbol{\theta}_2) = \frac{1}{n} E_1 \log \frac{f(\mathbf{y}; \boldsymbol{\theta}_1)}{f(\mathbf{y}; \boldsymbol{\theta}_2)}, \quad (2.51)$$

where E_1 denotes expectation with respect to the density determined by $\boldsymbol{\theta}_1$. For the Gaussian regression model, the parameters are $\boldsymbol{\theta} = (\boldsymbol{\beta}', \sigma^2)'$. Show that we obtain

$$I(\boldsymbol{\theta}_1; \boldsymbol{\theta}_2) = \frac{1}{2} \left(\frac{\sigma_1^2}{\sigma_2^2} - \log \frac{\sigma_1^2}{\sigma_2^2} - 1 \right) + \frac{1}{2} \frac{(\boldsymbol{\beta}_1 - \boldsymbol{\beta}_2)' Z' Z (\boldsymbol{\beta}_1 - \boldsymbol{\beta}_2)}{n \sigma_2^2} \quad (2.52)$$

in that case.

2.5 Model Selection Both selection criteria (2.15) and (2.16) are derived from information theoretic arguments, based on the well-known Kullback-Leibler discrimination information numbers (see Kullback and Leibler, 1951, Kullback, 1958). We give an argument due to Hurvich and Tsai (1989). We think of the measure (2.52) as measuring the discrepancy between the two densities, characterized by the parameter values $\theta'_1 = (\beta'_1, \sigma_1^2)'$ and $\theta'_2 = (\beta'_2, \sigma_2^2)'$. Now, if the true value of the parameter vector is θ_1 , we argue that the best model would be one that minimizes the discrepancy between the theoretical value and the sample, say $I(\theta_1; \hat{\theta})$. Because θ_1 will not be known, Hurvich and Tsai (1989) considered finding an unbiased estimator for $E_1[I(\beta_1, \sigma_1^2; \hat{\beta}, \hat{\sigma}^2)]$, where

$$I(\beta_1, \sigma_1^2; \hat{\beta}, \hat{\sigma}^2) = \frac{1}{2} \left(\frac{\sigma_1^2}{\hat{\sigma}^2} - \log \frac{\sigma_1^2}{\hat{\sigma}^2} - 1 \right) + \frac{1}{2} \frac{(\beta_1 - \hat{\beta})' Z' Z (\beta_1 - \hat{\beta})}{n \hat{\sigma}^2}$$

and β is a $k \times 1$ regression vector. Show that

$$E_1[I(\beta_1, \sigma_1^2; \hat{\beta}, \hat{\sigma}^2)] = \frac{1}{2} \left(-\log \sigma_1^2 + E_1 \log \hat{\sigma}^2 + \frac{n+k}{n-k-2} - 1 \right), \quad (2.53)$$

using the distributional properties of the regression coefficients and error variance. An unbiased estimator for $E_1 \log \hat{\sigma}^2$ is $\log \hat{\sigma}^2$. Hence, we have shown that the expectation of the above discrimination information is as claimed. As models with differing dimensions k are considered, only the second and third terms in (2.53) will vary and we only need unbiased estimators for those two terms. This gives the form of AICc quoted in (2.16) in the chapter. You will need the two distributional results

$$\frac{n \hat{\sigma}^2}{\sigma_1^2} \sim \chi_{n-k}^2 \quad \text{and} \quad \frac{(\hat{\beta} - \beta_1)' Z' Z (\hat{\beta} - \beta_1)}{\sigma_1^2} \sim \chi_k^2$$

The two quantities are distributed independently as chi-squared distributions with the indicated degrees of freedom. If $x \sim \chi_n^2$, $E(1/x) = 1/(n-2)$.

Section 2.3

2.6 Consider a process consisting of a linear trend with an additive noise term consisting of independent random variables w_t with zero means and variances σ_w^2 , that is,

$$x_t = \beta_0 + \beta_1 t + w_t,$$

where β_0, β_1 are fixed constants.

- Prove x_t is nonstationary.
- Prove that the first difference series $\nabla x_t = x_t - x_{t-1}$ is stationary by finding its mean and autocovariance function.
- Repeat part (b) if w_t is replaced by a general stationary process, say y_t , with mean function μ_y and autocovariance function $\gamma_y(h)$.

2.7 Show (2.27) is stationary.

2.8 The glacial varve record plotted in Figure 2.6 exhibits some nonstationarity that can be improved by transforming to logarithms and some additional nonstationarity that can be corrected by differencing the logarithms.

- Argue that the glacial varves series, say x_t , exhibits heteroscedasticity by computing the sample variance over the first half and the second half of the data. Argue that the transformation $y_t = \log x_t$ stabilizes the variance over the series. Plot the histograms of x_t and y_t to see whether the approximation to normality is improved by transforming the data.
- Plot the series y_t . Do any time intervals, of the order 100 years, exist where one can observe behavior comparable to that observed in the global temperature records in Figure 1.2?
- Examine the sample ACF of y_t and comment.
- Compute the difference $u_t = y_t - y_{t-1}$, examine its time plot and sample ACF, and argue that differencing the logged varve data produces a reasonably stationary series. Can you think of a practical interpretation for u_t ? Hint: For $|p|$ close to zero, $\log(1 + p) \approx p$; let $p = (y_t - y_{t-1})/y_{t-1}$.
- Based on the sample ACF of the differenced series computed in (c), argue that a generalization of the model given by Example 1.24 might be reasonable. Assume

$$u_t = \mu + w_t + \theta w_{t-1}$$

is stationary when the inputs w_t are assumed independent with mean 0 and variance σ_w^2 . Show that

$$\gamma_u(h) = \begin{cases} \sigma_w^2(1 + \theta^2) & \text{if } h = 0, \\ \theta \sigma_w^2 & \text{if } h = \pm 1, \\ 0 & \text{if } |h| > 1. \end{cases}$$

- Based on part (e), use $\hat{\rho}_u(1)$ and the estimate of the variance of u_t , $\hat{\gamma}_u(0)$, to derive estimates of θ and σ_w^2 . This is an application of the method of moments from classical statistics, where estimators of the parameters are derived by equating sample moments to theoretical moments.

2.9 In this problem, we will explore the periodic nature of S_t , the SOI series displayed in Figure 1.5.

- Detrend the series by fitting a regression of S_t on time t . Is there a significant trend in the sea surface temperature? Comment.
- Calculate the periodogram for the detrended series obtained in part (a). Identify the frequencies of the two main peaks (with an obvious one at the frequency of one cycle every 12 months). What is the probable El Niño cycle indicated by the minor peak?

2.10 Consider the model (2.43) used in Example 2.10,

$$x_t = \sum_{j=0}^{n/2} \beta_1(j/n) \cos(2\pi t j/n) + \beta_2(j/n) \sin(2\pi t j/n).$$

Verify that for any positive integer n and $j, k = 0, 1, \dots, \lfloor n/2 \rfloor$, where $\lfloor \cdot \rfloor$ denotes the greatest integer function:⁹

(a) Except for $j = 0$ or $j = n/2$,

$$\sum_{t=1}^n \cos^2(2\pi t j/n) = \sum_{t=1}^n \sin^2(2\pi t j/n) = n/2$$

(b) When $j = 0$ or $j = n/2$,

$$\sum_{t=1}^n \cos^2(2\pi t j/n) = n \quad \text{but} \quad \sum_{t=1}^n \sin^2(2\pi t j/n) = 0.$$

(c) For $j \neq k$,

$$\sum_{t=1}^n \cos(2\pi t j/n) \cos(2\pi t k/n) = \sum_{t=1}^n \sin(2\pi t j/n) \sin(2\pi t k/n) = 0.$$

Also, for any j and k ,

$$\sum_{t=1}^n \cos(2\pi t j/n) \sin(2\pi t k/n) = 0.$$

Section 2.4

2.11 Consider the two weekly time series [oil](#) and [gas](#). The oil series is in dollars per barrel, while the gas series is in cents per gallon; see Appendix R for details.

- Plot the data on the same graph. Which of the simulated series displayed in §1.3 do these series most resemble? Do you believe the series are stationary (explain your answer)?
- In economics, it is often the percentage change in price (termed *growth rate* or *return*), rather than the absolute price change, that is important. Argue that a transformation of the form $y_t = \nabla \log x_t$ might be applied to the data, where x_t is the oil or gas price series [see the hint in Problem 2.8(d)].
- Transform the data as described in part (b), plot the data on the same graph, look at the sample ACFs of the transformed data, and comment.

⁹ We'll do part of (a):

$$\begin{aligned} \sum_{t=1}^n \cos^2(2\pi t j/n) &= \frac{1}{4} \sum_{t=1}^n (e^{2\pi i t j/n} + e^{-2\pi i t j/n})(e^{2\pi i t j/n} + e^{-2\pi i t j/n}) \\ &= \frac{1}{4} \sum_{t=1}^n (e^{4\pi i t j/n} + 1 + 1 + e^{-4\pi i t j/n}) = \frac{n}{2}. \end{aligned}$$

- (d) Plot the CCF of the transformed data and comment. The small, but significant values when `gas` leads `oil` might be considered as feedback.
- (e) Exhibit scatterplots of the oil and gas growth rate series for up to three weeks of lead time of oil prices; include a nonparametric smoother in each plot and comment on the results (e.g., Are there outliers? Are the relationships linear?).
- (f) There have been a number of studies questioning whether gasoline prices respond more quickly when oil prices are rising than when oil prices are falling (“asymmetry”). We will attempt to explore this question here with simple lagged regression; we will ignore some obvious problems such as outliers and autocorrelated errors, so this will not be a definitive analysis. Let G_t and O_t denote the gas and oil growth rates.
 - (i) Fit the regression (and comment on the results)

$$G_t = \alpha_1 + \alpha_2 I_t + \beta_1 O_t + \beta_2 O_{t-1} + w_t,$$

where $I_t = 1$ if $O_t \geq 0$ and 0 otherwise (I_t is the indicator of no growth or positive growth in oil price). Hint:

```
poil = diff(log(oil))
pgas = diff(log(gas))
indi = ifelse(poil < 0, 0, 1)
mess = ts.intersect(pgas, poil, poilL = lag(poil,-1), indi)
summary(fit <- lm(pgas~ poil + poilL + indi, data=mess))
```

- (ii) What is the fitted model when there is negative growth in oil price at time t ? What is the fitted model when there is no or positive growth in oil price? Do these results support the asymmetry hypothesis?
- (iii) Analyze the residuals from the fit and comment.

2.12 Use two different smoothing techniques described in §2.3 to estimate the trend in the global temperature series `gtemp`. Comment.

Chapter 3

ARIMA Models

3.1 Introduction

In Chapters 1 and 2, we introduced autocorrelation and cross-correlation functions (ACFs and CCFs) as tools for clarifying relations that may occur within and between time series at various lags. In addition, we explained how to build linear models based on classical regression theory for exploiting the associations indicated by large values of the ACF or CCF. The time domain, or regression, methods of this chapter are appropriate when we are dealing with possibly nonstationary, shorter time series; these series are the rule rather than the exception in many applications. In addition, if the emphasis is on forecasting future values, then the problem is easily treated as a regression problem. This chapter develops a number of regression techniques for time series that are all related to classical ordinary and weighted or correlated least squares.

Classical regression is often insufficient for explaining all of the interesting dynamics of a time series. For example, the ACF of the residuals of the simple linear regression fit to the global temperature data (see [Example 2.4](#) of Chapter 2) reveals additional structure in the data that the regression did not capture. Instead, the introduction of correlation as a phenomenon that may be generated through lagged linear relations leads to proposing the autoregressive (AR) and autoregressive moving average (ARMA) models. Adding nonstationary models to the mix leads to the autoregressive integrated moving average (ARIMA) model popularized in the landmark work by Box and Jenkins (1970). The Box–Jenkins method for identifying a plausible ARIMA model is given in this chapter along with techniques for parameter estimation and forecasting for these models. A partial theoretical justification of the use of ARMA models is discussed in [Appendix B, §B.4](#).

3.2 Autoregressive Moving Average Models

The classical regression model of Chapter 2 was developed for the static case, namely, we only allow the dependent variable to be influenced by current values of the

independent variables. In the time series case, it is desirable to allow the dependent variable to be influenced by the past values of the independent variables and possibly by its own past values. If the present can be plausibly modeled in terms of only the past values of the independent inputs, we have the enticing prospect that forecasting will be possible.

INTRODUCTION TO AUTOREGRESSIVE MODELS

Autoregressive models are based on the idea that the current value of the series, x_t , can be explained as a function of p past values, $x_{t-1}, x_{t-2}, \dots, x_{t-p}$, where p determines the number of steps into the past needed to forecast the current value. As a typical case, recall [Example 1.10](#) in which data were generated using the model

$$x_t = x_{t-1} - .90x_{t-2} + w_t,$$

where w_t is white Gaussian noise with $\sigma_w^2 = 1$. We have now assumed the current value is a particular *linear* function of past values. The regularity that persists in [Figure 1.9](#) gives an indication that forecasting for such a model might be a distinct possibility, say, through some version such as

$$x_{n+1}^n = x_n - .90x_{n-1},$$

where the quantity on the left-hand side denotes the forecast at the next period $n + 1$ based on the observed data, $x_1, x_2, \dots, x_n$. We will make this notion more precise in our discussion of forecasting (§3.5).

The extent to which it might be possible to forecast a real data series from its own past values can be assessed by looking at the autocorrelation function and the lagged scatterplot matrices discussed in Chapter 2. For example, the lagged scatterplot matrix for the Southern Oscillation Index (SOI), shown in [Figure 2.7](#), gives a distinct indication that lags 1 and 2, for example, are linearly associated with the current value. The ACF shown in [Figure 1.15](#) shows relatively large positive values at lags 1, 2, 12, 24, and 36 and large negative values at 18, 30, and 42. We note also the possible relation between the SOI and Recruitment series indicated in the scatterplot matrix shown in [Figure 2.8](#). We will indicate in later sections on transfer function and vector AR modeling how to handle the dependence on values taken by other series.

The preceding discussion motivates the following definition.

Definition 3.1 An **autoregressive model** of order p , abbreviated **AR**(p), is of the form

$$x_t = \phi_1 x_{t-1} + \phi_2 x_{t-2} + \dots + \phi_p x_{t-p} + w_t, \quad (3.1)$$

where x_t is stationary, $w_t \sim wn(0, \sigma_w^2)$, and $\phi_1, \phi_2, \dots, \phi_p$ are constants ($\phi_p \neq 0$). The mean of x_t in (3.1) is zero. If the mean, μ , of x_t is not zero, replace x_t by $x_t - \mu$ in (3.1),

$$x_t - \mu = \phi_1 (x_{t-1} - \mu) + \phi_2 (x_{t-2} - \mu) + \dots + \phi_p (x_{t-p} - \mu) + w_t,$$

or write

$$x_t = \alpha + \phi_1 x_{t-1} + \phi_2 x_{t-2} + \dots + \phi_p x_{t-p} + w_t, \quad (3.2)$$

where $\alpha = \mu(1 - \phi_1 - \dots - \phi_p)$.

We note that (3.2) is similar to the regression model of §2.2, and hence the term auto (or self) regression. Some technical difficulties, however, develop from applying that model because the regressors, $x_{t-1}, \dots, x_{t-p}$, are random components, whereas z_t was assumed to be fixed. A useful form follows by using the backshift operator (2.29) to write the AR(p) model, (3.1), as

$$(1 - \phi_1 B - \phi_2 B^2 - \dots - \phi_p B^p)x_t = w_t, \quad (3.3)$$

or even more concisely as

$$\phi(B)x_t = w_t. \quad (3.4)$$

The properties of $\phi(B)$ are important in solving (3.4) for x_t . This leads to the following definition.

Definition 3.2 *The autoregressive operator is defined to be*

$$\phi(B) = 1 - \phi_1 B - \phi_2 B^2 - \dots - \phi_p B^p. \quad (3.5)$$

We initiate the investigation of AR models by considering the first-order model, AR(1), given by $x_t = \phi x_{t-1} + w_t$. Iterating backwards k times, we get

$$\begin{aligned} x_t &= \phi x_{t-1} + w_t = \phi(\phi x_{t-2} + w_{t-1}) + w_t \\ &= \phi^2 x_{t-2} + \phi w_{t-1} + w_t \\ &\vdots \\ &= \phi^k x_{t-k} + \sum_{j=0}^{k-1} \phi^j w_{t-j}. \end{aligned}$$

This method suggests that, by continuing to iterate backward, and provided that $|\phi| < 1$ and $\sup_t \text{var}(x_t) < \infty$, we can represent an AR(1) model as a linear process given by¹

$$x_t = \sum_{j=0}^{\infty} \phi^j w_{t-j}. \quad (3.6)$$

The AR(1) process defined by (3.6) is stationary with mean

$$E(x_t) = \sum_{j=0}^{\infty} \phi^j E(w_{t-j}) = 0,$$

and autocovariance function,

$$\begin{aligned} \gamma(h) &= \text{cov}(x_{t+h}, x_t) = E \left[\left(\sum_{j=0}^{\infty} \phi^j w_{t+h-j} \right) \left(\sum_{k=0}^{\infty} \phi^k w_{t-k} \right) \right] \\ &= E \left[\left(w_{t+h} + \dots + \phi^h w_t + \phi^{h+1} w_{t-1} + \dots \right) (w_t + \phi w_{t-1} + \dots) \right] \\ &= \sigma_w^2 \sum_{j=0}^{\infty} \phi^{h+j} \phi^j = \sigma_w^2 \phi^h \sum_{j=0}^{\infty} \phi^{2j} = \frac{\sigma_w^2 \phi^h}{1 - \phi^2}, \quad h \geq 0. \end{aligned} \quad (3.7)$$

¹ Note that $\lim_{k \rightarrow \infty} E \left(x_t - \sum_{j=0}^{k-1} \phi^j w_{t-j} \right)^2 = \lim_{k \rightarrow \infty} \phi^{2k} E \left(x_{t-k}^2 \right) = 0$, so (3.6) exists in the mean square sense (see Appendix A for a definition).

Recall that $\gamma(h) = \gamma(-h)$, so we will only exhibit the autocovariance function for $h \geq 0$. From (3.7), the ACF of an AR(1) is

$$\rho(h) = \frac{\gamma(h)}{\gamma(0)} = \phi^h, \quad h \geq 0, \quad (3.8)$$

and $\rho(h)$ satisfies the recursion

$$\rho(h) = \phi \rho(h-1), \quad h = 1, 2, \dots \quad (3.9)$$

We will discuss the ACF of a general AR(p) model in §3.4.

Example 3.1 The Sample Path of an AR(1) Process

Figure 3.1 shows a time plot of two AR(1) processes, one with $\phi = .9$ and one with $\phi = -.9$; in both cases, $\sigma_w^2 = 1$. In the first case, $\rho(h) = .9^h$, for $h \geq 0$, so observations close together in time are positively correlated with each other. This result means that observations at contiguous time points will tend to be close in value to each other; this fact shows up in the top of Figure 3.1 as a very smooth sample path for x_t . Now, contrast this with the case in which $\phi = -.9$, so that $\rho(h) = (-.9)^h$, for $h \geq 0$. This result means that observations at contiguous time points are negatively correlated but observations two time points apart are positively correlated. This fact shows up in the bottom of Figure 3.1, where, for example, if an observation, x_t , is positive, the next observation, x_{t+1} , is typically negative, and the next observation, x_{t+2} , is typically positive. Thus, in this case, the sample path is very choppy.

The following R code can be used to obtain a figure similar to Figure 3.1:

```
par(mfrow=c(2,1))
plot(arima.sim(list(order=c(1,0,0), ar=.9), n=100), ylab="x",
     main=(expression(AR(1)~~~phi==+.9)))
plot(arima.sim(list(order=c(1,0,0), ar=-.9), n=100), ylab="x",
     main=(expression(AR(1)~~~phi==-.9)))
```

Example 3.2 Explosive AR Models and Causality

In Example 1.18, it was discovered that the random walk $x_t = x_{t-1} + w_t$ is not stationary. We might wonder whether there is a stationary AR(1) process with $|\phi| > 1$. Such processes are called explosive because the values of the time series quickly become large in magnitude. Clearly, because $|\phi|^j$ increases without bound as $j \rightarrow \infty$, $\sum_{j=0}^{k-1} \phi^j w_{t-j}$ will not converge (in mean square) as $k \rightarrow \infty$, so the intuition used to get (3.6) will not work directly. We can, however, modify that argument to obtain a stationary model as follows. Write $x_{t+1} = \phi x_t + w_{t+1}$, in which case,

$$\begin{aligned} x_t &= \phi^{-1} x_{t+1} - \phi^{-1} w_{t+1} = \phi^{-1} (\phi^{-1} x_{t+2} - \phi^{-1} w_{t+2}) - \phi^{-1} w_{t+1} \\ &\vdots \\ &= \phi^{-k} x_{t+k} - \sum_{j=1}^{k-1} \phi^{-j} w_{t+j}, \end{aligned} \quad (3.10)$$




Fig. 3.1. Simulated AR(1) models: $\phi = .9$ (top); $\phi = -.9$ (bottom).

by iterating forward k steps. Because $|\phi|^{-1} < 1$, this result suggests the stationary future dependent AR(1) model

$$x_t = - \sum_{j=1}^{\infty} \phi^{-j} w_{t+j}. \quad (3.11)$$

The reader can verify that this is stationary and of the AR(1) form $x_t = \phi x_{t-1} + w_t$. Unfortunately, this model is useless because it requires us to know the future to be able to predict the future. When a process does not depend on the future, such as the AR(1) when $|\phi| < 1$, we will say the process is causal. In the explosive case of this example, the process is stationary, but it is also future dependent, and not causal.

Example 3.3 Every Explosion Has a Cause

Excluding explosive models from consideration is not a problem because the models have causal counterparts. For example, if

$$x_t = \phi x_{t-1} + w_t \quad \text{with} \quad |\phi| > 1$$

and $w_t \sim \text{iid } N(0, \sigma_w^2)$, then using (3.11), $\{x_t\}$ is a non-causal stationary Gaussian process with $E(x_t) = 0$ and

$$\begin{aligned}\gamma_x(h) &= \text{cov}(x_{t+h}, x_t) = \text{cov}\left(-\sum_{j=1}^{\infty} \phi^{-j} w_{t+h+j}, -\sum_{k=1}^{\infty} \phi^{-k} w_{t+k}\right) \\ &= \sigma_w^2 \phi^{-2} \phi^{-h} / (1 - \phi^{-2}).\end{aligned}$$

Thus, using (3.7), the causal process defined by

$$y_t = \phi^{-1} y_{t-1} + v_t$$

where $v_t \sim \text{iid } N(0, \sigma_w^2 \phi^{-2})$ is stochastically equal to the x_t process (i.e., all finite distributions of the processes are the same). For example, if $x_t = 2x_{t-1} + w_t$ with $\sigma_w^2 = 1$, then $y_t = \frac{1}{2}y_{t-1} + v_t$ with $\sigma_v^2 = 1/4$ is an equivalent causal process (see Problem 3.3). This concept generalizes to higher orders, but it is easier to show using Chapter 4 techniques; see Example 4.8.

The technique of iterating backward to get an idea of the stationary solution of AR models works well when $p = 1$, but not for larger orders. A general technique is that of matching coefficients. Consider the AR(1) model in operator form

$$\phi(B)x_t = w_t, \quad (3.12)$$

where $\phi(B) = 1 - \phi B$, and $|\phi| < 1$. Also, write the model in equation (3.6) using operator form as

$$x_t = \sum_{j=0}^{\infty} \psi_j w_{t-j} = \psi(B)w_t, \quad (3.13)$$

where $\psi(B) = \sum_{j=0}^{\infty} \psi_j B^j$ and $\psi_j = \phi^j$. Suppose we did not know that $\psi_j = \phi^j$. We could substitute $\psi(B)w_t$ from (3.13) for x_t in (3.12) to obtain

$$\phi(B)\psi(B)w_t = w_t. \quad (3.14)$$

The coefficients of B on the left-hand side of (3.14) must be equal to those on right-hand side of (3.14), which means

$$(1 - \phi B)(1 + \psi_1 B + \psi_2 B^2 + \cdots + \psi_j B^j + \cdots) = 1. \quad (3.15)$$

Reorganizing the coefficients in (3.15),

$$1 + (\psi_1 - \phi)B + (\psi_2 - \psi_1\phi)B^2 + \cdots + (\psi_j - \psi_{j-1}\phi)B^j + \cdots = 1,$$

we see that for each $j = 1, 2, \dots$, the coefficient of B^j on the left must be zero because it is zero on the right. The coefficient of B on the left is $(\psi_1 - \phi)$, and equating this to zero, $\psi_1 - \phi = 0$, leads to $\psi_1 = \phi$. Continuing, the coefficient of B^2 is $(\psi_2 - \psi_1\phi)$, so $\psi_2 = \phi^2$. In general,

$$\psi_j = \psi_{j-1}\phi,$$

with $\psi_0 = 1$, which leads to the solution $\psi_j = \phi^j$.

Another way to think about the operations we just performed is to consider the AR(1) model in operator form, $\phi(B)x_t = w_t$. Now multiply both sides by $\phi^{-1}(B)$ (assuming the inverse operator exists) to get

$$\phi^{-1}(B)\phi(B)x_t = \phi^{-1}(B)w_t,$$

or

$$x_t = \phi^{-1}(B)w_t.$$

We know already that

$$\phi^{-1}(B) = 1 + \phi B + \phi^2 B^2 + \cdots + \phi^j B^j + \cdots,$$

that is, $\phi^{-1}(B)$ is $\psi(B)$ in (3.13). Thus, we notice that working with operators is like working with polynomials. That is, consider the polynomial $\phi(z) = 1 - \phi z$, where z is a complex number and $|\phi| < 1$. Then,

$$\phi^{-1}(z) = \frac{1}{(1 - \phi z)} = 1 + \phi z + \phi^2 z^2 + \cdots + \phi^j z^j + \cdots, \quad |z| \leq 1,$$

and the coefficients of B^j in $\phi^{-1}(B)$ are the same as the coefficients of z^j in $\phi^{-1}(z)$. In other words, we may treat the backshift operator, B , as a complex number, z . These results will be generalized in our discussion of ARMA models. We will find the polynomials corresponding to the operators useful in exploring the general properties of ARMA models.

INTRODUCTION TO MOVING AVERAGE MODELS

As an alternative to the autoregressive representation in which the x_t on the left-hand side of the equation are assumed to be combined linearly, the moving average model of order q , abbreviated as MA(q), assumes the white noise w_t on the right-hand side of the defining equation are combined linearly to form the observed data.

Definition 3.3 *The moving average model of order q , or MA(q) model, is defined to be*

$$x_t = w_t + \theta_1 w_{t-1} + \theta_2 w_{t-2} + \cdots + \theta_q w_{t-q}, \quad (3.16)$$

where $w_t \sim wn(0, \sigma_w^2)$, and $\theta_1, \theta_2, \dots, \theta_q$ ($\theta_q \neq 0$) are parameters.²

The system is the same as the infinite moving average defined as the linear process (3.13), where $\psi_0 = 1$, $\psi_j = \theta_j$, for $j = 1, \dots, q$, and $\psi_j = 0$ for other values. We may also write the MA(q) process in the equivalent form

$$x_t = \theta(B)w_t, \quad (3.17)$$

using the following definition.

Definition 3.4 *The moving average operator is*

$$\theta(B) = 1 + \theta_1 B + \theta_2 B^2 + \cdots + \theta_q B^q. \quad (3.18)$$

Unlike the autoregressive process, the moving average process is stationary for any values of the parameters $\theta_1, \dots, \theta_q$; details of this result are provided in §3.4.

² Some texts and software packages write the MA model with negative coefficients; that is, $x_t = w_t - \theta_1 w_{t-1} - \theta_2 w_{t-2} - \cdots - \theta_q w_{t-q}$.

Example 3.4 The MA(1) Process

Consider the MA(1) model $x_t = w_t + \theta w_{t-1}$. Then, $E(x_t) = 0$,

$$\gamma(h) = \begin{cases} (1 + \theta^2)\sigma_w^2 & h = 0, \\ \theta\sigma_w^2 & h = 1, \\ 0 & h > 1, \end{cases}$$

and the ACF is

$$\rho(h) = \begin{cases} \frac{\theta}{(1+\theta^2)} & h = 1, \\ 0 & h > 1. \end{cases}$$

Note $|\rho(1)| \leq 1/2$ for all values of θ (Problem 3.1). Also, x_t is correlated with x_{t-1} , but not with $x_{t-2}, x_{t-3}, \dots$. Contrast this with the case of the AR(1) model in which the correlation between x_t and x_{t-k} is never zero. When $\theta = .5$, for example, x_t and x_{t-1} are positively correlated, and $\rho(1) = .4$. When $\theta = -.5$, x_t and x_{t-1} are negatively correlated, $\rho(1) = -.4$. Figure 3.2 shows a time plot of these two processes with $\sigma_w^2 = 1$. The series in where $\theta = .5$ is smoother than the series where $\theta = -.5$.

A figure similar to Figure 3.2 can be created in R as follows:

```
par(mfrow = c(2,1))
plot(arima.sim(list(order=c(0,0,1), ma=.5), n=100), ylab="x",
     main=(expression(MA(1)~~~theta==+.5)))
plot(arima.sim(list(order=c(0,0,1), ma=-.5), n=100), ylab="x",
     main=(expression(MA(1)~~~theta==-.5)))
```

Example 3.5 Non-uniqueness of MA Models and Invertibility

Using Example 3.4, we note that for an MA(1) model, $\rho(h)$ is the same for θ and $\frac{1}{\theta}$; try 5 and $\frac{1}{5}$, for example. In addition, the pair $\sigma_w^2 = 1$ and $\theta = 5$ yield the same autocovariance function as the pair $\sigma_w^2 = 25$ and $\theta = 1/5$, namely,

$$\gamma(h) = \begin{cases} 26 & h = 0, \\ 5 & h = 1, \\ 0 & h > 1. \end{cases}$$

Thus, the MA(1) processes

$$x_t = w_t + \frac{1}{5}w_{t-1}, \quad w_t \sim \text{iid } N(0, 25)$$

and

$$y_t = v_t + 5v_{t-1}, \quad v_t \sim \text{iid } N(0, 1)$$

are the same because of normality (i.e., all finite distributions are the same). We can only observe the time series, x_t or y_t , and not the noise, w_t or v_t , so we cannot distinguish between the models. Hence, we will have to choose only one of them. For convenience, by mimicking the criterion of causality for AR models, we will




Fig. 3.2. Simulated MA(1) models: $\theta = .5$ (top); $\theta = -.5$ (bottom).

choose the model with an infinite AR representation. Such a process is called an invertible process.

To discover which model is the invertible model, we can reverse the roles of x_t and w_t (because we are mimicking the AR case) and write the MA(1) model as $w_t = -\theta w_{t-1} + x_t$. Following the steps that led to (3.6), if $|\theta| < 1$, then $w_t = \sum_{j=0}^{\infty} (-\theta)^j x_{t-j}$, which is the desired infinite AR representation of the model. Hence, given a choice, we will choose the model with $\sigma_w^2 = 25$ and $\theta = 1/5$ because it is invertible.

As in the AR case, the polynomial, $\theta(z)$, corresponding to the moving average operators, $\theta(B)$, will be useful in exploring general properties of MA processes. For example, following the steps of equations (3.12)–(3.15), we can write the MA(1) model as $x_t = \theta(B)w_t$, where $\theta(B) = 1 + \theta B$. If $|\theta| < 1$, then we can write the model as $\pi(B)x_t = w_t$, where $\pi(B) = \theta^{-1}(B)$. Let $\theta(z) = 1 + \theta z$, for $|z| \leq 1$, then $\pi(z) = \theta^{-1}(z) = 1/(1 + \theta z) = \sum_{j=0}^{\infty} (-\theta)^j z^j$, and we determine that $\pi(B) = \sum_{j=0}^{\infty} (-\theta)^j B^j$.

AUTOREGRESSIVE MOVING AVERAGE MODELS

We now proceed with the general development of autoregressive, moving average, and mixed autoregressive moving average (ARMA), models for stationary time series.

Definition 3.5 A time series $\{x_t; t = 0, \pm 1, \pm 2, \dots\}$ is **ARMA**(p, q) if it is stationary and

$$x_t = \phi_1 x_{t-1} + \dots + \phi_p x_{t-p} + w_t + \theta_1 w_{t-1} + \dots + \theta_q w_{t-q}, \quad (3.19)$$

with $\phi_p \neq 0$, $\theta_q \neq 0$, and $\sigma_w^2 > 0$. The parameters p and q are called the autoregressive and the moving average orders, respectively. If x_t has a nonzero mean μ , we set $\alpha = \mu(1 - \phi_1 - \dots - \phi_p)$ and write the model as

$$x_t = \alpha + \phi_1 x_{t-1} + \dots + \phi_p x_{t-p} + w_t + \theta_1 w_{t-1} + \dots + \theta_q w_{t-q}, \quad (3.20)$$

where $w_t \sim wn(0, \sigma_w^2)$.

As previously noted, when $q = 0$, the model is called an autoregressive model of order p , AR(p), and when $p = 0$, the model is called a moving average model of order q , MA(q). To aid in the investigation of ARMA models, it will be useful to write them using the AR operator, (3.5), and the MA operator, (3.18). In particular, the ARMA(p, q) model in (3.19) can then be written in concise form as

$$\phi(B)x_t = \theta(B)w_t. \quad (3.21)$$

The concise form of the model points to a potential problem in that we can unnecessarily complicate the model by multiplying both sides by another operator, say

$$\eta(B)\phi(B)x_t = \eta(B)\theta(B)w_t,$$

without changing the dynamics. Consider the following example.

Example 3.6 Parameter Redundancy

Consider a white noise process $x_t = w_t$. Equivalently, we can write this as $.5x_{t-1} = .5w_{t-1}$ by shifting back one unit of time and multiplying by .5. Now, subtract the two representations to obtain

$$x_t - .5x_{t-1} = w_t - .5w_{t-1},$$

or

$$x_t = .5x_{t-1} - .5w_{t-1} + w_t, \quad (3.22)$$

which looks like an ARMA(1, 1) model. Of course, x_t is still white noise; nothing has changed in this regard [i.e., $x_t = w_t$ is the solution to (3.22)], but we have hidden the fact that x_t is white noise because of the parameter redundancy or over-parameterization. Write the parameter redundant model in operator form as $\phi(B)x_t = \theta(B)w_t$, or

$$(1 - .5B)x_t = (1 - .5B)w_t.$$

Apply the operator $\phi(B)^{-1} = (1 - .5B)^{-1}$ to both sides to obtain

$$x_t = (1 - .5B)^{-1}(1 - .5B)x_t = (1 - .5B)^{-1}(1 - .5B)w_t = w_t,$$

which is the original model. We can easily detect the problem of over-parameterization with the use of the operators or their associated polynomials. That is, write the AR

polynomial $\phi(z) = (1 - .5z)$, the MA polynomial $\theta(z) = (1 - .5z)$, and note that both polynomials have a common factor, namely $(1 - .5z)$. This common factor immediately identifies the parameter redundancy. Discarding the common factor in each leaves $\phi(z) = 1$ and $\theta(z) = 1$, from which we conclude $\phi(B) = 1$ and $\theta(B) = 1$, and we deduce that the model is actually white noise.

The consideration of parameter redundancy will be crucial when we discuss estimation for general ARMA models. As this example points out, we might fit an ARMA(1, 1) model to white noise data and find that the parameter estimates are significant. If we were unaware of parameter redundancy, we might claim the data are correlated when in fact they are not (Problem 3.20). Although we have not yet discussed estimation, we present the following demonstration of the problem. We generated 150 iid standard normals and then fit an ARMA(1, 1) to the data. Note that $\hat{\phi} = -.96$ and $\hat{\theta} = .95$, and both are significant. Below is the R code (note that the estimate called ‘intercept’ is really the estimate of the mean).

```
set.seed(8675309)           # Jenny, I got your number
x = rnorm(150)              # generate iid N(0,1)s
arima(x, order=c(1,0,1))    # estimation
Coefficients:
      ar1      ma1 intercept<= misnomer
    -0.9595  0.9527   0.0462
s.e.    0.1688  0.1750   0.0727
```

Thus, forgetting the (non-significant) mean estimate, the fitted model looks like

$$(1 + .96B)x_t = (1 + .95B)w_t.$$

Example 3.2, Example 3.5, and Example 3.6 point to a number of problems with the general definition of ARMA(p, q) models, as given by (3.19), or, equivalently, by (3.21). To summarize, we have seen the following problems:

- (i) parameter redundant models,
- (ii) stationary AR models that depend on the future, and
- (iii) MA models that are not unique.

To overcome these problems, we will require some additional restrictions on the model parameters. First, we make the following definitions.

Definition 3.6 *The AR and MA polynomials are defined as*

$$\phi(z) = 1 - \phi_1 z - \cdots - \phi_p z^p, \quad \phi_p \neq 0, \quad (3.23)$$

and

$$\theta(z) = 1 + \theta_1 z + \cdots + \theta_q z^q, \quad \theta_q \neq 0, \quad (3.24)$$

respectively, where z is a complex number.

To address the first problem, we will henceforth refer to an ARMA(p, q) model to mean that it is in its simplest form. That is, in addition to the original definition given in equation (3.19), we will also require that $\phi(z)$ and $\theta(z)$ have no common factors.

So, the process, $x_t = .5x_{t-1} - .5w_{t-1} + w_t$, discussed in [Example 3.6](#) is not referred to as an ARMA(1, 1) process because, in its reduced form, x_t is white noise.

To address the problem of future-dependent models, we formally introduce the concept of causality.

Definition 3.7 An ARMA(p, q) model is said to be **causal**, if the time series $\{x_t; t = 0, \pm 1, \pm 2, \dots\}$ can be written as a one-sided linear process:

$$x_t = \sum_{j=0}^{\infty} \psi_j w_{t-j} = \psi(B)w_t, \quad (3.25)$$

where $\psi(B) = \sum_{j=0}^{\infty} \psi_j B^j$, and $\sum_{j=0}^{\infty} |\psi_j| < \infty$; we set $\psi_0 = 1$.

In [Example 3.2](#), the AR(1) process, $x_t = \phi x_{t-1} + w_t$, is causal only when $|\phi| < 1$. Equivalently, the process is causal only when the root of $\phi(z) = 1 - \phi z$ is bigger than one in absolute value. That is, the root, say, z_0 , of $\phi(z)$ is $z_0 = 1/\phi$ (because $\phi(z_0) = 0$) and $|z_0| > 1$ because $|\phi| < 1$. In general, we have the following property.

Property 3.1 Causality of an ARMA(p, q) Process

An ARMA(p, q) model is causal if and only if $\phi(z) \neq 0$ for $|z| \leq 1$. The coefficients of the linear process given in (3.25) can be determined by solving

$$\psi(z) = \sum_{j=0}^{\infty} \psi_j z^j = \frac{\theta(z)}{\phi(z)}, \quad |z| \leq 1.$$

Another way to phrase [Property 3.1](#) is that an ARMA process is causal only when the roots of $\phi(z)$ lie outside the unit circle; that is, $\phi(z) = 0$ only when $|z| > 1$. Finally, to address the problem of uniqueness discussed in [Example 3.5](#), we choose the model that allows an infinite autoregressive representation.

Definition 3.8 An ARMA(p, q) model is said to be **invertible**, if the time series $\{x_t; t = 0, \pm 1, \pm 2, \dots\}$ can be written as

$$\pi(B)x_t = \sum_{j=0}^{\infty} \pi_j x_{t-j} = w_t, \quad (3.26)$$

where $\pi(B) = \sum_{j=0}^{\infty} \pi_j B^j$, and $\sum_{j=0}^{\infty} |\pi_j| < \infty$; we set $\pi_0 = 1$.

Analogous to [Property 3.1](#), we have the following property.

Property 3.2 Invertibility of an ARMA(p, q) Process

An ARMA(p, q) model is invertible if and only if $\theta(z) \neq 0$ for $|z| \leq 1$. The coefficients π_j of $\pi(B)$ given in (3.26) can be determined by solving

$$\pi(z) = \sum_{j=0}^{\infty} \pi_j z^j = \frac{\phi(z)}{\theta(z)}, \quad |z| \leq 1.$$

Another way to phrase [Property 3.2](#) is that an ARMA process is invertible only when the roots of $\theta(z)$ lie outside the unit circle; that is, $\theta(z) = 0$ only when $|z| > 1$. The proof of [Property 3.1](#) is given in [§B.2](#) (the proof of [Property 3.2](#) is similar). The following examples illustrate these concepts.

Example 3.7 Parameter Redundancy, Causality, Invertibility

Consider the process

$$x_t = .4x_{t-1} + .45x_{t-2} + w_t + w_{t-1} + .25w_{t-2},$$

or, in operator form,

$$(1 - .4B - .45B^2)x_t = (1 + B + .25B^2)w_t.$$

At first, x_t appears to be an ARMA(2, 2) process. But notice that

$$\phi(B) = 1 - .4B - .45B^2 = (1 + .5B)(1 - .9B)$$

and

$$\theta(B) = (1 + B + .25B^2) = (1 + .5B)^2$$

have a common factor that can be canceled. After cancellation, the operators are $\phi(B) = (1 - .9B)$ and $\theta(B) = (1 + .5B)$, so the model is an ARMA(1, 1) model, $(1 - .9B)x_t = (1 + .5B)w_t$, or

$$x_t = .9x_{t-1} + .5w_{t-1} + w_t. \quad (3.27)$$

The model is causal because $\phi(z) = (1 - .9z) = 0$ when $z = 10/9$, which is outside the unit circle. The model is also invertible because the root of $\theta(z) = (1 + .5z)$ is $z = -2$, which is outside the unit circle.

To write the model as a linear process, we can obtain the ψ -weights using **Property 3.1**, $\phi(z)\psi(z) = \theta(z)$, or

$$(1 - .9z)(1 + \psi_1z + \psi_2z^2 + \cdots + \psi_jz^j + \cdots) = 1 + .5z.$$

Rearranging, we get

$$1 + (\psi_1 - .9)z + (\psi_2 - .9\psi_1)z^2 + \cdots + (\psi_j - .9\psi_{j-1})z^j + \cdots = 1 + .5z.$$

Matching the coefficients of z on the left and right sides we get $\psi_1 - .9 = .5$ and $\psi_j - .9\psi_{j-1} = 0$ for $j > 1$. Thus, $\psi_j = 1.4(.9)^{j-1}$ for $j \geq 1$ and (3.27) can be written as

$$x_t = w_t + 1.4 \sum_{j=1}^{\infty} .9^{j-1} w_{t-j}.$$

The values of ψ_j may be calculated in R as follows:

```
ARMAtoMA(ar = .9, ma = .5, 10) # first 10 psi-weights
[1] 1.40 1.26 1.13 1.02 0.92 0.83 0.74 0.67 0.60 0.54
```

The invertible representation using **Property 3.1** is obtained by matching coefficients in $\theta(z)\pi(z) = \phi(z)$,

$$(1 + .5z)(1 + \pi_1z + \pi_2z^2 + \pi_3z^3 + \cdots) = 1 - .9z.$$

In this case, the π -weights are given by $\pi_j = (-1)^j 1.4 (.5)^{j-1}$, for $j \geq 1$, and hence, we can also write (3.27) as




Fig. 3.3. Causal region for an AR(2) in terms of the parameters.

$$x_t = -1.4 \sum_{j=1}^{\infty} (-.5)^{j-1} x_{t-j} + w_t.$$

The values of π_j may be calculated in R as follows by reversing the roles of w_t and x_t ; i.e., write the model as $w_t = -.5w_{t-1} + x_t - .9x_{t-1}$:

```
ARMAtoMA(ar = -.5, ma = -.9, 10) # first 10 pi-weights
[1] -1.400 .700 -.350 .175 -.087 .044 -.022 .011 -.006 .003
```

Example 3.8 Causal Conditions for an AR(2) Process

For an AR(1) model, $(1 - \phi B)x_t = w_t$, to be causal, the root of $\phi(z) = 1 - \phi z$ must lie outside of the unit circle. In this case, the root (or zero) occurs at $z_0 = 1/\phi$ [i.e., $\phi(z_0) = 0$], so it is easy to go from the causal requirement on the root, $|1/\phi| > 1$, to a requirement on the parameter, $|\phi| < 1$. It is not so easy to establish this relationship for higher order models.

For example, the AR(2) model, $(1 - \phi_1 B - \phi_2 B^2)x_t = w_t$, is causal when the two roots of $\phi(z) = 1 - \phi_1 z - \phi_2 z^2$ lie outside of the unit circle. Using the quadratic formula, this requirement can be written as

$$\left| \frac{\phi_1 \pm \sqrt{\phi_1^2 + 4\phi_2}}{-2\phi_2} \right| > 1.$$

The roots of $\phi(z)$ may be real and distinct, real and equal, or a complex conjugate pair. If we denote those roots by z_1 and z_2 , we can write $\phi(z) = (1 - z_1^{-1}z)(1 - z_2^{-1}z)$; note that $\phi(z_1) = \phi(z_2) = 0$. The model can be written in operator form as $(1 - z_1^{-1}B)(1 - z_2^{-1}B)x_t = w_t$. From this representation, it follows that $\phi_1 = (z_1^{-1} + z_2^{-1})$ and $\phi_2 = -(z_1 z_2)^{-1}$. This relationship and the fact that $|z_1| > 1$ and $|z_2| > 1$ can be used to establish the following equivalent condition for causality:

$$\phi_1 + \phi_2 < 1, \quad \phi_2 - \phi_1 < 1, \quad \text{and} \quad |\phi_2| < 1. \quad (3.28)$$

This causality condition specifies a triangular region in the parameter space; see [Figure 3.3](#). We leave the details of the equivalence to the reader ([Problem 3.5](#)).

3.3 Difference Equations

The study of the behavior of ARMA processes and their ACFs is greatly enhanced by a basic knowledge of difference equations, simply because they are difference equations. This topic is also useful in the study of time domain models and stochastic processes in general. We will give a brief and heuristic account of the topic along with some examples of the usefulness of the theory. For details, the reader is referred to Mickens (1990).

Suppose we have a sequence of numbers $u_0, u_1, u_2, \dots$ such that

$$u_n - \alpha u_{n-1} = 0, \quad \alpha \neq 0, \quad n = 1, 2, \dots \quad (3.29)$$

For example, recall (3.9) in which we showed that the ACF of an AR(1) process is a sequence, $\rho(h)$, satisfying

$$\rho(h) - \phi \rho(h-1) = 0, \quad h = 1, 2, \dots$$

Equation (3.29) represents a homogeneous difference equation of order 1. To solve the equation, we write:

$$\begin{aligned} u_1 &= \alpha u_0 \\ u_2 &= \alpha u_1 = \alpha^2 u_0 \\ &\vdots \\ u_n &= \alpha u_{n-1} = \alpha^n u_0. \end{aligned}$$

Given an initial condition $u_0 = c$, we may solve (3.29), namely, $u_n = \alpha^n c$.

In operator notation, (3.29) can be written as $(1 - \alpha B)u_n = 0$. The polynomial associated with (3.29) is $\alpha(z) = 1 - \alpha z$, and the root, say, z_0 , of this polynomial is $z_0 = 1/\alpha$; that is $\alpha(z_0) = 0$. We know a solution (in fact, *the* solution) to (3.29), with initial condition $u_0 = c$, is

$$u_n = \alpha^n c = \left(z_0^{-1}\right)^n c. \quad (3.30)$$

That is, the solution to the difference equation (3.29) depends only on the initial condition and the inverse of the root to the associated polynomial $\alpha(z)$.

Now suppose that the sequence satisfies

$$u_n - \alpha_1 u_{n-1} - \alpha_2 u_{n-2} = 0, \quad \alpha_2 \neq 0, \quad n = 2, 3, \dots \quad (3.31)$$

This equation is a homogeneous difference equation of order 2. The corresponding polynomial is

$$\alpha(z) = 1 - \alpha_1 z - \alpha_2 z^2,$$

which has two roots, say, z_1 and z_2 ; that is, $\alpha(z_1) = \alpha(z_2) = 0$. We will consider two cases. First suppose $z_1 \neq z_2$. Then the general solution to (3.31) is

$$u_n = c_1 z_1^{-n} + c_2 z_2^{-n}, \quad (3.32)$$

where c_1 and c_2 depend on the initial conditions. The claim that is a solution can be verified by direct substitution of (3.32) into (3.31):

$$\begin{aligned} (c_1 z_1^{-n} + c_2 z_2^{-n}) - \alpha_1 (c_1 z_1^{-(n-1)} + c_2 z_2^{-(n-1)}) - \alpha_2 (c_1 z_1^{-(n-2)} + c_2 z_2^{-(n-2)}) \\ = c_1 z_1^{-n} (1 - \alpha_1 z_1 - \alpha_2 z_1^2) + c_2 z_2^{-n} (1 - \alpha_1 z_2 - \alpha_2 z_2^2) \\ = c_1 z_1^{-n} \alpha(z_1) + c_2 z_2^{-n} \alpha(z_2) = 0. \end{aligned}$$

Given two initial conditions u_0 and u_1 , we may solve for c_1 and c_2 :

$$u_0 = c_1 + c_2 \quad \text{and} \quad u_1 = c_1 z_1^{-1} + c_2 z_2^{-1},$$

where z_1 and z_2 can be solved for in terms of α_1 and α_2 using the quadratic formula, for example.

When the roots are equal, $z_1 = z_2 (= z_0)$, a general solution to (3.31) is

$$u_n = z_0^{-n} (c_1 + c_2 n). \quad (3.33)$$

This claim can also be verified by direct substitution of (3.33) into (3.31):

$$\begin{aligned} z_0^{-n} (c_1 + c_2 n) - \alpha_1 (z_0^{-(n-1)} [c_1 + c_2 (n-1)]) - \alpha_2 (z_0^{-(n-2)} [c_1 + c_2 (n-2)]) \\ = z_0^{-n} (c_1 + c_2 n) (1 - \alpha_1 z_0 - \alpha_2 z_0^2) + c_2 z_0^{-n+1} (\alpha_1 + 2\alpha_2 z_0) \\ = c_2 z_0^{-n+1} (\alpha_1 + 2\alpha_2 z_0). \end{aligned}$$

To show that $(\alpha_1 + 2\alpha_2 z_0) = 0$, write $1 - \alpha_1 z - \alpha_2 z^2 = (1 - z_0^{-1} z)^2$, and take derivatives with respect to z on both sides of the equation to obtain $(\alpha_1 + 2\alpha_2 z) = 2z_0^{-1} (1 - z_0^{-1} z)$. Thus, $(\alpha_1 + 2\alpha_2 z_0) = 2z_0^{-1} (1 - z_0^{-1} z_0) = 0$, as was to be shown. Finally, given two initial conditions, u_0 and u_1 , we can solve for c_1 and c_2 :

$$u_0 = c_1 \quad \text{and} \quad u_1 = (c_1 + c_2) z_0^{-1}.$$

It can also be shown that these solutions are unique.

To summarize these results, in the case of distinct roots, the solution to the homogeneous difference equation of degree two was

$$\begin{aligned} u_n = z_1^{-n} \times (\text{a polynomial in } n \text{ of degree } m_1 - 1) \\ + z_2^{-n} \times (\text{a polynomial in } n \text{ of degree } m_2 - 1), \end{aligned} \quad (3.34)$$

where m_1 is the multiplicity of the root z_1 and m_2 is the multiplicity of the root z_2 . In this example, of course, $m_1 = m_2 = 1$, and we called the polynomials of degree zero c_1 and c_2 , respectively. In the case of the repeated root, the solution was

$$u_n = z_0^{-n} \times (\text{a polynomial in } n \text{ of degree } m_0 - 1), \quad (3.35)$$

where m_0 is the multiplicity of the root z_0 ; that is, $m_0 = 2$. In this case, we wrote the polynomial of degree one as $c_1 + c_2 n$. In both cases, we solved for c_1 and c_2 given two initial conditions, u_0 and u_1 .

Example 3.9 The ACF of an AR(2) Process

Suppose $x_t = \phi_1 x_{t-1} + \phi_2 x_{t-2} + w_t$ is a causal AR(2) process. Multiply each side of the model by x_{t-h} for $h > 0$, and take expectation:

$$E(x_t x_{t-h}) = \phi_1 E(x_{t-1} x_{t-h}) + \phi_2 E(x_{t-2} x_{t-h}) + E(w_t x_{t-h}).$$

The result is

$$\gamma(h) = \phi_1 \gamma(h-1) + \phi_2 \gamma(h-2), \quad h = 1, 2, \dots \quad (3.36)$$

In (3.36), we used the fact that $E(x_t) = 0$ and for $h > 0$,

$$E(w_t x_{t-h}) = E\left(w_t \sum_{j=0}^{\infty} \psi_j w_{t-h-j}\right) = 0.$$

Divide (3.36) through by $\gamma(0)$ to obtain the difference equation for the ACF of the process:

$$\rho(h) - \phi_1 \rho(h-1) - \phi_2 \rho(h-2) = 0, \quad h = 1, 2, \dots \quad (3.37)$$

The initial conditions are $\rho(0) = 1$ and $\rho(-1) = \phi_1/(1 - \phi_2)$, which is obtained by evaluating (3.37) for $h = 1$ and noting that $\rho(1) = \rho(-1)$.

Using the results for the homogeneous difference equation of order two, let z_1 and z_2 be the roots of the associated polynomial, $\phi(z) = 1 - \phi_1 z - \phi_2 z^2$. Because the model is causal, we know the roots are outside the unit circle: $|z_1| > 1$ and $|z_2| > 1$. Now, consider the solution for three cases:

(i) When z_1 and z_2 are real and distinct, then

$$\rho(h) = c_1 z_1^{-h} + c_2 z_2^{-h},$$

so $\rho(h) \rightarrow 0$ exponentially fast as $h \rightarrow \infty$.

(ii) When $z_1 = z_2 (= z_0)$ are real and equal, then

$$\rho(h) = z_0^{-h} (c_1 + c_2 h),$$

so $\rho(h) \rightarrow 0$ exponentially fast as $h \rightarrow \infty$.

(iii) When $z_1 = \bar{z}_2$ are a complex conjugate pair, then $c_2 = \bar{c}_1$ (because $\rho(h)$ is real), and

$$\rho(h) = c_1 z_1^{-h} + \bar{c}_1 \bar{z}_1^{-h}.$$

Write c_1 and z_1 in polar coordinates, for example, $z_1 = |z_1|e^{i\theta}$, where θ is the angle whose tangent is the ratio of the imaginary part and the real part of z_1




Fig. 3.4. Simulated AR(2) model, $n = 144$ with $\phi_1 = 1.5$ and $\phi_2 = -.75$.

(sometimes called $\arg(z_1)$; the range of θ is $[-\pi, \pi]$). Then, using the fact that $e^{i\alpha} + e^{-i\alpha} = 2 \cos(\alpha)$, the solution has the form

$$\rho(h) = a|z_1|^{-h} \cos(h\theta + b),$$

where a and b are determined by the initial conditions. Again, $\rho(h)$ dampens to zero exponentially fast as $h \rightarrow \infty$, but it does so in a sinusoidal fashion. The implication of this result is shown in the next example.

Example 3.10 An AR(2) with Complex Roots

Figure 3.4 shows $n = 144$ observations from the AR(2) model

$$x_t = 1.5x_{t-1} - .75x_{t-2} + w_t,$$

with $\sigma_w^2 = 1$, and with complex roots chosen so the process exhibits pseudo-cyclic behavior at the rate of one cycle every 12 time points. The autoregressive polynomial for this model is $\phi(z) = 1 - 1.5z + .75z^2$. The roots of $\phi(z)$ are $1 \pm i/\sqrt{3}$, and $\theta = \tan^{-1}(1/\sqrt{3}) = 2\pi/12$ radians per unit time. To convert the angle to cycles per unit time, divide by 2π to get $1/12$ cycles per unit time. The ACF for this model is shown in §3.4, Figure 3.5.

To calculate the roots of the polynomial and solve for $\arg$ in R:

```
z = c(1, -1.5, .75)      # coefficients of the polynomial
(a = polyroot(z)[1])     # print one root: 1+0.57735i = 1 + i/sqrt(3)
arg = Arg(a)/(2*pi)       # arg in cycles/pt
1/arg                     # = 12, the pseudo period
```

To reproduce Figure 3.4:

```
set.seed(8675309)
ar2 = arima.sim(list(order=c(2,0,0), ar=c(1.5,-.75)), n = 144)
plot(ar2, axes=FALSE, xlab="Time")
axis(2); axis(1, at=seq(0,144,by=12)); box()
abline(v=seq(0,144,by=12), lty=2)
```

To calculate and display the ACF for this model:

```
ACF = ARMAacf(ar=c(1.5,-.75), ma=0, 50)
plot(ACF, type="h", xlab="lag")
abline(h=0)
```

We now exhibit the solution for the general homogeneous difference equation of order p :

$$u_n - \alpha_1 u_{n-1} - \cdots - \alpha_p u_{n-p} = 0, \quad \alpha_p \neq 0, \quad n = p, p+1, \dots \quad (3.38)$$

The associated polynomial is

$$\alpha(z) = 1 - \alpha_1 z - \cdots - \alpha_p z^p.$$

Suppose $\alpha(z)$ has r distinct roots, z_1 with multiplicity m_1 , z_2 with multiplicity m_2 , $\dots$, and z_r with multiplicity m_r , such that $m_1 + m_2 + \cdots + m_r = p$. The general solution to the difference equation (3.38) is

$$u_n = z_1^{-n} P_1(n) + z_2^{-n} P_2(n) + \cdots + z_r^{-n} P_r(n), \quad (3.39)$$

where $P_j(n)$, for $j = 1, 2, \dots, r$, is a polynomial in n , of degree $m_j - 1$. Given p initial conditions $u_0, \dots, u_{p-1}$, we can solve for the $P_j(n)$ explicitly.

Example 3.11 The ψ -weights for an ARMA Model

For a causal ARMA(p, q) model, $\phi(B)x_t = \theta(B)w_t$, where the zeros of $\phi(z)$ are outside the unit circle, recall that we may write

$$x_t = \sum_{j=0}^{\infty} \psi_j w_{t-j},$$

where the ψ -weights are determined using [Property 3.1](#).

For the pure MA(q) model, $\psi_0 = 1$, $\psi_j = \theta_j$, for $j = 1, \dots, q$, and $\psi_j = 0$, otherwise. For the general case of ARMA(p, q) models, the task of solving for the ψ -weights is much more complicated, as was demonstrated in [Example 3.7](#). The use of the theory of homogeneous difference equations can help here. To solve for the ψ -weights in general, we must match the coefficients in $\phi(z)\psi(z) = \theta(z)$:

$$(1 - \phi_1 z - \phi_2 z^2 - \cdots)(\psi_0 + \psi_1 z + \psi_2 z^2 + \cdots) = (1 + \theta_1 z + \theta_2 z^2 + \cdots).$$

The first few values are

$$\begin{aligned} \psi_0 &= 1 \\ \psi_1 - \phi_1 \psi_0 &= \theta_1 \\ \psi_2 - \phi_1 \psi_1 - \phi_2 \psi_0 &= \theta_2 \\ \psi_3 - \phi_1 \psi_2 - \phi_2 \psi_1 - \phi_3 \psi_0 &= \theta_3 \\ &\vdots \end{aligned}$$

where we would take $\phi_j = 0$ for $j > p$, and $\theta_j = 0$ for $j > q$. The ψ -weights satisfy the homogeneous difference equation given by

$$\psi_j - \sum_{k=1}^p \phi_k \psi_{j-k} = 0, \quad j \geq \max(p, q + 1), \quad (3.40)$$

with initial conditions

$$\psi_j - \sum_{k=1}^j \phi_k \psi_{j-k} = \theta_j, \quad 0 \leq j < \max(p, q + 1). \quad (3.41)$$

The general solution depends on the roots of the AR polynomial $\phi(z) = 1 - \phi_1 z - \cdots - \phi_p z^p$, as seen from (3.40). The specific solution will, of course, depend on the initial conditions.

Consider the ARMA process given in (3.27), $x_t = .9x_{t-1} + .5w_{t-1} + w_t$. Because $\max(p, q + 1) = 2$, using (3.41), we have $\psi_0 = 1$ and $\psi_1 = .9 + .5 = 1.4$. By (3.40), for $j = 2, 3, \dots$, the ψ -weights satisfy $\psi_j - .9\psi_{j-1} = 0$. The general solution is $\psi_j = c .9^j$. To find the specific solution, use the initial condition $\psi_1 = 1.4$, so $1.4 = .9c$ or $c = 1.4/.9$. Finally, $\psi_j = 1.4(.9)^{j-1}$, for $j \geq 1$, as we saw in Example 3.7.

To view, for example, the first 50 ψ -weights in R, use:

```
ARMAtoMA(ar=.9, ma=.5, 50)      # for a list
plot(ARMAtoMA(ar=.9, ma=.5, 50)) # for a graph
```

3.4 Autocorrelation and Partial Autocorrelation

We begin by exhibiting the ACF of an MA(q) process, $x_t = \theta(B)w_t$, where $\theta(B) = 1 + \theta_1 B + \cdots + \theta_q B^q$. Because x_t is a finite linear combination of white noise terms, the process is stationary with mean

$$E(x_t) = \sum_{j=0}^q \theta_j E(w_{t-j}) = 0,$$

where we have written $\theta_0 = 1$, and with autocovariance function

$$\begin{aligned} \gamma(h) &= \text{cov}(x_{t+h}, x_t) = \text{cov}\left(\sum_{j=0}^q \theta_j w_{t+h-j}, \sum_{k=0}^q \theta_k w_{t-k}\right) \\ &= \begin{cases} \sigma_w^2 \sum_{j=0}^{q-h} \theta_j \theta_{j+h}, & 0 \leq h \leq q \\ 0 & h > q. \end{cases} \end{aligned} \quad (3.42)$$

Recall that $\gamma(h) = \gamma(-h)$, so we will only display the values for $h \geq 0$. Note that $\gamma(q)$ cannot be zero because $\theta_q \neq 0$. The cutting off of $\gamma(h)$ after q lags is the signature of the MA(q) model. Dividing (3.42) by $\gamma(0)$ yields the ACF of an MA(q):

$$\rho(h) = \begin{cases} \frac{\sum_{j=0}^{q-h} \theta_j \theta_{j+h}}{1 + \theta_1^2 + \cdots + \theta_q^2} & 1 \leq h \leq q \\ 0 & h > q. \end{cases} \quad (3.43)$$

For a causal ARMA(p, q) model, $\phi(B)x_t = \theta(B)w_t$, where the zeros of $\phi(z)$ are outside the unit circle, write

$$x_t = \sum_{j=0}^{\infty} \psi_j w_{t-j}. \quad (3.44)$$

It follows immediately that $E(x_t) = 0$. Also, the autocovariance function of x_t can be written as

$$\gamma(h) = \text{cov}(x_{t+h}, x_t) = \sigma_w^2 \sum_{j=0}^{\infty} \psi_j \psi_{j+h}, \quad h \geq 0. \quad (3.45)$$

We could then use (3.40) and (3.41) to solve for the ψ -weights. In turn, we could solve for $\gamma(h)$, and the ACF $\rho(h) = \gamma(h)/\gamma(0)$. As in [Example 3.9](#), it is also possible to obtain a homogeneous difference equation directly in terms of $\gamma(h)$. First, we write

$$\begin{aligned} \gamma(h) &= \text{cov}(x_{t+h}, x_t) = \text{cov}\left(\sum_{j=1}^p \phi_j x_{t+h-j} + \sum_{j=0}^q \theta_j w_{t+h-j}, x_t\right) \\ &= \sum_{j=1}^p \phi_j \gamma(h-j) + \sigma_w^2 \sum_{j=h}^q \theta_j \psi_{j-h}, \quad h \geq 0, \end{aligned} \quad (3.46)$$

where we have used the fact that, for $h \geq 0$,

$$\text{cov}(w_{t+h-j}, x_t) = \text{cov}\left(w_{t+h-j}, \sum_{k=0}^{\infty} \psi_k w_{t-k}\right) = \psi_{j-h} \sigma_w^2.$$

From (3.46), we can write a general homogeneous equation for the ACF of a causal ARMA process:

$$\gamma(h) - \phi_1 \gamma(h-1) - \cdots - \phi_p \gamma(h-p) = 0, \quad h \geq \max(p, q+1), \quad (3.47)$$

with initial conditions

$$\gamma(h) - \sum_{j=1}^p \phi_j \gamma(h-j) = \sigma_w^2 \sum_{j=h}^q \theta_j \psi_{j-h}, \quad 0 \leq h < \max(p, q+1). \quad (3.48)$$

Dividing (3.47) and (3.48) through by $\gamma(0)$ will allow us to solve for the ACF, $\rho(h) = \gamma(h)/\gamma(0)$.

Example 3.12 The ACF of an AR(p)

In [Example 3.9](#) we considered the case where $p = 2$. For the general case, it follows immediately from (3.47) that

$$\rho(h) - \phi_1 \rho(h-1) - \cdots - \phi_p \rho(h-p) = 0, \quad h \geq p. \quad (3.49)$$

Let $z_1, \dots, z_r$ denote the roots of $\phi(z)$, each with multiplicity $m_1, \dots, m_r$, respectively, where $m_1 + \cdots + m_r = p$. Then, from (3.39), the general solution is

$$\rho(h) = z_1^{-h} P_1(h) + z_2^{-h} P_2(h) + \cdots + z_r^{-h} P_r(h), \quad h \geq p, \quad (3.50)$$

where $P_j(h)$ is a polynomial in h of degree $m_j - 1$.

Recall that for a causal model, all of the roots are outside the unit circle, $|z_i| > 1$, for $i = 1, \dots, r$. If all the roots are real, then $\rho(h)$ dampens exponentially fast to zero as $h \rightarrow \infty$. If some of the roots are complex, then they will be in conjugate pairs and $\rho(h)$ will dampen, in a sinusoidal fashion, exponentially fast to zero as $h \rightarrow \infty$. In the case of complex roots, the time series will appear to be cyclic in nature. This, of course, is also true for ARMA models in which the AR part has complex roots.

Example 3.13 The ACF of an ARMA(1, 1)

Consider the ARMA(1, 1) process $x_t = \phi x_{t-1} + \theta w_{t-1} + w_t$, where $|\phi| < 1$. Based on (3.47), the autocovariance function satisfies

$$\gamma(h) - \phi \gamma(h-1) = 0, \quad h = 2, 3, \dots,$$

and it follows from (3.29)–(3.30) that the general solution is

$$\gamma(h) = c \phi^h, \quad h = 1, 2, \dots \quad (3.51)$$

To obtain the initial conditions, we use (3.48):

$$\gamma(0) = \phi \gamma(1) + \sigma_w^2 [1 + \theta \phi + \theta^2] \quad \text{and} \quad \gamma(1) = \phi \gamma(0) + \sigma_w^2 \theta.$$

Solving for $\gamma(0)$ and $\gamma(1)$, we obtain:

$$\gamma(0) = \sigma_w^2 \frac{1 + 2\theta\phi + \theta^2}{1 - \phi^2} \quad \text{and} \quad \gamma(1) = \sigma_w^2 \frac{(1 + \theta\phi)(\phi + \theta)}{1 - \phi^2}.$$

To solve for c , note that from (3.51), $\gamma(1) = c \phi$ or $c = \gamma(1)/\phi$. Hence, the specific solution for $h \geq 1$ is

$$\gamma(h) = \frac{\gamma(1)}{\phi} \phi^h = \sigma_w^2 \frac{(1 + \theta\phi)(\phi + \theta)}{1 - \phi^2} \phi^{h-1}.$$

Finally, dividing through by $\gamma(0)$ yields the ACF

$$\rho(h) = \frac{(1 + \theta\phi)(\phi + \theta)}{1 + 2\theta\phi + \theta^2} \phi^{h-1}, \quad h \geq 1. \quad (3.52)$$

Notice that the general pattern of $\rho(h)$ in (3.52) is not different from that of an AR(1) given in (3.8). Hence, it is unlikely that we will be able to tell the difference between an ARMA(1,1) and an AR(1) based solely on an ACF estimated from a sample. This consideration will lead us to the partial autocorrelation function.

THE PARTIAL AUTOCORRELATION FUNCTION (PACF)

We have seen in (3.43), for MA(q) models, the ACF will be zero for lags greater than q . Moreover, because $\theta_q \neq 0$, the ACF will not be zero at lag q . Thus, the ACF provides a considerable amount of information about the order of the dependence when the process is a moving average process. If the process, however, is ARMA or AR, the ACF alone tells us little about the orders of dependence. Hence, it is worthwhile pursuing a function that will behave like the ACF of MA models, but for AR models, namely, the partial autocorrelation function (PACF).

Recall that if X , Y , and Z are random variables, then the partial correlation between X and Y given Z is obtained by regressing X on Z to obtain $\widehat{X}$, regressing Y on Z to obtain $\widehat{Y}$, and then calculating

$$\rho_{XY|Z} = \text{corr}\{X - \widehat{X}, Y - \widehat{Y}\}.$$

The idea is that $\rho_{XY|Z}$ measures the correlation between X and Y with the linear effect of Z removed (or partialled out). If the variables are multivariate normal, then this definition coincides with $\rho_{XY|Z} = \text{corr}(X, Y | Z)$.

To motivate the idea for time series, consider a causal AR(1) model, $x_t = \phi x_{t-1} + w_t$. Then,

$$\begin{aligned} \gamma_x(2) &= \text{cov}(x_t, x_{t-2}) = \text{cov}(\phi x_{t-1} + w_t, x_{t-2}) \\ &= \text{cov}(\phi^2 x_{t-2} + \phi w_{t-1} + w_t, x_{t-2}) = \phi^2 \gamma_x(0). \end{aligned}$$

This result follows from causality because x_{t-2} involves $\{w_{t-2}, w_{t-3}, \dots\}$, which are all uncorrelated with w_t and w_{t-1} . The correlation between x_t and x_{t-2} is not zero, as it would be for an MA(1), because x_t is dependent on x_{t-2} through x_{t-1} . Suppose we break this chain of dependence by removing (or partial out) the effect x_{t-1} . That is, we consider the correlation between $x_t - \phi x_{t-1}$ and $x_{t-2} - \phi x_{t-1}$, because it is the correlation between x_t and x_{t-2} with the linear dependence of each on x_{t-1} removed. In this way, we have broken the dependence chain between x_t and x_{t-2} . In fact,

$$\text{cov}(x_t - \phi x_{t-1}, x_{t-2} - \phi x_{t-1}) = \text{cov}(w_t, x_{t-2} - \phi x_{t-1}) = 0.$$

Hence, the tool we need is partial autocorrelation, which is the correlation between x_s and x_t with the linear effect of everything “in the middle” removed.

To formally define the PACF for mean-zero stationary time series, let $\widehat{x}_{t+h}$, for $h \geq 2$, denote the regression³ of x_{t+h} on $\{x_{t+h-1}, x_{t+h-2}, \dots, x_{t+1}\}$, which we write as

$$\widehat{x}_{t+h} = \beta_1 x_{t+h-1} + \beta_2 x_{t+h-2} + \dots + \beta_{h-1} x_{t+1}. \quad (3.53)$$

No intercept term is needed in (3.53) because the mean of x_t is zero (otherwise, replace x_t by $x_t - \mu_x$ in this discussion). In addition, let $\widehat{x}_t$ denote the regression of x_t on $\{x_{t+1}, x_{t+2}, \dots, x_{t+h-1}\}$, then

³ The term regression here refers to regression in the population sense. That is, $\widehat{x}_{t+h}$ is the linear combination of $\{x_{t+h-1}, x_{t+h-2}, \dots, x_{t+1}\}$ that minimizes the mean squared error $E(x_{t+h} - \sum_{j=1}^{h-1} \alpha_j x_{t+j})^2$.

$$\widehat{x}_t = \beta_1 x_{t+1} + \beta_2 x_{t+2} + \cdots + \beta_{h-1} x_{t+h-1}. \quad (3.54)$$

Because of stationarity, the coefficients, $\beta_1, \dots, \beta_{h-1}$ are the same in (3.53) and (3.54); we will explain this result in the next section.

Definition 3.9 *The partial autocorrelation function (PACF) of a stationary process, x_t , denoted ϕ_{hh} , for $h = 1, 2, \dots$, is*

$$\phi_{11} = \text{corr}(x_{t+1}, x_t) = \rho(1) \quad (3.55)$$

and

$$\phi_{hh} = \text{corr}(x_{t+h} - \widehat{x}_{t+h}, x_t - \widehat{x}_t), \quad h \geq 2. \quad (3.56)$$

Both $(x_{t+h} - \widehat{x}_{t+h})$ and $(x_t - \widehat{x}_t)$ are uncorrelated with $\{x_{t+1}, \dots, x_{t+h-1}\}$. The PACF, ϕ_{hh} , is the correlation between x_{t+h} and x_t with the linear dependence of $\{x_{t+1}, \dots, x_{t+h-1}\}$ on each, removed. If the process x_t is Gaussian, then $\phi_{hh} = \text{corr}(x_{t+h}, x_t \mid x_{t+1}, \dots, x_{t+h-1})$; that is, ϕ_{hh} is the correlation coefficient between x_{t+h} and x_t in the bivariate distribution of (x_{t+h}, x_t) conditional on $\{x_{t+1}, \dots, x_{t+h-1}\}$.

Example 3.14 The PACF of an AR(1)

Consider the PACF of the AR(1) process given by $x_t = \phi x_{t-1} + w_t$, with $|\phi| < 1$. By definition, $\phi_{11} = \rho(1) = \phi$. To calculate ϕ_{22} , consider the regression of x_{t+2} on x_{t+1} , say, $\widehat{x}_{t+2} = \beta x_{t+1}$. We choose β to minimize

$$E(x_{t+2} - \widehat{x}_{t+2})^2 = E(x_{t+2} - \beta x_{t+1})^2 = \gamma(0) - 2\beta\gamma(1) + \beta^2\gamma(0).$$

Taking derivatives with respect to β and setting the result equal to zero, we have $\beta = \gamma(1)/\gamma(0) = \rho(1) = \phi$. Next, consider the regression of x_t on x_{t+1} , say $\widehat{x}_t = \beta x_{t+1}$. We choose β to minimize

$$E(x_t - \widehat{x}_t)^2 = E(x_t - \beta x_{t+1})^2 = \gamma(0) - 2\beta\gamma(1) + \beta^2\gamma(0).$$

This is the same equation as before, so $\beta = \phi$. Hence,

$$\begin{aligned} \phi_{22} &= \text{corr}(x_{t+2} - \widehat{x}_{t+2}, x_t - \widehat{x}_t) = \text{corr}(x_{t+2} - \phi x_{t+1}, x_t - \phi x_{t+1}) \\ &= \text{corr}(w_{t+2}, x_t - \phi x_{t+1}) = 0 \end{aligned}$$

by causality. Thus, $\phi_{22} = 0$. In the next example, we will see that in this case, $\phi_{hh} = 0$ for all $h > 1$.

Example 3.15 The PACF of an AR(p)

The model implies $x_{t+h} = \sum_{j=1}^p \phi_j x_{t+h-j} + w_{t+h}$, where the roots of $\phi(z)$ are outside the unit circle. When $h > p$, the regression of x_{t+h} on $\{x_{t+1}, \dots, x_{t+h-1}\}$, is

$$\widehat{x}_{t+h} = \sum_{j=1}^p \phi_j x_{t+h-j}.$$




Fig. 3.5. The ACF and PACF of an AR(2) model with $\phi_1 = 1.5$ and $\phi_2 = -.75$.

We have not proved this obvious result yet, but we will prove it in the next section. Thus, when $h > p$,

$$\phi_{hh} = \text{corr}(x_{t+h} - \hat{x}_{t+h}, x_t - \hat{x}_t) = \text{corr}(w_{t+h}, x_t - \hat{x}_t) = 0,$$

because, by causality, $x_t - \hat{x}_t$ depends only on $\{w_{t+h-1}, w_{t+h-2}, \dots\}$; recall equation (3.54). When $h \leq p$, ϕ_{pp} is not zero, and $\phi_{11}, \dots, \phi_{p-1,p-1}$ are not necessarily zero. We will see later that, in fact, $\phi_{pp} = \phi_p$. Figure 3.5 shows the ACF and the PACF of the AR(2) model presented in Example 3.10.

To reproduce Figure 3.5 in R, use the following commands:

```
ACF = ARMAacf(ar=c(1.5,-.75), ma=0, 24)[-1]
PACF = ARMAacf(ar=c(1.5,-.75), ma=0, 24, pacf=TRUE)
par(mfrow=c(1,2))
plot(ACF, type="h", xlab="lag", ylim=c(-.8,1)); abline(h=0)
plot(PACF, type="h", xlab="lag", ylim=c(-.8,1)); abline(h=0)
```

Example 3.16 The PACF of an Invertible MA(q)

For an invertible MA(q), we can write $x_t = -\sum_{j=1}^{\infty} \pi_j x_{t-j} + w_t$. Moreover, no finite representation exists. From this result, it should be apparent that the PACF will never cut off, as in the case of an AR(p).

For an MA(1), $x_t = w_t + \theta w_{t-1}$, with $|\theta| < 1$, calculations similar to Example 3.14 will yield $\phi_{22} = -\theta^2/(1 + \theta^2 + \theta^4)$. For the MA(1) in general, we can show that

$$\phi_{hh} = -\frac{(-\theta)^h(1 - \theta^2)}{1 - \theta^{2(h+1)}}, \quad h \geq 1.$$

In the next section, we will discuss methods of calculating the PACF. The PACF for MA models behaves much like the ACF for AR models. Also, the PACF for AR

Table 3.1. Behavior of the ACF and PACF for ARMA Models

	$AR(p)$	$MA(q)$	$ARMA(p, q)$
ACF	Tails off	Cuts off after lag q	Tails off
PACF	Cuts off after lag p	Tails off	Tails off




Fig. 3.6. ACF and PACF of the Recruitment series. Note that the lag axes are in terms of season (12 months in this case).

models behaves much like the ACF for MA models. Because an invertible ARMA model has an infinite AR representation, the PACF will not cut off. We may summarize these results in [Table 3.1](#).

Example 3.17 Preliminary Analysis of the Recruitment Series

We consider the problem of modeling the Recruitment series shown in [Figure 1.5](#). There are 453 months of observed recruitment ranging over the years 1950-1987. The ACF and the PACF given in [Figure 3.6](#) are consistent with the behavior of an AR(2). The ACF has cycles corresponding roughly to a 12-month period, and the PACF has large values for $h = 1, 2$ and then is essentially zero for higher order lags. Based on [Table 3.1](#), these results suggest that a second-order ($p = 2$) autoregressive model might provide a good fit. Although we will discuss estimation in detail in §3.6, we ran a regression (see §2.2) using the data triplets $\{(x; z_1, z_2) : (x_3; x_2, x_1), (x_4; x_3, x_2), \dots, (x_{453}; x_{452}, x_{451})\}$ to fit a model of the form

$$x_t = \phi_0 + \phi_1 x_{t-1} + \phi_2 x_{t-2} + w_t$$

for $t = 3, 4, \dots, 453$. The values of the estimates were $\widehat{\phi}_0 = 6.74_{(1.11)}$, $\widehat{\phi}_1 = 1.35_{(.04)}$, $\widehat{\phi}_2 = -.46_{(.04)}$, and $\widehat{\sigma}_w^2 = 89.72$, where the estimated standard errors are in parentheses.

The following R code can be used for this analysis. We use the script `acf2` to print and plot the ACF and PACF; see Appendix R for details.

```
acf2(rec, 48)      # will produce values and a graphic
(regr = ar.ols(rec, order=2, demean=FALSE, intercept=TRUE))
regr$asy.se.coef  # standard errors of the estimates
```

3.5 Forecasting

In forecasting, the goal is to predict future values of a time series, x_{n+m} , $m = 1, 2, \dots$, based on the data collected to the present, $\mathbf{x} = \{x_n, x_{n-1}, \dots, x_1\}$. Throughout this section, we will assume x_t is stationary and the model parameters are known. The problem of forecasting when the model parameters are unknown will be discussed in the next section; also, see [Problem 3.26](#). The minimum mean square error predictor of x_{n+m} is

$$x_{n+m}^n = E(x_{n+m} \mid \mathbf{x}) \quad (3.57)$$

because the conditional expectation minimizes the mean square error

$$E [x_{n+m} - g(\mathbf{x})]^2, \quad (3.58)$$

where $g(\mathbf{x})$ is a function of the observations $\mathbf{x}$; see [Problem 3.14](#).

First, we will restrict attention to predictors that are linear functions of the data, that is, predictors of the form

$$x_{n+m}^n = \alpha_0 + \sum_{k=1}^n \alpha_k x_k, \quad (3.59)$$

where $\alpha_0, \alpha_1, \dots, \alpha_n$ are real numbers. We note that the α s depend on n and m , but for now we drop the dependence from the notation. For example, if $n = m = 1$, then x_2^1 is the one-step-ahead linear forecast of x_2 given x_1 . In terms of (3.59), $x_2^1 = \alpha_0 + \alpha_1 x_1$. But if $n = 2$, x_3^2 is the one-step-ahead linear forecast of x_3 given x_1 and x_2 . In terms of (3.59), $x_3^2 = \alpha_0 + \alpha_1 x_1 + \alpha_2 x_2$, and in general, the α s in x_2^1 and x_3^2 will be different.

Linear predictors of the form (3.59) that minimize the mean square prediction error (3.58) are called best linear predictors (BLPs). As we shall see, linear prediction depends only on the second-order moments of the process, which are easy to estimate from the data. Much of the material in this section is enhanced by the theoretical material presented in [Appendix B](#). For example, [Theorem B.3](#) states that if the process is Gaussian, minimum mean square error predictors and best linear predictors are the same. The following property, which is based on the Projection Theorem, [Theorem B.1](#) of Appendix B, is a key result.

Property 3.3 Best Linear Prediction for Stationary Processes

Given data $x_1, \dots, x_n$, the best linear predictor, $x_{n+m}^n = \alpha_0 + \sum_{k=1}^n \alpha_k x_k$, of x_{n+m} , for $m \geq 1$, is found by solving

$$E[(x_{n+m} - x_{n+m}^n) x_k] = 0, \quad k = 0, 1, \dots, n, \quad (3.60)$$

where $x_0 = 1$, for $\alpha_0, \alpha_1, \dots, \alpha_n$.

The equations specified in (3.60) are called the prediction equations, and they are used to solve for the coefficients $\{\alpha_0, \alpha_1, \dots, \alpha_n\}$. If $E(x_t) = \mu$, the first equation ($k = 0$) of (3.60) implies

$$E(x_{n+m}^n) = E(x_{n+m}) = \mu.$$

Thus, taking expectation in (3.59), we have

$$\mu = \alpha_0 + \sum_{k=1}^n \alpha_k \mu \quad \text{or} \quad \alpha_0 = \mu \left(1 - \sum_{k=1}^n \alpha_k\right).$$

Hence, the form of the BLP is

$$x_{n+m}^n = \mu + \sum_{k=1}^n \alpha_k (x_k - \mu).$$

Thus, until we discuss estimation, there is no loss of generality in considering the case that $\mu = 0$, in which case, $\alpha_0 = 0$.

First, consider one-step-ahead prediction. That is, given $\{x_1, \dots, x_n\}$, we wish to forecast the value of the time series at the next time point, x_{n+1} . The BLP of x_{n+1} is of the form

$$x_{n+1}^n = \phi_{n1}x_n + \phi_{n2}x_{n-1} + \dots + \phi_{nn}x_1, \quad (3.61)$$

where we now display the dependence of the coefficients on n ; in this case, α_k in (3.59) is $\phi_{n,n+1-k}$ in (3.61), for $k = 1, \dots, n$. Using **Property 3.3**, the coefficients $\{\phi_{n1}, \phi_{n2}, \dots, \phi_{nn}\}$ satisfy

$$E\left[\left(x_{n+1} - \sum_{j=1}^n \phi_{nj}x_{n+1-j}\right)x_{n+1-k}\right] = 0, \quad k = 1, \dots, n,$$

or

$$\sum_{j=1}^n \phi_{nj}\gamma(k-j) = \gamma(k), \quad k = 1, \dots, n. \quad (3.62)$$

The prediction equations (3.62) can be written in matrix notation as

$$\Gamma_n \boldsymbol{\phi}_n = \boldsymbol{\gamma}_n, \quad (3.63)$$

where $\Gamma_n = \{\gamma(k-j)\}_{j,k=1}^n$ is an $n \times n$ matrix, $\boldsymbol{\phi}_n = (\phi_{n1}, \dots, \phi_{nn})'$ is an $n \times 1$ vector, and $\boldsymbol{\gamma}_n = (\gamma(1), \dots, \gamma(n))'$ is an $n \times 1$ vector.

The matrix Γ_n is nonnegative definite. If Γ_n is singular, there are many solutions to (3.63), but, by the Projection Theorem (Theorem B.1), x_{n+1}^n is unique. If Γ_n is nonsingular, the elements of ϕ_n are unique, and are given by

$$\phi_n = \Gamma_n^{-1} \gamma_n. \quad (3.64)$$

For ARMA models, the fact that $\sigma_w^2 > 0$ and $\gamma(h) \rightarrow 0$ as $h \rightarrow \infty$ is enough to ensure that Γ_n is positive definite (Problem 3.12). It is sometimes convenient to write the one-step-ahead forecast in vector notation

$$x_{n+1}^n = \phi_n' \mathbf{x}, \quad (3.65)$$

where $\mathbf{x} = (x_n, x_{n-1}, \dots, x_1)'$.

The mean square one-step-ahead prediction error is

$$P_{n+1}^n = E(x_{n+1} - x_{n+1}^n)^2 = \gamma(0) - \gamma_n' \Gamma_n^{-1} \gamma_n. \quad (3.66)$$

To verify (3.66) using (3.64) and (3.65),

$$\begin{aligned} E(x_{n+1} - x_{n+1}^n)^2 &= E(x_{n+1} - \phi_n' \mathbf{x})^2 = E(x_{n+1} - \gamma_n' \Gamma_n^{-1} \mathbf{x})^2 \\ &= E(x_{n+1}^2 - 2\gamma_n' \Gamma_n^{-1} \mathbf{x} x_{n+1} + \gamma_n' \Gamma_n^{-1} \mathbf{x} \mathbf{x}' \Gamma_n^{-1} \gamma_n) \\ &= \gamma(0) - 2\gamma_n' \Gamma_n^{-1} \gamma_n + \gamma_n' \Gamma_n^{-1} \Gamma_n \Gamma_n^{-1} \gamma_n \\ &= \gamma(0) - \gamma_n' \Gamma_n^{-1} \gamma_n. \end{aligned}$$

Example 3.18 Prediction for an AR(2)

Suppose we have a causal AR(2) process $x_t = \phi_1 x_{t-1} + \phi_2 x_{t-2} + w_t$, and one observation x_1 . Then, using equation (3.64), the one-step-ahead prediction of x_2 based on x_1 is

$$x_2^1 = \phi_{11} x_1 = \frac{\gamma(1)}{\gamma(0)} x_1 = \rho(1) x_1.$$

Now, suppose we want the one-step-ahead prediction of x_3 based on two observations x_1 and x_2 ; i.e., $x_3^2 = \phi_{21} x_2 + \phi_{22} x_1$. We could use (3.62)

$$\begin{aligned} \phi_{21} \gamma(0) + \phi_{22} \gamma(1) &= \gamma(1) \\ \phi_{21} \gamma(1) + \phi_{22} \gamma(0) &= \gamma(2) \end{aligned}$$

to solve for ϕ_{21} and ϕ_{22} , or use the matrix form in (3.64) and solve

$$\begin{pmatrix} \phi_{21} \\ \phi_{22} \end{pmatrix} = \begin{pmatrix} \gamma(0) & \gamma(1) \\ \gamma(1) & \gamma(0) \end{pmatrix}^{-1} \begin{pmatrix} \gamma(1) \\ \gamma(2) \end{pmatrix},$$

but, it should be apparent from the model that $x_3^2 = \phi_1 x_2 + \phi_2 x_1$. Because $\phi_1 x_2 + \phi_2 x_1$ satisfies the prediction equations (3.60),

$$E\{[x_3 - (\phi_1 x_2 + \phi_2 x_1)]x_1\} = E(w_3 x_1) = 0,$$

$$E\{[x_3 - (\phi_1 x_2 + \phi_2 x_1)]x_2\} = E(w_3 x_2) = 0,$$

it follows that, indeed, $x_3^2 = \phi_1 x_2 + \phi_2 x_1$, and by the uniqueness of the coefficients in this case, that $\phi_{21} = \phi_1$ and $\phi_{22} = \phi_2$. Continuing in this way, it is easy to verify that, for $n \geq 2$,

$$x_{n+1}^n = \phi_1 x_n + \phi_2 x_{n-1}.$$

That is, $\phi_{n1} = \phi_1$, $\phi_{n2} = \phi_2$, and $\phi_{nj} = 0$, for $j = 3, 4, \dots, n$.

From [Example 3.18](#), it should be clear ([Problem 3.42](#)) that, if the time series is a causal AR(p) process, then, for $n \geq p$,

$$x_{n+1}^n = \phi_1 x_n + \phi_2 x_{n-1} + \dots + \phi_p x_{n-p+1}. \quad (3.67)$$

For ARMA models in general, the prediction equations will not be as simple as the pure AR case. In addition, for n large, the use of (3.64) is prohibitive because it requires the inversion of a large matrix. There are, however, iterative solutions that do not require any matrix inversion. In particular, we mention the recursive solution due to Levinson (1947) and Durbin (1960).

Property 3.4 The Durbin–Levinson Algorithm

Equations (3.64) and (3.66) can be solved iteratively as follows:

$$\phi_{00} = 0, \quad P_1^0 = \gamma(0). \quad (3.68)$$

For $n \geq 1$,

$$\phi_{nn} = \frac{\rho(n) - \sum_{k=1}^{n-1} \phi_{n-1,k} \rho(n-k)}{1 - \sum_{k=1}^{n-1} \phi_{n-1,k} \rho(k)}, \quad P_{n+1}^n = P_n^{n-1} (1 - \phi_{nn}^2), \quad (3.69)$$

where, for $n \geq 2$,

$$\phi_{nk} = \phi_{n-1,k} - \phi_{nn} \phi_{n-1,n-k}, \quad k = 1, 2, \dots, n-1. \quad (3.70)$$

The proof of [Property 3.4](#) is left as an exercise; see [Problem 3.13](#).

Example 3.19 Using the Durbin–Levinson Algorithm

To use the algorithm, start with $\phi_{00} = 0$, $P_1^0 = \gamma(0)$. Then, for $n = 1$,

$$\phi_{11} = \rho(1), \quad P_2^1 = \gamma(0)[1 - \phi_{11}^2].$$

For $n = 2$,

$$\begin{aligned} \phi_{22} &= \frac{\rho(2) - \phi_{11} \rho(1)}{1 - \phi_{11} \rho(1)}, \quad \phi_{21} = \phi_{11} - \phi_{22} \phi_{11}, \\ P_3^2 &= P_2^1 [1 - \phi_{22}^2] = \gamma(0)[1 - \phi_{11}^2][1 - \phi_{22}^2]. \end{aligned}$$

For $n = 3$,

$$\begin{aligned}\phi_{33} &= \frac{\rho(3) - \phi_{21} \rho(2) - \phi_{22} \rho(1)}{1 - \phi_{21} \rho(1) - \phi_{22} \rho(2)}, \\ \phi_{32} &= \phi_{22} - \phi_{33} \phi_{21}, \quad \phi_{31} = \phi_{21} - \phi_{33} \phi_{22}, \\ P_4^3 &= P_3^2 [1 - \phi_{33}^2] = \gamma(0) [1 - \phi_{11}^2] [1 - \phi_{22}^2] [1 - \phi_{33}^2],\end{aligned}$$

and so on. Note that, in general, the standard error of the one-step-ahead forecast is the square root of

$$P_{n+1}^n = \gamma(0) \prod_{j=1}^n [1 - \phi_{jj}^2]. \quad (3.71)$$

An important consequence of the Durbin–Levinson algorithm is (see **Problem 3.13**) as follows.

Property 3.5 Iterative Solution for the PACF

The PACF of a stationary process x_t , can be obtained iteratively via (3.69) as ϕ_{nn} , for $n = 1, 2, \dots$

Using **Property 3.5** and putting $n = p$ in (3.61) and (3.67), it follows that for an AR(p) model,

$$\begin{aligned}x_{p+1}^p &= \phi_{p1} x_p + \phi_{p2} x_{p-1} + \dots + \phi_{pp} x_1 \\ &= \phi_1 x_p + \phi_2 x_{p-1} + \dots + \phi_p x_1.\end{aligned} \quad (3.72)$$

Result (3.72) shows that for an AR(p) model, the partial autocorrelation coefficient at lag p , ϕ_{pp} , is also the last coefficient in the model, ϕ_p , as was claimed in **Example 3.15**.

Example 3.20 The PACF of an AR(2)

We will use the results of **Example 3.19** and **Property 3.5** to calculate the first three values, ϕ_{11} , ϕ_{22} , ϕ_{33} , of the PACF. Recall from **Example 3.9** that $\rho(h) - \phi_1 \rho(h-1) - \phi_2 \rho(h-2) = 0$ for $h \geq 1$. When $h = 1, 2, 3$, we have $\rho(1) = \phi_1 / (1 - \phi_2)$, $\rho(2) = \phi_1 \rho(1) + \phi_2$, $\rho(3) - \phi_1 \rho(2) - \phi_2 \rho(1) = 0$. Thus,

$$\begin{aligned}\phi_{11} &= \rho(1) = \frac{\phi_1}{1 - \phi_2} \\ \phi_{22} &= \frac{\rho(2) - \rho(1)^2}{1 - \rho(1)^2} = \frac{\left[\phi_1 \left(\frac{\phi_1}{1 - \phi_2} \right) + \phi_2 \right] - \left(\frac{\phi_1}{1 - \phi_2} \right)^2}{1 - \left(\frac{\phi_1}{1 - \phi_2} \right)^2} = \phi_2 \\ \phi_{21} &= \rho(1) [1 - \phi_2] = \phi_1 \\ \phi_{33} &= \frac{\rho(3) - \phi_1 \rho(2) - \phi_2 \rho(1)}{1 - \phi_1 \rho(1) - \phi_2 \rho(2)} = 0.\end{aligned}$$

Notice that, as shown in (3.72), $\phi_{22} = \phi_2$ for an AR(2) model.

So far, we have concentrated on one-step-ahead prediction, but **Property 3.3** allows us to calculate the BLP of x_{n+m} for any $m \geq 1$. Given data, $\{x_1, \dots, x_n\}$, the m -step-ahead predictor is

$$x_{n+m}^n = \phi_{n1}^{(m)} x_n + \phi_{n2}^{(m)} x_{n-1} + \cdots + \phi_{nn}^{(m)} x_1, \quad (3.73)$$

where $\{\phi_{n1}^{(m)}, \phi_{n2}^{(m)}, \dots, \phi_{nn}^{(m)}\}$ satisfy the prediction equations,

$$\sum_{j=1}^n \phi_{nj}^{(m)} E(x_{n+1-j} x_{n+1-k}) = E(x_{n+m} x_{n+1-k}), \quad k = 1, \dots, n,$$

or

$$\sum_{j=1}^n \phi_{nj}^{(m)} \gamma(k-j) = \gamma(m+k-1), \quad k = 1, \dots, n. \quad (3.74)$$

The prediction equations can again be written in matrix notation as

$$\Gamma_n \boldsymbol{\phi}_n^{(m)} = \boldsymbol{\gamma}_n^{(m)}, \quad (3.75)$$

where $\boldsymbol{\gamma}_n^{(m)} = (\gamma(m), \dots, \gamma(m+n-1))'$, and $\boldsymbol{\phi}_n^{(m)} = (\phi_{n1}^{(m)}, \dots, \phi_{nn}^{(m)})'$ are $n \times 1$ vectors.

The mean square m -step-ahead prediction error is

$$P_{n+m}^n = E(x_{n+m} - x_{n+m}^n)^2 = \gamma(0) - \boldsymbol{\gamma}_n^{(m)'} \Gamma_n^{-1} \boldsymbol{\gamma}_n^{(m)}. \quad (3.76)$$

Another useful algorithm for calculating forecasts was given by Brockwell and Davis (1991, Chapter 5). This algorithm follows directly from applying the projection theorem ([Theorem B.1](#)) to the innovations, $x_t - x_t^{t-1}$, for $t = 1, \dots, n$, using the fact that the innovations $x_t - x_t^{t-1}$ and $x_s - x_s^{s-1}$ are uncorrelated for $s \neq t$ (see [Problem 3.43](#)). We present the case in which x_t is a mean-zero stationary time series.

Property 3.6 The Innovations Algorithm

The one-step-ahead predictors, x_{t+1}^t , and their mean-squared errors, P_{t+1}^t , can be calculated iteratively as

$$x_1^0 = 0, \quad P_1^0 = \gamma(0)$$

$$x_{t+1}^t = \sum_{j=1}^t \theta_{tj} (x_{t+1-j} - x_{t+1-j}^{t-j}), \quad t = 1, 2, \dots \quad (3.77)$$

$$P_{t+1}^t = \gamma(0) - \sum_{j=0}^{t-1} \theta_{t,t-j}^2 P_{j+1}^j \quad t = 1, 2, \dots, \quad (3.78)$$

where, for $j = 0, 1, \dots, t-1$,

$$\theta_{t,t-j} = \left(\gamma(t-j) - \sum_{k=0}^{j-1} \theta_{j,j-k} \theta_{t,t-k} P_{k+1}^k \right) / P_{j+1}^j. \quad (3.79)$$

Given data $x_1, \dots, x_n$, the innovations algorithm can be calculated successively for $t = 1$, then $t = 2$ and so on, in which case the calculation of x_{n+1}^n and P_{n+1}^n is made at the final step $t = n$. The m -step-ahead predictor and its mean-square error based on the innovations algorithm ([Problem 3.43](#)) are given by

$$x_{n+m}^n = \sum_{j=m}^{n+m-1} \theta_{n+m-1,j} (x_{n+m-j} - x_{n+m-j}^{n+m-j-1}), \quad (3.80)$$

$$P_{n+m}^n = \gamma(0) - \sum_{j=m}^{n+m-1} \theta_{n+m-1,j}^2 P_{n+m-j}^{n+m-j-1}, \quad (3.81)$$

where the $\theta_{n+m-1,j}$ are obtained by continued iteration of (3.79).

Example 3.21 Prediction for an MA(1)

The innovations algorithm lends itself well to prediction for moving average processes. Consider an MA(1) model, $x_t = w_t + \theta w_{t-1}$. Recall that $\gamma(0) = (1 + \theta^2)\sigma_w^2$, $\gamma(1) = \theta\sigma_w^2$, and $\gamma(h) = 0$ for $h > 1$. Then, using [Property 3.6](#), we have

$$\theta_{n1} = \theta\sigma_w^2/P_n^{n-1}$$

$$\theta_{nj} = 0, \quad j = 2, \dots, n$$

$$P_1^0 = (1 + \theta^2)\sigma_w^2$$

$$P_{n+1}^n = (1 + \theta^2 - \theta\theta_{n1})\sigma_w^2.$$

Finally, from (3.77), the one-step-ahead predictor is

$$x_{n+1}^n = \theta (x_n - x_n^{n-1}) \sigma_w^2 / P_n^{n-1}.$$

FORECASTING ARMA PROCESSES

The general prediction equations (3.60) provide little insight into forecasting for ARMA models in general. There are a number of different ways to express these forecasts, and each aids in understanding the special structure of ARMA prediction. Throughout, we assume x_t is a causal and invertible ARMA(p, q) process, $\phi(B)x_t = \theta(B)w_t$, where $w_t \sim \text{iid } N(0, \sigma_w^2)$. In the non-zero mean case, $E(x_t) = \mu_x$, simply replace x_t with $x_t - \mu_x$ in the model. First, we consider two types of forecasts. We write x_{n+m}^n to mean the minimum mean square error predictor of x_{n+m} based on the data $\{x_n, \dots, x_1\}$, that is,

$$x_{n+m}^n = E(x_{n+m} \mid x_n, \dots, x_1).$$

For ARMA models, it is easier to calculate the predictor of x_{n+m} , assuming we have the complete history of the process $\{x_n, x_{n-1}, \dots, x_1, x_0, x_{-1}, \dots\}$. We will denote the predictor of x_{n+m} *based on the infinite past* as

$$\tilde{x}_{n+m} = E(x_{n+m} \mid x_n, x_{n-1}, \dots, x_1, x_0, x_{-1}, \dots).$$

In general, x_{n+m}^n and $\tilde{x}_{n+m}$ are not the same, but the idea here is that, for large samples, $\tilde{x}_{n+m}$ will provide a good approximation to x_{n+m}^n .

Now, write x_{n+m} in its causal and invertible forms:

$$x_{n+m} = \sum_{j=0}^{\infty} \psi_j w_{n+m-j}, \quad \psi_0 = 1 \quad (3.82)$$

$$w_{n+m} = \sum_{j=0}^{\infty} \pi_j x_{n+m-j}, \quad \pi_0 = 1. \quad (3.83)$$

Then, taking conditional expectations in (3.82), we have

$$\tilde{x}_{n+m} = \sum_{j=0}^{\infty} \psi_j \tilde{w}_{n+m-j} = \sum_{j=m}^{\infty} \psi_j w_{n+m-j}, \quad (3.84)$$

because, by causality and invertibility,

$$\tilde{w}_t = E(w_t \mid x_n, x_{n-1}, \dots, x_0, x_{-1}, \dots) = \begin{cases} 0 & t > n \\ w_t & t \leq n. \end{cases}$$

Similarly, taking conditional expectations in (3.83), we have

$$0 = \tilde{x}_{n+m} + \sum_{j=1}^{\infty} \pi_j \tilde{x}_{n+m-j},$$

or

$$\tilde{x}_{n+m} = - \sum_{j=1}^{m-1} \pi_j \tilde{x}_{n+m-j} - \sum_{j=m}^{\infty} \pi_j x_{n+m-j}, \quad (3.85)$$

using the fact $E(x_t \mid x_n, x_{n-1}, \dots, x_0, x_{-1}, \dots) = x_t$, for $t \leq n$. Prediction is accomplished recursively using (3.85), starting with the one-step-ahead predictor, $m = 1$, and then continuing for $m = 2, 3, \dots$. Using (3.84), we can write

$$x_{n+m} - \tilde{x}_{n+m} = \sum_{j=0}^{m-1} \psi_j w_{n+m-j},$$

so the mean-square prediction error can be written as

$$P_{n+m}^n = E(x_{n+m} - \tilde{x}_{n+m})^2 = \sigma_w^2 \sum_{j=0}^{m-1} \psi_j^2. \quad (3.86)$$

Also, we note, for a fixed sample size, n , the prediction errors are correlated. That is, for $k \geq 1$,

$$E\{(x_{n+m} - \tilde{x}_{n+m})(x_{n+m+k} - \tilde{x}_{n+m+k})\} = \sigma_w^2 \sum_{j=0}^{m-1} \psi_j \psi_{j+k}. \quad (3.87)$$

Example 3.22 Long-Range Forecasts

Consider forecasting an ARMA process with mean μ_x . Replacing x_{n+m} with $x_{n+m} - \mu_x$ in (3.82), and taking conditional expectation as in (3.84), we deduce that the m -step-ahead forecast can be written as

$$\tilde{x}_{n+m} = \mu_x + \sum_{j=m}^{\infty} \psi_j w_{n+m-j}. \quad (3.88)$$

Noting that the ψ -weights dampen to zero exponentially fast, it is clear that

$$\tilde{x}_{n+m} \rightarrow \mu_x \quad (3.89)$$

exponentially fast (in the mean square sense) as $m \rightarrow \infty$. Moreover, by (3.86), the mean square prediction error

$$P_{n+m}^n \rightarrow \sigma_w^2 \sum_{j=0}^{\infty} \psi_j^2 = \gamma_x(0) = \sigma_x^2, \quad (3.90)$$

exponentially fast as $m \rightarrow \infty$; recall (3.45).

It should be clear from (3.89) and (3.90) that ARMA forecasts quickly settle to the mean with a constant prediction error as the forecast horizon, m , grows. This effect can be seen in Figure 3.7 on page 115 where the Recruitment series is forecast for 24 months; see Example 3.24.

When n is small, the general prediction equations (3.60) can be used easily. When n is large, we would use (3.85) by truncating, because we do not observe $x_0, x_{-1}, x_{-2}, \dots$, and only the data $x_1, x_2, \dots, x_n$ are available. In this case, we can truncate (3.85) by setting $\sum_{j=n+m}^{\infty} \pi_j x_{n+m-j} = 0$. The truncated predictor is then written as

$$\tilde{x}_{n+m}^n = - \sum_{j=1}^{m-1} \pi_j \tilde{x}_{n+m-j}^n - \sum_{j=m}^{n+m-1} \pi_j x_{n+m-j}, \quad (3.91)$$

which is also calculated recursively, $m = 1, 2, \dots$. The mean square prediction error, in this case, is approximated using (3.86).

For AR(p) models, and when $n > p$, equation (3.67) yields the exact predictor, x_{n+m}^n , of x_{n+m} , and there is no need for approximations. That is, for $n > p$, $\tilde{x}_{n+m}^n = x_{n+m}^n = x_{n+m}^n$. Also, in this case, the one-step-ahead prediction error is $E(x_{n+1} - x_{n+1}^n)^2 = \sigma_w^2$. For pure MA(q) or ARMA(p, q) models, truncated prediction has a fairly nice form.

Property 3.7 Truncated Prediction for ARMA

For ARMA(p, q) models, the truncated predictors for $m = 1, 2, \dots$, are

$$\tilde{x}_{n+m}^n = \phi_1 \tilde{x}_{n+m-1}^n + \dots + \phi_p \tilde{x}_{n+m-p}^n + \theta_1 \tilde{w}_{n+m-1}^n + \dots + \theta_q \tilde{w}_{n+m-q}^n, \quad (3.92)$$

where $\tilde{x}_t^n = x_t$ for $1 \leq t \leq n$ and $\tilde{x}_t^n = 0$ for $t \leq 0$. The truncated prediction errors are given by: $\tilde{w}_t^n = 0$ for $t \leq 0$ or $t > n$, and

$$\tilde{w}_t^n = \phi(B) \tilde{x}_t^n - \theta_1 \tilde{w}_{t-1}^n - \dots - \theta_q \tilde{w}_{t-q}^n$$

for $1 \leq t \leq n$.

Example 3.23 Forecasting an ARMA(1, 1) Series

Given data $x_1, \dots, x_n$, for forecasting purposes, write the model as

$$x_{n+1} = \phi x_n + w_{n+1} + \theta w_n.$$

Then, based on (3.92), the one-step-ahead truncated forecast is

$$\tilde{x}_{n+1}^n = \phi x_n + 0 + \theta \tilde{w}_n^n.$$

For $m \geq 2$, we have

$$\tilde{x}_{n+m}^n = \phi \tilde{x}_{n+m-1}^n,$$

which can be calculated recursively, $m = 2, 3, \dots$

To calculate $\tilde{w}_n^n$, which is needed to initialize the successive forecasts, the model can be written as $w_t = x_t - \phi x_{t-1} - \theta w_{t-1}$ for $t = 1, \dots, n$. For truncated forecasting using (3.92), put $\tilde{w}_0^n = 0$, $x_0 = 0$, and then iterate the errors forward in time

$$\tilde{w}_t^n = x_t - \phi x_{t-1} - \theta \tilde{w}_{t-1}^n, \quad t = 1, \dots, n.$$

The approximate forecast variance is computed from (3.86) using the ψ -weights determined as in Example 3.11. In particular, the ψ -weights satisfy $\psi_j = (\phi + \theta)\phi^{j-1}$, for $j \geq 1$. This result gives

$$P_{n+m}^n = \sigma_w^2 \left[1 + (\phi + \theta)^2 \sum_{j=1}^{m-1} \phi^{2(j-1)} \right] = \sigma_w^2 \left[1 + \frac{(\phi + \theta)^2 (1 - \phi^{2(m-1)})}{(1 - \phi^2)} \right].$$

To assess the precision of the forecasts, prediction intervals are typically calculated along with the forecasts. In general, $(1 - \alpha)$ prediction intervals are of the form

$$x_{n+m}^n \pm c_{\alpha/2} \sqrt{P_{n+m}^n}, \quad (3.93)$$

where $c_{\alpha/2}$ is chosen to get the desired degree of confidence. For example, if the process is Gaussian, then choosing $c_{\alpha/2} = 2$ will yield an approximate 95% prediction interval for x_{n+m} . If we are interested in establishing prediction intervals over more than one time period, then $c_{\alpha/2}$ should be adjusted appropriately, for example, by using Bonferroni's inequality [see (4.58) in Chapter 4 or Johnson and Wichern, 1992, Chapter 5].

Example 3.24 Forecasting the Recruitment Series

Using the parameter estimates as the actual parameter values, Figure 3.7 shows the result of forecasting the Recruitment series given in Example 3.17 over a 24-month horizon, $m = 1, 2, \dots, 24$. The actual forecasts are calculated as

$$x_{n+m}^n = 6.74 + 1.35x_{n+m-1}^n - .46x_{n+m-2}^n$$

for $n = 453$ and $m = 1, 2, \dots, 12$. Recall that $x_t^s = x_t$ when $t \leq s$. The forecasts errors P_{n+m}^n are calculated using (3.86). Recall that $\hat{\sigma}_w^2 = 89.72$, and using (3.40)




Fig. 3.7. Twenty-four month forecasts for the Recruitment series. The actual data shown are from about January 1980 to September 1987, and then the forecasts plus and minus one standard error are displayed.

from **Example 3.11**, we have $\psi_j = 1.35\psi_{j-1} - .46\psi_{j-2}$ for $j \geq 2$, where $\psi_0 = 1$ and $\psi_1 = 1.35$. Thus, for $n = 453$,

$$\begin{aligned} P_{n+1}^n &= 89.72, \\ P_{n+2}^n &= 89.72(1 + 1.35^2), \\ P_{n+3}^n &= 89.72(1 + 1.35^2 + [1.35^2 - .46]^2), \end{aligned}$$

and so on.

Note how the forecast levels off quickly and the prediction intervals are wide, even though in this case the forecast limits are only based on one standard error; that is, $x_{n+m}^n \pm \sqrt{P_{n+m}^n}$.

To reproduce the analysis and **Figure 3.7**, use the following commands:

```
regr = ar.ols(rec, order=2, demean=FALSE, intercept=TRUE)
fore = predict(regr, n.ahead=24)
ts.plot(rec, fore$pred, col=1:2, xlim=c(1980,1990), ylab="Recruitment")
lines(fore$pred, type="p", col=2)
lines(fore$pred+fore$se, lty="dashed", col=4)
lines(fore$pred-fore$se, lty="dashed", col=4)
```

We complete this section with a brief discussion of backcasting. In backcasting, we want to predict x_{1-m} , for $m = 1, 2, \dots$, based on the data $\{x_1, \dots, x_n\}$. Write the backcast as

$$x_{1-m}^n = \sum_{j=1}^n \alpha_j x_j. \quad (3.94)$$

Analogous to (3.74), the prediction equations (assuming $\mu_x = 0$) are

$$\sum_{j=1}^n \alpha_j E(x_j x_k) = E(x_{1-m} x_k), \quad k = 1, \dots, n, \quad (3.95)$$

or

$$\sum_{j=1}^n \alpha_j \gamma(k-j) = \gamma(m+k-1), \quad k = 1, \dots, n. \quad (3.96)$$

These equations are precisely the prediction equations for forward prediction. That is, $\alpha_j \equiv \phi_{nj}^{(m)}$, for $j = 1, \dots, n$, where the $\phi_{nj}^{(m)}$ are given by (3.75). Finally, the backcasts are given by

$$x_{1-m}^n = \phi_{n1}^{(m)} x_1 + \dots + \phi_{nn}^{(m)} x_n, \quad m = 1, 2, \dots \quad (3.97)$$

Example 3.25 Backcasting an ARMA(1, 1)

Consider an ARMA(1, 1) process, $x_t = \phi x_{t-1} + \theta w_{t-1} + w_t$; we will call this the *forward model*. We have just seen that best linear prediction backward in time is the same as best linear prediction forward in time for stationary models. Assuming the models are Gaussian, we also have that minimum mean square error prediction backward in time is the same as forward in time for ARMA models.⁴ Thus, the process can equivalently be generated by the *backward model*,

$$x_t = \phi x_{t+1} + \theta v_{t+1} + v_t,$$

where $\{v_t\}$ is a Gaussian white noise process with variance σ_w^2 . We may write $x_t = \sum_{j=0}^{\infty} \psi_j v_{t+j}$, where $\psi_0 = 1$; this means that x_t is uncorrelated with $\{v_{t-1}, v_{t-2}, \dots\}$, in analogy to the forward model.

Given data $\{x_1, \dots, x_n\}$, truncate $v_n^n = E(v_n | x_1, \dots, x_n)$ to zero and then iterate backward. That is, put $\tilde{v}_n^n = 0$, as an initial approximation, and then generate the errors backward

$$\tilde{v}_t^n = x_t - \phi x_{t+1} - \theta \tilde{v}_{t+1}^n, \quad t = (n-1), (n-2), \dots, 1.$$

Then,

$$\tilde{x}_0^n = \phi x_1 + \theta \tilde{v}_1^n + \tilde{v}_0^n = \phi x_1 + \theta \tilde{v}_1^n,$$

because $\tilde{v}_t^n = 0$ for $t \leq 0$. Continuing, the general truncated backcasts are given by

$$\tilde{x}_{1-m}^n = \phi \tilde{x}_{2-m}^n, \quad m = 2, 3, \dots$$

To backcast data in R, simply reverse the data, fit the model and predict. In the following, we backcasted a simulated ARMA(1,1) process; see Figure 3.8.

```
set.seed(90210)
x = arima.sim(list(order = c(1,0,1), ar = .9, ma=.5), n = 100)
xr = replace(x, TRUE, rev(x)) # xr is the reversed data
pxr = predict(arima(xr, order=c(1,0,1)), 10) # predict the reversed data
pxrp = rev(pxr$pred) # reorder the predictors (for plotting)
pxrse = rev(pxr$se) # reorder the SEs
nx = ts(c(pxrp, x), start=-9) # attach the backcasts to the data
plot(nx, ylab=expression(X[~t]), main='Backcasting')
lines(-9:0, nx[1:10], col=2, type='o', pch=20)
lines(-9:0, nx[1:10] + pxrse, col=4, lty=2)
lines(-9:0, nx[1:10] - pxrse, col=4, lty=2)
```

⁴ In the stationary Gaussian case, (a) the distribution of $\{x_{n+1}, x_n, \dots, x_1\}$ is the same as (b) the distribution of $\{x_0, x_1, \dots, x_n\}$. In forecasting we use (a) to obtain $E(x_{n+1} | x_n, \dots, x_1)$; in backcasting we use (b) to obtain $E(x_0 | x_1, \dots, x_n)$. Because (a) and (b) are the same, the two problems are equivalent.




Fig. 3.8. Display for Example 3.25; backcasts from a simulated ARMA(1, 1).

3.6 Estimation

Throughout this section, we assume we have n observations, $x_1, \dots, x_n$, from a causal and invertible Gaussian ARMA(p, q) process in which, initially, the order parameters, p and q , are known. Our goal is to estimate the parameters, $\phi_1, \dots, \phi_p, \theta_1, \dots, \theta_q$, and σ_w^2 . We will discuss the problem of determining p and q later in this section.

We begin with method of moments estimators. The idea behind these estimators is that of equating population moments to sample moments and then solving for the parameters in terms of the sample moments. We immediately see that, if $E(x_t) = \mu$, then the method of moments estimator of μ is the sample average, $\bar{x}$. Thus, while discussing method of moments, we will assume $\mu = 0$. Although the method of moments can produce good estimators, they can sometimes lead to suboptimal estimators. We first consider the case in which the method leads to optimal (efficient) estimators, that is, AR(p) models.

When the process is AR(p),

$$x_t = \phi_1 x_{t-1} + \dots + \phi_p x_{t-p} + w_t,$$

the first $p + 1$ equations of (3.47) and (3.48) lead to the following:

Definition 3.10 The Yule–Walker equations are given by

$$\gamma(h) = \phi_1 \gamma(h-1) + \dots + \phi_p \gamma(h-p), \quad h = 1, 2, \dots, p, \quad (3.98)$$

$$\sigma_w^2 = \gamma(0) - \phi_1 \gamma(1) - \dots - \phi_p \gamma(p). \quad (3.99)$$

In matrix notation, the Yule–Walker equations are

$$\Gamma_p \boldsymbol{\phi} = \boldsymbol{\gamma}_p, \quad \sigma_w^2 = \gamma(0) - \boldsymbol{\phi}' \boldsymbol{\gamma}_p, \quad (3.100)$$

where $\Gamma_p = \{\gamma(k-j)\}_{j,k=1}^p$ is a $p \times p$ matrix, $\boldsymbol{\phi} = (\phi_1, \dots, \phi_p)'$ is a $p \times 1$ vector, and $\boldsymbol{\gamma}_p = (\gamma(1), \dots, \gamma(p))'$ is a $p \times 1$ vector. Using the method of moments, we replace $\gamma(h)$ in (3.100) by $\hat{\gamma}(h)$ [see equation (1.34)] and solve

$$\widehat{\boldsymbol{\phi}} = \widehat{\Gamma}_p^{-1} \widehat{\boldsymbol{\gamma}}_p, \quad \widehat{\sigma}_w^2 = \widehat{\gamma}(0) - \widehat{\boldsymbol{\gamma}}_p' \widehat{\Gamma}_p^{-1} \widehat{\boldsymbol{\gamma}}_p. \quad (3.101)$$

These estimators are typically called the Yule–Walker estimators. For calculation purposes, it is sometimes more convenient to work with the sample ACF. By factoring $\widehat{\gamma}(0)$ in (3.101), we can write the Yule–Walker estimates as

$$\widehat{\boldsymbol{\phi}} = \widehat{\mathbf{R}}_p^{-1} \widehat{\boldsymbol{\rho}}_p, \quad \widehat{\sigma}_w^2 = \widehat{\gamma}(0) \left[1 - \widehat{\boldsymbol{\rho}}_p' \widehat{\mathbf{R}}_p^{-1} \widehat{\boldsymbol{\rho}}_p \right], \quad (3.102)$$

where $\widehat{\mathbf{R}}_p = \{\widehat{\rho}(k-j)\}_{j,k=1}^p$ is a $p \times p$ matrix and $\widehat{\boldsymbol{\rho}}_p = (\widehat{\rho}(1), \dots, \widehat{\rho}(p))'$ is a $p \times 1$ vector.

For AR(p) models, if the sample size is large, the Yule–Walker estimators are approximately normally distributed, and $\widehat{\sigma}_w^2$ is close to the true value of σ_w^2 . We state these results in **Property 3.8**; for details, see Appendix B, §B.3.

Property 3.8 Large Sample Results for Yule–Walker Estimators

The asymptotic ($n \rightarrow \infty$) behavior of the Yule–Walker estimators in the case of causal AR(p) processes is as follows:

$$\sqrt{n} (\widehat{\boldsymbol{\phi}} - \boldsymbol{\phi}) \xrightarrow{d} N(\mathbf{0}, \sigma_w^2 \Gamma_p^{-1}), \quad \widehat{\sigma}_w^2 \xrightarrow{P} \sigma_w^2. \quad (3.103)$$

The Durbin–Levinson algorithm, (3.68)–(3.70), can be used to calculate $\widehat{\boldsymbol{\phi}}$ without inverting $\widehat{\Gamma}_p$ or $\widehat{\mathbf{R}}_p$, by replacing $\gamma(h)$ by $\widehat{\gamma}(h)$ in the algorithm. In running the algorithm, we will iteratively calculate the $h \times 1$ vector, $\widehat{\boldsymbol{\phi}}_h = (\widehat{\phi}_{h1}, \dots, \widehat{\phi}_{hh})'$, for $h = 1, 2, \dots$. Thus, in addition to obtaining the desired forecasts, the Durbin–Levinson algorithm yields $\widehat{\phi}_{hh}$, the sample PACF. Using (3.103), we can show the following property.

Property 3.9 Large Sample Distribution of the PACF

For a causal AR(p) process, asymptotically ($n \rightarrow \infty$),

$$\sqrt{n} \widehat{\phi}_{hh} \xrightarrow{d} N(0, 1), \quad \text{for } h > p. \quad (3.104)$$

Example 3.26 Yule–Walker Estimation for an AR(2) Process

The data shown in **Figure 3.4** were $n = 144$ simulated observations from the AR(2) model

$$x_t = 1.5x_{t-1} - .75x_{t-2} + w_t,$$

where $w_t \sim \text{iid } N(0, 1)$. For these data, $\widehat{\gamma}(0) = 8.903$, $\widehat{\rho}(1) = .849$, and $\widehat{\rho}(2) = .519$. Thus,

$$\widehat{\boldsymbol{\phi}} = \begin{pmatrix} \widehat{\phi}_1 \\ \widehat{\phi}_2 \end{pmatrix} = \begin{bmatrix} 1 & .849 \\ .849 & 1 \end{bmatrix}^{-1} \begin{pmatrix} .849 \\ .519 \end{pmatrix} = \begin{pmatrix} 1.463 \\ -.723 \end{pmatrix}$$

and

$$\widehat{\sigma}_w^2 = 8.903 \left[1 - (.849, .519) \begin{pmatrix} 1.463 \\ -.723 \end{pmatrix} \right] = 1.187.$$

By **Property 3.8**, the asymptotic variance–covariance matrix of $\widehat{\phi}$,

$$\frac{1}{144} \frac{1.187}{8.903} \begin{bmatrix} 1 & .849 \\ .849 & 1 \end{bmatrix}^{-1} = \begin{bmatrix} .058^2 & -.003 \\ -.003 & .058^2 \end{bmatrix},$$

can be used to get confidence regions for, or make inferences about $\widehat{\phi}$ and its components. For example, an approximate 95% confidence interval for ϕ_2 is $-.723 \pm 2(.058)$, or $(-.838, -.608)$, which contains the true value of $\phi_2 = -.75$.

For these data, the first three sample partial autocorrelations are $\widehat{\phi}_{11} = \widehat{\rho}(1) = .849$, $\widehat{\phi}_{22} = \widehat{\phi}_2 = -.721$, and $\widehat{\phi}_{33} = -.085$. According to **Property 3.9**, the asymptotic standard error of ϕ_{33} is $1/\sqrt{144} = .083$, and the observed value, $-.085$, is about only one standard deviation from $\phi_{33} = 0$.

Example 3.27 Yule–Walker Estimation of the Recruitment Series

In **Example 3.17** we fit an AR(2) model to the recruitment series using regression. Below are the results of fitting the same model using Yule-Walker estimation in R, which are nearly identical to the values in **Example 3.17**.

```
rec.yw = ar.yw(rec, order=2)
rec.yw$x.mean # = 62.26 (mean estimate)
rec.yw$ar      # = 1.33, -.44 (parameter estimates)
sqrt(diag(rec.yw$sasy.var.coef)) # = .04, .04 (standard errors)
rec.yw$var.pred # = 94.80 (error variance estimate)
```

To obtain the 24 month ahead predictions and their standard errors, and then plot the results as in **Example 3.24**, use the R commands:

```
rec.pr = predict(rec.yw, n.ahead=24)
U = rec.pr$pred + rec.pr$se
L = rec.pr$pred - rec.pr$se
minx = min(rec,L); maxx = max(rec,U)
ts.plot(rec, rec.pr$pred, xlim=c(1980,1990), ylim=c(minx,maxx))
lines(rec.pr$pred, col="red", type="o")
lines(U, col="blue", lty="dashed")
lines(L, col="blue", lty="dashed")
```

In the case of $AR(p)$ models, the Yule–Walker estimators given in (3.102) are optimal in the sense that the asymptotic distribution, (3.103), is the best asymptotic normal distribution. This is because, given initial conditions, $AR(p)$ models are linear models, and the Yule–Walker estimators are essentially least squares estimators. If we use method of moments for MA or ARMA models, we will not get optimal estimators because such processes are nonlinear in the parameters.

Example 3.28 Method of Moments Estimation for an MA(1)

Consider the time series

$$x_t = w_t + \theta w_{t-1},$$

where $|\theta| < 1$. The model can then be written as

$$x_t = \sum_{j=1}^{\infty} (-\theta)^j x_{t-j} + w_t,$$

which is nonlinear in θ . The first two population autocovariances are $\gamma(0) = \sigma_w^2(1 + \theta^2)$ and $\gamma(1) = \sigma_w^2\theta$, so the estimate of θ is found by solving:

$$\widehat{\rho}(1) = \frac{\widehat{\gamma}(1)}{\widehat{\gamma}(0)} = \frac{\widehat{\theta}}{1 + \widehat{\theta}^2}.$$

Two solutions exist, so we would pick the invertible one. If $|\widehat{\rho}(1)| \leq \frac{1}{2}$, the solutions are real, otherwise, a real solution does not exist. Even though $|\rho(1)| < \frac{1}{2}$ for an invertible MA(1), it may happen that $|\widehat{\rho}(1)| \geq \frac{1}{2}$ because it is an estimator. For example, the following simulation in R produces a value of $\widehat{\rho}(1) = .507$ when the true value is $\rho(1) = .9/(1 + .9^2) = .497$.

```
set.seed(2)
ma1 = arima.sim(list(order = c(0,0,1), ma = 0.9), n = 50)
acf(ma1, plot=FALSE)[1] # = .507 (lag 1 sample ACF)
```

When $|\widehat{\rho}(1)| < \frac{1}{2}$, the invertible estimate is

$$\widehat{\theta} = \frac{1 - \sqrt{1 - 4\widehat{\rho}(1)^2}}{2\widehat{\rho}(1)}.$$

It can be shown that⁵

$$\widehat{\theta} \sim \text{AN}\left(\theta, \frac{1 + \theta^2 + 4\theta^4 + \theta^6 + \theta^8}{n(1 - \theta^2)^2}\right);$$

AN is read *asymptotically normal* and is defined in **Definition A.5** of Appendix A. The maximum likelihood estimator (which we discuss next) of θ , in this case, has an asymptotic variance of $(1 - \theta^2)/n$. When $\theta = .5$, for example, the ratio of the asymptotic variance of the method of moments estimator to the maximum likelihood estimator of θ is about 3.5. That is, for large samples, the variance of the method of moments estimator is about 3.5 times larger than the variance of the MLE of θ when $\theta = .5$.

MAXIMUM LIKELIHOOD AND LEAST SQUARES ESTIMATION

To fix ideas, we first focus on the causal AR(1) case. Let

$$x_t = \mu + \phi(x_{t-1} - \mu) + w_t \quad (3.105)$$

where $|\phi| < 1$ and $w_t \sim \text{iid } N(0, \sigma_w^2)$. Given data $x_1, x_2, \dots, x_n$, we seek the likelihood

$$L(\mu, \phi, \sigma_w^2) = f(x_1, x_2, \dots, x_n \mid \mu, \phi, \sigma_w^2).$$

In the case of an AR(1), we may write the likelihood as

$$L(\mu, \phi, \sigma_w^2) = f(x_1)f(x_2 \mid x_1) \cdots f(x_n \mid x_{n-1}),$$

⁵ The result follows from **Theorem A.7** given in Appendix A and the delta method. See the proof of **Theorem A.7** for details on the delta method.

where we have dropped the parameters in the densities, $f(\cdot)$, to ease the notation. Because $x_t \mid x_{t-1} \sim N(\mu + \phi(x_{t-1} - \mu), \sigma_w^2)$, we have

$$f(x_t \mid x_{t-1}) = f_w[(x_t - \mu) - \phi(x_{t-1} - \mu)],$$

where $f_w(\cdot)$ is the density of w_t , that is, the normal density with mean zero and variance σ_w^2 . We may then write the likelihood as

$$L(\mu, \phi, \sigma_w) = f(x_1) \prod_{t=2}^n f_w[(x_t - \mu) - \phi(x_{t-1} - \mu)].$$

To find $f(x_1)$, we can use the causal representation

$$x_1 = \mu + \sum_{j=0}^{\infty} \phi^j w_{1-j}$$

to see that x_1 is normal, with mean μ and variance $\sigma_w^2/(1 - \phi^2)$. Finally, for an AR(1), the likelihood is

$$L(\mu, \phi, \sigma_w^2) = (2\pi\sigma_w^2)^{-n/2} (1 - \phi^2)^{1/2} \exp \left[-\frac{S(\mu, \phi)}{2\sigma_w^2} \right], \quad (3.106)$$

where

$$S(\mu, \phi) = (1 - \phi^2)(x_1 - \mu)^2 + \sum_{t=2}^n [(x_t - \mu) - \phi(x_{t-1} - \mu)]^2. \quad (3.107)$$

Typically, $S(\mu, \phi)$ is called the unconditional sum of squares. We could have also considered the estimation of μ and ϕ using unconditional least squares, that is, estimation by minimizing $S(\mu, \phi)$.

Taking the partial derivative of the log of (3.106) with respect to σ_w^2 and setting the result equal to zero, we see that for any given values of μ and ϕ in the parameter space, $\sigma_w^2 = n^{-1} S(\mu, \phi)$ maximizes the likelihood. Thus, the maximum likelihood estimate of σ_w^2 is

$$\hat{\sigma}_w^2 = n^{-1} S(\hat{\mu}, \hat{\phi}), \quad (3.108)$$

where $\hat{\mu}$ and $\hat{\phi}$ are the MLEs of μ and ϕ , respectively. If we replace n in (3.108) by $n - 2$, we would obtain the unconditional least squares estimate of σ_w^2 .

If, in (3.106), we take logs, replace σ_w^2 by $\hat{\sigma}_w^2$, and ignore constants, $\hat{\mu}$ and $\hat{\phi}$ are the values that minimize the criterion function

$$l(\mu, \phi) = \log [n^{-1} S(\mu, \phi)] - n^{-1} \log(1 - \phi^2); \quad (3.109)$$

that is, $l(\mu, \phi) \propto -2 \log L(\mu, \phi, \hat{\sigma}_w^2)$.⁶ Because (3.107) and (3.109) are complicated functions of the parameters, the minimization of $l(\mu, \phi)$ or $S(\mu, \phi)$ is accomplished

⁶ The criterion function is sometimes called the profile or concentrated likelihood.

numerically. In the case of AR models, we have the advantage that, conditional on initial values, they are linear models. That is, we can drop the term in the likelihood that causes the nonlinearity. Conditioning on x_1 , the conditional likelihood becomes

$$\begin{aligned} L(\mu, \phi, \sigma_w^2 \mid x_1) &= \prod_{t=2}^n f_w [(x_t - \mu) - \phi(x_{t-1} - \mu)] \\ &= (2\pi\sigma_w^2)^{-(n-1)/2} \exp \left[-\frac{S_c(\mu, \phi)}{2\sigma_w^2} \right], \end{aligned} \quad (3.110)$$

where the conditional sum of squares is

$$S_c(\mu, \phi) = \sum_{t=2}^n [(x_t - \mu) - \phi(x_{t-1} - \mu)]^2. \quad (3.111)$$

The conditional MLE of σ_w^2 is

$$\hat{\sigma}_w^2 = S_c(\hat{\mu}, \hat{\phi}) / (n - 1), \quad (3.112)$$

and $\hat{\mu}$ and $\hat{\phi}$ are the values that minimize the conditional sum of squares, $S_c(\mu, \phi)$. Letting $\alpha = \mu(1 - \phi)$, the conditional sum of squares can be written as

$$S_c(\mu, \phi) = \sum_{t=2}^n [x_t - (\alpha + \phi x_{t-1})]^2. \quad (3.113)$$

The problem is now the linear regression problem stated in §2.2. Following the results from least squares estimation, we have $\hat{\alpha} = \bar{x}_{(2)} - \hat{\phi}\bar{x}_{(1)}$, where $\bar{x}_{(1)} = (n - 1)^{-1} \sum_{t=1}^{n-1} x_t$, and $\bar{x}_{(2)} = (n - 1)^{-1} \sum_{t=2}^n x_t$, and the conditional estimates are then

$$\hat{\mu} = \frac{\bar{x}_{(2)} - \hat{\phi}\bar{x}_{(1)}}{1 - \hat{\phi}} \quad (3.114)$$

$$\hat{\phi} = \frac{\sum_{t=2}^n (x_t - \bar{x}_{(2)})(x_{t-1} - \bar{x}_{(1)})}{\sum_{t=2}^n (x_{t-1} - \bar{x}_{(1)})^2}. \quad (3.115)$$

From (3.114) and (3.115), we see that $\hat{\mu} \approx \bar{x}$ and $\hat{\phi} \approx \hat{\rho}(1)$. That is, the Yule–Walker estimators and the conditional least squares estimators are approximately the same. The only difference is the inclusion or exclusion of terms involving the endpoints, x_1 and x_n . We can also adjust the estimate of σ_w^2 in (3.112) to be equivalent to the least squares estimator, that is, divide $S_c(\hat{\mu}, \hat{\phi})$ by $(n - 3)$ instead of $(n - 1)$ in (3.112).

For general AR(p) models, maximum likelihood estimation, unconditional least squares, and conditional least squares follow analogously to the AR(1) example. For general ARMA models, it is difficult to write the likelihood as an explicit function of the parameters. Instead, it is advantageous to write the likelihood in terms of the innovations, or one-step-ahead prediction errors, $x_t - x_t^{t-1}$. This will also be useful in Chapter 6 when we study state-space models.

For a normal ARMA(p, q) model, let $\boldsymbol{\beta} = (\mu, \phi_1, \dots, \phi_p, \theta_1, \dots, \theta_q)'$ be the $(p + q + 1)$ -dimensional vector of the model parameters. The likelihood can be written as

$$L(\boldsymbol{\beta}, \sigma_w^2) = \prod_{t=1}^n f(x_t \mid x_{t-1}, \dots, x_1).$$

The conditional distribution of x_t given $x_{t-1}, \dots, x_1$ is Gaussian with mean x_t^{t-1} and variance P_t^{t-1} . Recall from (3.71) that $P_t^{t-1} = \gamma(0) \prod_{j=1}^{t-1} (1 - \phi_{jj}^2)$. For ARMA models, $\gamma(0) = \sigma_w^2 \sum_{j=0}^{\infty} \psi_j^2$, in which case we may write

$$P_t^{t-1} = \sigma_w^2 \left\{ \left[\sum_{j=0}^{\infty} \psi_j^2 \right] \left[\prod_{j=1}^{t-1} (1 - \phi_{jj}^2) \right] \right\} \stackrel{\text{def}}{=} \sigma_w^2 r_t,$$

where r_t is the term in the braces. Note that the r_t terms are functions only of the regression parameters and that they may be computed recursively as $r_{t+1} = (1 - \phi_{tt}^2) r_t$ with initial condition $r_1 = \sum_{j=0}^{\infty} \psi_j^2$. The likelihood of the data can now be written as

$$L(\boldsymbol{\beta}, \sigma_w^2) = (2\pi\sigma_w^2)^{-n/2} [r_1(\boldsymbol{\beta}) r_2(\boldsymbol{\beta}) \cdots r_n(\boldsymbol{\beta})]^{-1/2} \exp \left[-\frac{S(\boldsymbol{\beta})}{2\sigma_w^2} \right], \quad (3.116)$$

where

$$S(\boldsymbol{\beta}) = \sum_{t=1}^n \left[\frac{(x_t - x_t^{t-1}(\boldsymbol{\beta}))^2}{r_t(\boldsymbol{\beta})} \right]. \quad (3.117)$$

Both x_t^{t-1} and r_t are functions of $\boldsymbol{\beta}$ alone, and we make that fact explicit in (3.116)-(3.117). Given values for $\boldsymbol{\beta}$ and σ_w^2 , the likelihood may be evaluated using the techniques of §3.5. Maximum likelihood estimation would now proceed by maximizing (3.116) with respect to $\boldsymbol{\beta}$ and σ_w^2 . As in the AR(1) example, we have

$$\hat{\sigma}_w^2 = n^{-1} S(\hat{\boldsymbol{\beta}}), \quad (3.118)$$

where $\hat{\boldsymbol{\beta}}$ is the value of $\boldsymbol{\beta}$ that minimizes the concentrated likelihood

$$l(\boldsymbol{\beta}) = \log [n^{-1} S(\boldsymbol{\beta})] + n^{-1} \sum_{t=1}^n \log r_t(\boldsymbol{\beta}). \quad (3.119)$$

For the AR(1) model (3.105) discussed previously, recall that $x_1^0 = \mu$ and $x_t^{t-1} = \mu + \phi(x_{t-1} - \mu)$, for $t = 2, \dots, n$. Also, using the fact that $\phi_{11} = \phi$ and $\phi_{hh} = 0$ for $h > 1$, we have $r_1 = \sum_{j=0}^{\infty} \phi^{2j} = (1 - \phi^2)^{-1}$, $r_2 = (1 - \phi^2)^{-1}(1 - \phi^2) = 1$, and in general, $r_t = 1$ for $t = 2, \dots, n$. Hence, the likelihood presented in (3.106) is identical to the innovations form of the likelihood given by (3.116). Moreover, the generic $S(\boldsymbol{\beta})$ in (3.117) is $S(\mu, \phi)$ given in (3.107) and the generic $l(\boldsymbol{\beta})$ in (3.119) is $l(\mu, \phi)$ in (3.109).

Unconditional least squares would be performed by minimizing (3.117) with respect to $\boldsymbol{\beta}$. Conditional least squares estimation would involve minimizing (3.117)

with respect to $\boldsymbol{\beta}$ but where, to ease the computational burden, the predictions and their errors are obtained by conditioning on initial values of the data. In general, numerical optimization routines are used to obtain the actual estimates and their standard errors.

Example 3.29 The Newton–Raphson and Scoring Algorithms

Two common numerical optimization routines for accomplishing maximum likelihood estimation are Newton–Raphson and scoring. We will give a brief account of the mathematical ideas here. The actual implementation of these algorithms is much more complicated than our discussion might imply. For details, the reader is referred to any of the *Numerical Recipes* books, for example, Press et al. (1993).

Let $l(\boldsymbol{\beta})$ be a criterion function of k parameters $\boldsymbol{\beta} = (\beta_1, \dots, \beta_k)$ that we wish to minimize with respect to $\boldsymbol{\beta}$. For example, consider the likelihood function given by (3.109) or by (3.119). Suppose $l(\widehat{\boldsymbol{\beta}})$ is the extremum that we are interested in finding, and $\widehat{\boldsymbol{\beta}}$ is found by solving $\partial l(\boldsymbol{\beta})/\partial \beta_j = 0$, for $j = 1, \dots, k$. Let $l^{(1)}(\boldsymbol{\beta})$ denote the $k \times 1$ vector of partials

$$l^{(1)}(\boldsymbol{\beta}) = \left(\frac{\partial l(\boldsymbol{\beta})}{\partial \beta_1}, \dots, \frac{\partial l(\boldsymbol{\beta})}{\partial \beta_k} \right)'.$$

Note, $l^{(1)}(\widehat{\boldsymbol{\beta}}) = \mathbf{0}$, the $k \times 1$ zero vector. Let $l^{(2)}(\boldsymbol{\beta})$ denote the $k \times k$ matrix of second-order partials

$$l^{(2)}(\boldsymbol{\beta}) = \left\{ -\frac{\partial^2 l(\boldsymbol{\beta})}{\partial \beta_i \partial \beta_j} \right\}_{i,j=1}^k,$$

and assume $l^{(2)}(\boldsymbol{\beta})$ is nonsingular. Let $\boldsymbol{\beta}_{(0)}$ be an initial estimator of $\boldsymbol{\beta}$. Then, using a Taylor expansion, we have the following approximation:

$$\mathbf{0} = l^{(1)}(\widehat{\boldsymbol{\beta}}) \approx l^{(1)}(\boldsymbol{\beta}_{(0)}) - l^{(2)}(\boldsymbol{\beta}_{(0)}) [\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}_{(0)}].$$

Setting the right-hand side equal to zero and solving for $\widehat{\boldsymbol{\beta}}$ [call the solution $\boldsymbol{\beta}_{(1)}$], we get

$$\boldsymbol{\beta}_{(1)} = \boldsymbol{\beta}_{(0)} + [l^{(2)}(\boldsymbol{\beta}_{(0)})]^{-1} l^{(1)}(\boldsymbol{\beta}_{(0)}).$$

The Newton–Raphson algorithm proceeds by iterating this result, replacing $\boldsymbol{\beta}_{(0)}$ by $\boldsymbol{\beta}_{(1)}$ to get $\boldsymbol{\beta}_{(2)}$, and so on, until convergence. Under a set of appropriate conditions, the sequence of estimators, $\boldsymbol{\beta}_{(1)}, \boldsymbol{\beta}_{(2)}, \dots$, will converge to $\widehat{\boldsymbol{\beta}}$, the MLE of $\boldsymbol{\beta}$.

For maximum likelihood estimation, the criterion function used is $l(\boldsymbol{\beta})$ given by (3.119); $l^{(1)}(\boldsymbol{\beta})$ is called the score vector, and $l^{(2)}(\boldsymbol{\beta})$ is called the Hessian. In the method of scoring, we replace $l^{(2)}(\boldsymbol{\beta})$ by $E[l^{(2)}(\boldsymbol{\beta})]$, the information matrix. Under appropriate conditions, the inverse of the information matrix is the asymptotic variance–covariance matrix of the estimator $\widehat{\boldsymbol{\beta}}$. This is sometimes approximated by the inverse of the Hessian at $\widehat{\boldsymbol{\beta}}$. If the derivatives are difficult to obtain, it is possible to use quasi-maximum likelihood estimation where numerical techniques are used to approximate the derivatives.

Example 3.30 MLE for the Recruitment Series

So far, we have fit an AR(2) model to the Recruitment series using ordinary least squares (Example 3.17) and using Yule–Walker (Example 3.27). The following is an R session used to fit an AR(2) model via maximum likelihood estimation to the Recruitment series; these results can be compared to the results in Example 3.17 and Example 3.27.

```
rec.mle = ar.mle(rec, order=2)
rec.mle$x.mean    # 62.26
rec.mle$ar        # 1.35, -.46
sqrt(diag(rec.mle$asy.var.coef)) # .04, .04
rec.mle$var.pred  # 89.34
```

We now discuss least squares for ARMA(p, q) models via Gauss–Newton. For general and complete details of the Gauss–Newton procedure, the reader is referred to Fuller (1996). As before, write $\boldsymbol{\beta} = (\phi_1, \dots, \phi_p, \theta_1, \dots, \theta_q)'$, and for the ease of discussion, we will put $\mu = 0$. We write the model in terms of the errors

$$w_t(\boldsymbol{\beta}) = x_t - \sum_{j=1}^p \phi_j x_{t-j} - \sum_{k=1}^q \theta_k w_{t-k}(\boldsymbol{\beta}), \quad (3.120)$$

emphasizing the dependence of the errors on the parameters.

For conditional least squares, we approximate the residual sum of squares by conditioning on $x_1, \dots, x_p$ (if $p > 0$) and $w_p = w_{p-1} = w_{p-2} = \dots = w_{1-q} = 0$ (if $q > 0$), in which case, given $\boldsymbol{\beta}$, we may evaluate (3.120) for $t = p+1, p+2, \dots, n$. Using this conditioning argument, the conditional error sum of squares is

$$S_c(\boldsymbol{\beta}) = \sum_{t=p+1}^n w_t^2(\boldsymbol{\beta}). \quad (3.121)$$

Minimizing $S_c(\boldsymbol{\beta})$ with respect to $\boldsymbol{\beta}$ yields the conditional least squares estimates. If $q = 0$, the problem is linear regression and no iterative technique is needed to minimize $S_c(\phi_1, \dots, \phi_p)$. If $q > 0$, the problem becomes nonlinear regression and we will have to rely on numerical optimization.

When n is large, conditioning on a few initial values will have little influence on the final parameter estimates. In the case of small to moderate sample sizes, one may wish to rely on unconditional least squares. The unconditional least squares problem is to choose $\boldsymbol{\beta}$ to minimize the unconditional sum of squares, which we have generically denoted by $S(\boldsymbol{\beta})$ in this section. The unconditional sum of squares can be written in various ways, and one useful form in the case of ARMA(p, q) models is derived in Box et al. (1994, Appendix A7.3). They showed (see Problem 3.19) the unconditional sum of squares can be written as

$$S(\boldsymbol{\beta}) = \sum_{t=-\infty}^n \widehat{w}_t^2(\boldsymbol{\beta}), \quad (3.122)$$

where $\widehat{w}_t(\boldsymbol{\beta}) = E(w_t | x_1, \dots, x_n)$. When $t \leq 0$, the $\widehat{w}_t(\boldsymbol{\beta})$ are obtained by backcasting. As a practical matter, we approximate $S(\boldsymbol{\beta})$ by starting the sum at $t = -M + 1$,

where M is chosen large enough to guarantee $\sum_{t=-\infty}^{-M} \widehat{w}_t^2(\boldsymbol{\beta}) \approx 0$. In the case of unconditional least squares estimation, a numerical optimization technique is needed even when $q = 0$.

To employ Gauss–Newton, let $\boldsymbol{\beta}_{(0)} = (\phi_1^{(0)}, \dots, \phi_p^{(0)}, \theta_1^{(0)}, \dots, \theta_q^{(0)})'$ be an initial estimate of $\boldsymbol{\beta}$. For example, we could obtain $\boldsymbol{\beta}_{(0)}$ by method of moments. The first-order Taylor expansion of $w_t(\boldsymbol{\beta})$ is

$$w_t(\boldsymbol{\beta}) \approx w_t(\boldsymbol{\beta}_{(0)}) - (\boldsymbol{\beta} - \boldsymbol{\beta}_{(0)})' \mathbf{z}_t(\boldsymbol{\beta}_{(0)}), \quad (3.123)$$

where

$$\mathbf{z}_t(\boldsymbol{\beta}_{(0)}) = \left(-\frac{\partial w_t(\boldsymbol{\beta}_{(0)})}{\partial \beta_1}, \dots, -\frac{\partial w_t(\boldsymbol{\beta}_{(0)})}{\partial \beta_{p+q}} \right)', \quad t = 1, \dots, n.$$

The linear approximation of $S_c(\boldsymbol{\beta})$ is

$$Q(\boldsymbol{\beta}) = \sum_{t=p+1}^n \left[w_t(\boldsymbol{\beta}_{(0)}) - (\boldsymbol{\beta} - \boldsymbol{\beta}_{(0)})' \mathbf{z}_t(\boldsymbol{\beta}_{(0)}) \right]^2 \quad (3.124)$$

and this is the quantity that we will minimize. For approximate unconditional least squares, we would start the sum in (3.124) at $t = -M + 1$, for a large value of M , and work with the backcasted values.

Using the results of ordinary least squares (§2.2), we know

$$(\widehat{\boldsymbol{\beta} - \boldsymbol{\beta}_{(0)}}) = \left(n^{-1} \sum_{t=p+1}^n \mathbf{z}_t(\boldsymbol{\beta}_{(0)}) \mathbf{z}_t'(\boldsymbol{\beta}_{(0)}) \right)^{-1} \left(n^{-1} \sum_{t=p+1}^n \mathbf{z}_t(\boldsymbol{\beta}_{(0)}) w_t(\boldsymbol{\beta}_{(0)}) \right) \quad (3.125)$$

minimizes $Q(\boldsymbol{\beta})$. From (3.125), we write the one-step Gauss–Newton estimate as

$$\boldsymbol{\beta}_{(1)} = \boldsymbol{\beta}_{(0)} + \Delta(\boldsymbol{\beta}_{(0)}), \quad (3.126)$$

where $\Delta(\boldsymbol{\beta}_{(0)})$ denotes the right-hand side of (3.125). Gauss–Newton estimation is accomplished by replacing $\boldsymbol{\beta}_{(0)}$ by $\boldsymbol{\beta}_{(1)}$ in (3.126). This process is repeated by calculating, at iteration $j = 2, 3, \dots$,

$$\boldsymbol{\beta}_{(j)} = \boldsymbol{\beta}_{(j-1)} + \Delta(\boldsymbol{\beta}_{(j-1)})$$

until convergence.

Example 3.31 Gauss–Newton for an MA(1)

Consider an invertible MA(1) process, $x_t = w_t + \theta w_{t-1}$. Write the truncated errors as

$$w_t(\theta) = x_t - \theta w_{t-1}(\theta), \quad t = 1, \dots, n, \quad (3.127)$$

where we condition on $w_0(\theta) = 0$. Taking derivatives,

$$-\frac{\partial w_t(\theta)}{\partial \theta} = w_{t-1}(\theta) + \theta \frac{\partial w_{t-1}(\theta)}{\partial \theta}, \quad t = 1, \dots, n, \quad (3.128)$$

where $\partial w_0(\theta)/\partial \theta = 0$. Using the notation of (3.123), we can also write (3.128) as




Fig. 3.9. ACF and PACF of transformed glacial varves.

$$z_t(\theta) = w_{t-1}(\theta) - \theta z_{t-1}(\theta), \quad t = 1, \dots, n, \quad (3.129)$$

where $z_0(\theta) = 0$.

Let $\theta_{(0)}$ be an initial estimate of θ , for example, the estimate given in [Example 3.28](#). Then, the Gauss–Newton procedure for conditional least squares is given by

$$\theta_{(j+1)} = \theta_{(j)} + \frac{\sum_{t=1}^n z_t(\theta_{(j)}) w_t(\theta_{(j)})}{\sum_{t=1}^n z_t^2(\theta_{(j)})}, \quad j = 0, 1, 2, \dots, \quad (3.130)$$

where the values in (3.130) are calculated recursively using (3.127) and (3.129). The calculations are stopped when $|\theta_{(j+1)} - \theta_{(j)}|$, or $|Q(\theta_{(j+1)}) - Q(\theta_{(j)})|$, are smaller than some preset amount.

Example 3.32 Fitting the Glacial Varve Series

Consider the series of glacial varve thicknesses from Massachusetts for $n = 634$ years, as analyzed in [Example 2.6](#) and in [Problem 2.8](#), where it was argued that a first-order moving average model might fit the logarithmically transformed and differenced varve series, say,

$$\nabla \log(x_t) = \log(x_t) - \log(x_{t-1}) = \log\left(\frac{x_t}{x_{t-1}}\right),$$

which can be interpreted as being approximately the percentage change in the thickness.

The sample ACF and PACF, shown in [Figure 3.9](#), confirm the tendency of $\nabla \log(x_t)$ to behave as a first-order moving average process as the ACF has only a




Fig. 3.10. Conditional sum of squares versus values of the moving average parameter for the glacial varve example, [Example 3.32](#). Vertical lines indicate the values of the parameter obtained via Gauss–Newton; see [Table 3.2](#) for the actual values.

significant peak at lag one and the PACF decreases exponentially. Using [Table 3.1](#), this sample behavior fits that of the MA(1) very well.

The results of eleven iterations of the Gauss–Newton procedure, [\(3.130\)](#), starting with $\theta_{(0)} = -.10$ are given in [Table 3.2](#). The final estimate is $\hat{\theta} = \theta_{(11)} = -.773$; interim values and the corresponding value of the conditional sum of squares, $S_c(\theta)$ given in [\(3.121\)](#), are also displayed in the table. The final estimate of the error variance is $\hat{\sigma}_w^2 = 148.98/632 = .236$ with 632 degrees of freedom (one is lost in differencing). The value of the sum of the squared derivatives at convergence is $\sum_{t=1}^n z_t^2(\theta_{(11)}) = 369.73$, and consequently, the estimated standard error of $\hat{\theta}$ is $\sqrt{.236/369.73} = .025$;⁷ this leads to a t -value of $-.773/.025 = -30.92$ with 632 degrees of freedom.

[Figure 3.10](#) displays the conditional sum of squares, $S_c(\theta)$ as a function of θ , as well as indicating the values of each step of the Gauss–Newton algorithm. Note that the Gauss–Newton procedure takes large steps toward the minimum initially, and then takes very small steps as it gets close to the minimizing value. When there is only one parameter, as in this case, it would be easy to evaluate $S_c(\theta)$ on a grid of points, and then choose the appropriate value of θ from the grid search. It would be difficult, however, to perform grid searches when there are many parameters.

In the general case of causal and invertible ARMA(p, q) models, maximum likelihood estimation and conditional and unconditional least squares estimation (and Yule–Walker estimation in the case of AR models) all lead to optimal estimators. The

⁷ To estimate the standard error, we are using the standard regression results from [\(2.6\)](#) as an approximation

Table 3.2. Gauss–Newton Results for **Example 3.32**

j	$\theta_{(j)}$	$S_c(\theta_{(j)})$	$\sum_{t=1}^n z_t^2(\theta_{(j)})$
0	−0.100	195.0010	183.3464
1	−0.250	177.7614	163.3038
2	−0.400	165.0027	161.6279
3	−0.550	155.6723	182.6432
4	−0.684	150.2896	247.4942
5	−0.736	149.2283	304.3125
6	−0.757	149.0272	337.9200
7	−0.766	148.9885	355.0465
8	−0.770	148.9812	363.2813
9	−0.771	148.9804	365.4045
10	−0.772	148.9799	367.5544
11	−0.773	148.9799	369.7314

proof of this general result can be found in a number of texts on theoretical time series analysis (for example, Brockwell and Davis, 1991, or Hannan, 1970, to mention a few). We will denote the ARMA coefficient parameters by $\boldsymbol{\beta} = (\phi_1, \dots, \phi_p, \theta_1, \dots, \theta_q)'$.

Property 3.10 Large Sample Distribution of the Estimators

Under appropriate conditions, for causal and invertible ARMA processes, the maximum likelihood, the unconditional least squares, and the conditional least squares estimators, each initialized by the method of moments estimator, all provide optimal estimators of σ_w^2 and $\boldsymbol{\beta}$, in the sense that $\widehat{\sigma}_w^2$ is consistent, and the asymptotic distribution of $\widehat{\boldsymbol{\beta}}$ is the best asymptotic normal distribution. In particular, as $n \rightarrow \infty$,

$$\sqrt{n}(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}) \xrightarrow{d} N(\mathbf{0}, \sigma_w^2 \boldsymbol{\Gamma}_{p,q}^{-1}). \quad (3.131)$$

The asymptotic variance–covariance matrix of the estimator $\widehat{\boldsymbol{\beta}}$ is the inverse of the information matrix. In particular, the $(p+q) \times (p+q)$ matrix $\boldsymbol{\Gamma}_{p,q}$, has the form

$$\boldsymbol{\Gamma}_{p,q} = \begin{pmatrix} \boldsymbol{\Gamma}_{\phi\phi} & \boldsymbol{\Gamma}_{\phi\theta} \\ \boldsymbol{\Gamma}_{\theta\phi} & \boldsymbol{\Gamma}_{\theta\theta} \end{pmatrix}. \quad (3.132)$$

The $p \times p$ matrix $\boldsymbol{\Gamma}_{\phi\phi}$ is given by (3.100), that is, the ij -th element of $\boldsymbol{\Gamma}_{\phi\phi}$, for $i, j = 1, \dots, p$, is $\gamma_x(i-j)$ from an $\text{AR}(p)$ process, $\phi(B)x_t = w_t$. Similarly, $\boldsymbol{\Gamma}_{\theta\theta}$ is a $q \times q$ matrix with the ij -th element, for $i, j = 1, \dots, q$, equal to $\gamma_y(i-j)$ from an $\text{AR}(q)$ process, $\theta(B)y_t = w_t$. The $p \times q$ matrix $\boldsymbol{\Gamma}_{\phi\theta} = \{\gamma_{xy}(i-j)\}$, for $i = 1, \dots, p$; $j = 1, \dots, q$; that is, the ij -th element is the cross-covariance between the two AR processes given by $\phi(B)x_t = w_t$ and $\theta(B)y_t = w_t$. Finally, $\boldsymbol{\Gamma}_{\theta\phi} = \boldsymbol{\Gamma}_{\phi\theta}'$ is $q \times p$.

Further discussion of **Property 3.10**, including a proof for the case of least squares estimators for $\text{AR}(p)$ processes, can be found in Appendix B, §B.3.

Example 3.33 Some Specific Asymptotic Distributions

The following are some specific cases of **Property 3.10**.

AR(1): $\gamma_x(0) = \sigma_w^2/(1 - \phi^2)$, so $\sigma_w^2 \Gamma_{1,0}^{-1} = (1 - \phi^2)$. Thus,

$$\hat{\phi} \sim \text{AN} \left[\phi, n^{-1}(1 - \phi^2) \right]. \quad (3.133)$$

AR(2): The reader can verify that

$$\gamma_x(0) = \left(\frac{1 - \phi_2}{1 + \phi_2} \right) \frac{\sigma_w^2}{(1 - \phi_2)^2 - \phi_1^2}$$

and $\gamma_x(1) = \phi_1 \gamma_x(0) + \phi_2 \gamma_x(1)$. From these facts, we can compute $\Gamma_{2,0}^{-1}$. In particular, we have

$$\begin{pmatrix} \hat{\phi}_1 \\ \hat{\phi}_2 \end{pmatrix} \sim \text{AN} \left[\begin{pmatrix} \phi_1 \\ \phi_2 \end{pmatrix}, n^{-1} \begin{pmatrix} 1 - \phi_2^2 & -\phi_1(1 + \phi_2) \\ \text{sym} & 1 - \phi_2^2 \end{pmatrix} \right]. \quad (3.134)$$

MA(1): In this case, write $\theta(B)y_t = w_t$, or $y_t + \theta y_{t-1} = w_t$. Then, analogous to the AR(1) case, $\gamma_y(0) = \sigma_w^2/(1 - \theta^2)$, so $\sigma_w^2 \Gamma_{0,1}^{-1} = (1 - \theta^2)$. Thus,

$$\hat{\theta} \sim \text{AN} \left[\theta, n^{-1}(1 - \theta^2) \right]. \quad (3.135)$$

MA(2): Write $y_t + \theta_1 y_{t-1} + \theta_2 y_{t-2} = w_t$, so, analogous to the AR(2) case, we have

$$\begin{pmatrix} \hat{\theta}_1 \\ \hat{\theta}_2 \end{pmatrix} \sim \text{AN} \left[\begin{pmatrix} \theta_1 \\ \theta_2 \end{pmatrix}, n^{-1} \begin{pmatrix} 1 - \theta_2^2 & \theta_1(1 + \theta_2) \\ \text{sym} & 1 - \theta_2^2 \end{pmatrix} \right]. \quad (3.136)$$

ARMA(1,1): To calculate $\Gamma_{\phi\theta}$, we must find $\gamma_{xy}(0)$, where $x_t - \phi x_{t-1} = w_t$ and $y_t + \theta y_{t-1} = w_t$. We have

$$\begin{aligned} \gamma_{xy}(0) &= \text{cov}(x_t, y_t) = \text{cov}(\phi x_{t-1} + w_t, -\theta y_{t-1} + w_t) \\ &= -\phi\theta \gamma_{xy}(0) + \sigma_w^2. \end{aligned}$$

Solving, we find, $\gamma_{xy}(0) = \sigma_w^2/(1 + \phi\theta)$. Thus,

$$\begin{pmatrix} \hat{\phi} \\ \hat{\theta} \end{pmatrix} \sim \text{AN} \left[\begin{pmatrix} \phi \\ \theta \end{pmatrix}, n^{-1} \begin{bmatrix} (1 - \phi^2)^{-1} & (1 + \phi\theta)^{-1} \\ \text{sym} & (1 - \theta^2)^{-1} \end{bmatrix}^{-1} \right]. \quad (3.137)$$

Example 3.34 Overfitting Caveat

The asymptotic behavior of the parameter estimators gives us an additional insight into the problem of fitting ARMA models to data. For example, suppose a time series follows an AR(1) process and we decide to fit an AR(2) to the data. Do any problems occur in doing this? More generally, why not simply fit large-order AR models to make sure that we capture the dynamics of the process? After all, if the process is truly an AR(1), the other autoregressive parameters will not be

significant. The answer is that if we overfit, we obtain less efficient, or less precise parameter estimates. For example, if we fit an AR(1) to an AR(1) process, for large n , $\text{var}(\widehat{\phi}_1) \approx n^{-1}(1 - \phi_1^2)$. But, if we fit an AR(2) to the AR(1) process, for large n , $\text{var}(\widehat{\phi}_1) \approx n^{-1}(1 - \phi_2^2) = n^{-1}$ because $\phi_2 = 0$. Thus, the variance of ϕ_1 has been inflated, making the estimator less precise.

We do want to mention, however, that overfitting can be used as a diagnostic tool. For example, if we fit an AR(2) model to the data and are satisfied with that model, then adding one more parameter and fitting an AR(3) should lead to approximately the same model as in the AR(2) fit. We will discuss model diagnostics in more detail in §3.8.

The reader might wonder, for example, why the asymptotic distributions of $\widehat{\phi}$ from an AR(1) and $\widehat{\theta}$ from an MA(1) are of the same form; compare (3.133) to (3.135). It is possible to explain this unexpected result heuristically using the intuition of linear regression. That is, for the normal regression model presented in §2.2 with no intercept term, $x_t = \beta z_t + w_t$, we know $\widehat{\beta}$ is normally distributed with mean β , and from (2.6),

$$\text{var} \left\{ \sqrt{n} (\widehat{\beta} - \beta) \right\} = n\sigma_w^2 \left(\sum_{t=1}^n z_t^2 \right)^{-1} = \sigma_w^2 \left(n^{-1} \sum_{t=1}^n z_t^2 \right)^{-1}.$$

For the causal AR(1) model given by $x_t = \phi x_{t-1} + w_t$, the intuition of regression tells us to expect that, for n large,

$$\sqrt{n} (\widehat{\phi} - \phi)$$

is approximately normal with mean zero and with variance given by

$$\sigma_w^2 \left(n^{-1} \sum_{t=2}^n x_{t-1}^2 \right)^{-1}.$$

Now, $n^{-1} \sum_{t=2}^n x_{t-1}^2$ is the sample variance (recall that the mean of x_t is zero) of the x_t , so as n becomes large we would expect it to approach $\text{var}(x_t) = \gamma(0) = \sigma_w^2 / (1 - \phi^2)$. Thus, the large sample variance of $\sqrt{n} (\widehat{\phi} - \phi)$ is

$$\sigma_w^2 \gamma_x(0)^{-1} = \sigma_w^2 \left(\frac{\sigma_w^2}{1 - \phi^2} \right)^{-1} = (1 - \phi^2);$$

that is, (3.133) holds.

In the case of an MA(1), we may use the discussion of **Example 3.31** to write an approximate regression model for the MA(1). That is, consider the approximation (3.129) as the regression model

$$z_t(\widehat{\theta}) = -\theta z_{t-1}(\widehat{\theta}) + w_{t-1},$$

where now, $z_{t-1}(\widehat{\theta})$ as defined in [Example 3.31](#), plays the role of the regressor. Continuing with the analogy, we would expect the asymptotic distribution of $\sqrt{n}(\widehat{\theta} - \theta)$ to be normal, with mean zero, and approximate variance

$$\sigma_w^2 \left(n^{-1} \sum_{t=2}^n z_{t-1}^2(\widehat{\theta}) \right)^{-1}.$$

As in the AR(1) case, $n^{-1} \sum_{t=2}^n z_{t-1}^2(\widehat{\theta})$ is the sample variance of the $z_t(\widehat{\theta})$ so, for large n , this should be $\text{var}\{z_t(\theta)\} = \gamma_z(0)$, say. But note, as seen from [\(3.129\)](#), $z_t(\theta)$ is approximately an AR(1) process with parameter $-\theta$. Thus,

$$\sigma_w^2 \gamma_z(0)^{-1} = \sigma_w^2 \left(\frac{\sigma_w^2}{1 - (-\theta)^2} \right)^{-1} = (1 - \theta^2),$$

which agrees with [\(3.135\)](#). Finally, the asymptotic distributions of the AR parameter estimates and the MA parameter estimates are of the same form because in the MA case, the “regressors” are the differential processes $z_t(\theta)$ that have AR structure, and it is this structure that determines the asymptotic variance of the estimators. For a rigorous account of this approach for the general case, see Fuller (1996, Theorem 5.5.4).

In [Example 3.32](#), the estimated standard error of $\widehat{\theta}$ was .025. In that example, we used regression results to estimate the standard error as the square root of

$$n^{-1} \widehat{\sigma}_w^2 \left(n^{-1} \sum_{t=1}^n z_t^2(\widehat{\theta}) \right)^{-1} = \frac{\widehat{\sigma}_w^2}{\sum_{t=1}^n z_t^2(\widehat{\theta})},$$

where $n = 632$, $\widehat{\sigma}_w^2 = .236$, $\sum_{t=1}^n z_t^2(\widehat{\theta}) = 369.73$ and $\widehat{\theta} = -.773$. Using [\(3.135\)](#), we could have also calculated this value using the asymptotic approximation, the square root of $(1 - (-.773)^2)/632$, which is also .025.

If n is small, or if the parameters are close to the boundaries, the asymptotic approximations can be quite poor. The bootstrap can be helpful in this case; for a broad treatment of the bootstrap, see Efron and Tibshirani (1994). We discuss the case of an AR(1) here and leave the general discussion for Chapter 6. For now, we give a simple example of the bootstrap for an AR(1) process.

Example 3.35 Bootstrapping an AR(1)

We consider an AR(1) model with a regression coefficient near the boundary of causality and an error process that is symmetric but not normal. Specifically, consider the causal model

$$x_t = \mu + \phi(x_{t-1} - \mu) + w_t, \quad (3.138)$$

where $\mu = 50$, $\phi = .95$, and w_t are iid double exponential with location zero, and scale parameter $\beta = 2$. The density of w_t is given by




Fig. 3.11. One hundred observations generated from the model in [Example 3.35](#).

$$f(w) = \frac{1}{2\beta} \exp\{-|w|/\beta\} \quad -\infty < w < \infty.$$

In this example, $E(w_t) = 0$ and $\text{var}(w_t) = 2\beta^2 = 8$. [Figure 3.11](#) shows $n = 100$ simulated observations from this process. This particular realization is interesting; the data look like they were generated from a nonstationary process with three different mean levels. In fact, the data were generated from a well-behaved, albeit non-normal, stationary and causal model. To show the advantages of the bootstrap, we will act as if we do not know the actual error distribution and we will proceed as if it were normal; of course, this means, for example, that the normal based MLE of ϕ will not be the actual MLE because the data are not normal.

Using the data shown in [Figure 3.11](#), we obtained the Yule–Walker estimates $\hat{\mu} = 40.05$, $\hat{\phi} = .96$, and $s_w^2 = 15.30$, where s_w^2 is the estimate of $\text{var}(w_t)$. Based on [Property 3.10](#), we would say that $\hat{\phi}$ is approximately normal with mean ϕ (which we supposedly do not know) and variance $(1 - \phi^2)/100$, which we would approximate by $(1 - .96^2)/100 = .03^2$.

To assess the finite sample distribution of $\hat{\phi}$ when $n = 100$, we simulated 1000 realizations of this AR(1) process and estimated the parameters via Yule–Walker. The finite sampling density of the Yule–Walker estimate of ϕ , based on the 1000 repeated simulations, is shown in [Figure 3.12](#). Clearly the sampling distribution is not close to normality for this sample size. The mean of the distribution shown in [Figure 3.12](#) is .89, and the variance of the distribution is .05²; these values are considerably different than the asymptotic values. Some of the quantiles of the finite sample distribution are .79 (5%), .86 (25%), .90 (50%), .93 (75%), and .95 (95%). The R code to perform the simulation and plot the histogram is as follows:

```
set.seed(111)
phi.yw = rep(NA, 1000)
for (i in 1:1000){
```




Fig. 3.12. Finite sample density of the Yule-Walker estimate of ϕ in [Example 3.35](#).

```
e = rexp(150, rate=.5); u = runif(150,-1,1); de = e*sign(u)
x = 50 + arima.sim(n=100,list(ar=.95), innov=de, n.start=50)
phi.yw[i] = ar.yw(x, order=1)$ar }
hist(phi.yw, prob=TRUE, main="")
lines(density(phi.yw, bw=.015))
```

Before discussing the bootstrap, we first investigate the sample innovation process, $x_t - x_t^{t-1}$, with corresponding variances P_t^{t-1} . For the AR(1) model in this example,

$$x_t^{t-1} = \mu + \phi(x_{t-1} - \mu), \quad t = 2, \dots, 100.$$

From this, it follows that

$$P_t^{t-1} = E(x_t - x_t^{t-1})^2 = \sigma_w^2, \quad t = 2, \dots, 100.$$

When $t = 1$, we have

$$x_1^0 = \mu \quad \text{and} \quad P_1^0 = \sigma_w^2 / (1 - \phi^2).$$

Thus, the innovations have zero mean but different variances; in order that all of the innovations have the same variance, σ_w^2 , we will write them as

$$\begin{aligned} \epsilon_1 &= (x_1 - \mu) \sqrt{(1 - \phi^2)} \\ \epsilon_t &= (x_t - \mu) - \phi(x_{t-1} - \mu), \quad \text{for } t = 2, \dots, 100. \end{aligned} \quad (3.139)$$

From these equations, we can write the model in terms of the ϵ_t as

$$\begin{aligned} x_1 &= \mu + \epsilon_1 / \sqrt{(1 - \phi^2)} \\ x_t &= \mu + \phi(x_{t-1} - \mu) + \epsilon_t \quad \text{for } t = 2, \dots, 100. \end{aligned} \quad (3.140)$$

Next, replace the parameters with their estimates in (3.139), that is, $\hat{\mu} = 40.048$ and $\hat{\phi} = .957$, and denote the resulting sample innovations as $\{\hat{\epsilon}_1, \dots, \hat{\epsilon}_{100}\}$. To

obtain one bootstrap sample, first randomly sample, with replacement, $n = 100$ values from the set of sample innovations; call the sampled values $\{\epsilon_1^*, \dots, \epsilon_{100}^*\}$. Now, generate a bootstrapped data set sequentially by setting

$$\begin{aligned} x_1^* &= 40.048 + \epsilon_1^* / \sqrt{(1 - .957^2)} \\ x_t^* &= 40.048 + .957(x_{t-1}^* - 40.048) + \epsilon_t^*, \quad t = 2, \dots, n. \end{aligned} \quad (3.141)$$

Next, estimate the parameters as if the data were x_t^* . Call these estimates $\hat{\mu}(1)$, $\hat{\phi}(1)$, and $s_w^2(1)$. Repeat this process a large number, B , of times, generating a collection of bootstrapped parameter estimates, $\{\hat{\mu}(b), \hat{\phi}(b), s_w^2(b), b = 1, \dots, B\}$. We can then approximate the finite sample distribution of an estimator from the bootstrapped parameter values. For example, we can approximate the distribution of $\hat{\phi} - \phi$ by the empirical distribution of $\hat{\phi}(b) - \hat{\phi}$, for $b = 1, \dots, B$.

Figure 3.13 shows the bootstrap histogram of 200 bootstrapped estimates of ϕ using the data shown in Figure 3.11. In addition, Figure 3.13 shows a density estimate based on the bootstrap histogram, as well as the asymptotic normal density that would have been used based on Property 3.10. Clearly, the bootstrap distribution of $\hat{\phi}$ is closer to the distribution of $\hat{\phi}$ shown in Figure 3.12 than to the asymptotic normal approximation. In particular, the mean of the distribution of $\hat{\phi}(b)$ is .92 with a variance of .05². Some quantiles of this distribution are .83 (5%), .90 (25%), .93 (50%), .95 (75%), and .98 (95%).

To perform a similar bootstrap exercise in R, use the following commands. We note that the R estimation procedure is conditional on the first observation, so the first residual is not returned. To get around this problem, we simply fix the first observation and bootstrap the remaining data. The simulated data are available in the file `arlboot`, but you can simulate your own data as was done in the code that produced Figure 3.12.

```
x = arlboot
m = mean(x)      # estimate of mu
fit = ar.yw(x, order=1)
phi = fit$ar      # estimate of phi
nboot = 200      # number of bootstrap replicates
resids = fit$resid[-1] # the first resid is NA
x.star = x       # initialize x*
phi.star.yw = rep(NA, nboot)
for (i in 1:nboot) {
  resid.star = sample(resids, replace=TRUE)
  for (t in 1:99){ x.star[t+1] = m + phi*(x.star[t]-m) + resid.star[t] }
  phi.star.yw[i] = ar.yw(x.star, order=1)$ar }
hist(phi.star.yw, 10, main="", prob=TRUE, ylim=c(0,14), xlim=c(.75,1.05))
lines(density(phi.star.yw, bw=.02))
u = seq(.75, 1.05, by=.001)
lines(u, dnorm(u, mean=.96, sd=.03), lty="dashed", lwd=2)
```




Fig. 3.13. Bootstrap histogram of $\hat{\phi}$ based on 200 bootstraps; a density estimate based on the histogram (solid line) and the corresponding asymptotic normal density (dashed line).

3.7 Integrated Models for Nonstationary Data

In Chapters 1 and 2, we saw that if x_t is a random walk, $x_t = x_{t-1} + w_t$, then by differencing x_t , we find that $\nabla x_t = w_t$ is stationary. In many situations, time series can be thought of as being composed of two components, a nonstationary trend component and a zero-mean stationary component. For example, in §2.2 we considered the model

$$x_t = \mu_t + y_t, \quad (3.142)$$

where $\mu_t = \beta_0 + \beta_1 t$ and y_t is stationary. Differencing such a process will lead to a stationary process:

$$\nabla x_t = x_t - x_{t-1} = \beta_1 + y_t - y_{t-1} = \beta_1 + \nabla y_t.$$

Another model that leads to first differencing is the case in which μ_t in (3.142) is stochastic and slowly varying according to a random walk. That is,

$$\mu_t = \mu_{t-1} + v_t$$

where v_t is stationary. In this case,

$$\nabla x_t = v_t + \nabla y_t,$$

is stationary. If μ_t in (3.142) is a k -th order polynomial, $\mu_t = \sum_{j=0}^k \beta_j t^j$, then (Problem 3.27) the differenced series $\nabla^k y_t$ is stationary. Stochastic trend models can also lead to higher order differencing. For example, suppose

$$\mu_t = \mu_{t-1} + v_t \quad \text{and} \quad v_t = v_{t-1} + e_t,$$

where e_t is stationary. Then, $\nabla x_t = v_t + \nabla y_t$ is not stationary, but

$$\nabla^2 x_t = e_t + \nabla^2 y_t$$

is stationary.

The integrated ARMA, or ARIMA, model is a broadening of the class of ARMA models to include differencing.

Definition 3.11 A process x_t is said to be **ARIMA**(p, d, q) if

$$\nabla^d x_t = (1 - B)^d x_t$$

is ARMA(p, q). In general, we will write the model as

$$\phi(B)(1 - B)^d x_t = \theta(B)w_t. \quad (3.143)$$

If $E(\nabla^d x_t) = \mu$, we write the model as

$$\phi(B)(1 - B)^d x_t = \delta + \theta(B)w_t,$$

where $\delta = \mu(1 - \phi_1 - \cdots - \phi_p)$.

Because of the nonstationarity, care must be taken when deriving forecasts. For the sake of completeness, we discuss this issue briefly here, but we stress the fact that both the theoretical and computational aspects of the problem are best handled via state-space models. We discuss the theoretical details in Chapter 6. For information on the state-space based computational aspects in R, see the ARIMA help files ([?arima](#) and [?predict.Arima](#)); our scripts [sarima](#) and [sarima.for](#) are basically front ends for these R scripts.

It should be clear that, since $y_t = \nabla^d x_t$ is ARMA, we can use §3.5 methods to obtain forecasts of y_t , which in turn lead to forecasts for x_t . For example, if $d = 1$, given forecasts y_{n+m}^n for $m = 1, 2, \dots$, we have $y_{n+m}^n = x_{n+m}^n - x_{n+m-1}^n$, so that

$$x_{n+m}^n = y_{n+m}^n + x_{n+m-1}^n$$

with initial condition $x_{n+1}^n = y_{n+1}^n + x_n$ (noting $x_n^n = x_n$).

It is a little more difficult to obtain the prediction errors P_{n+m}^n , but for large n , the approximation used in §3.5, equation (3.86), works well. That is, the mean-squared prediction error can be approximated by

$$P_{n+m}^n = \sigma_w^2 \sum_{j=0}^{m-1} \psi_j^{*2}, \quad (3.144)$$

where ψ_j^* is the coefficient of z^j in $\psi^*(z) = \theta(z)/\phi(z)(1 - z)^d$.

To better understand integrated models, we examine the properties of some simple cases; [Problem 3.29](#) covers the ARIMA(1, 1, 0) case.

Example 3.36 Random Walk with Drift

To fix ideas, we begin by considering the random walk with drift model first presented in [Example 1.11](#), that is,

$$x_t = \delta + x_{t-1} + w_t,$$

for $t = 1, 2, \dots$, and $x_0 = 0$. Technically, the model is not ARIMA, but we could include it trivially as an ARIMA(0, 1, 0) model. Given data $x_1, \dots, x_n$, the one-step-ahead forecast is given by

$$x_{n+1}^n = E(x_{n+1} \mid x_n, \dots, x_1) = E(\delta + x_n + w_{n+1} \mid x_n, \dots, x_1) = \delta + x_n.$$

The two-step-ahead forecast is given by $x_{n+2}^n = \delta + x_{n+1}^n = 2\delta + x_n$, and consequently, the m -step-ahead forecast, for $m = 1, 2, \dots$, is

$$x_{n+m}^n = m\delta + x_n, \quad (3.145)$$

To obtain the forecast errors, it is convenient to recall equation (1.4), i.e., $x_n = n\delta + \sum_{j=1}^n w_j$, in which case we may write

$$x_{n+m} = (n+m)\delta + \sum_{j=1}^{n+m} w_j = m\delta + x_n + \sum_{j=n+1}^{n+m} w_j.$$

From this it follows that the m -step-ahead prediction error is given by

$$P_{n+m}^n = E(x_{n+m} - x_{n+m}^n)^2 = E\left(\sum_{j=n+1}^{n+m} w_j\right)^2 = m\sigma_w^2. \quad (3.146)$$

Hence, unlike the stationary case (see [Example 3.22](#)), as the forecast horizon grows, the prediction errors, (3.146), increase without bound and the forecasts follow a straight line with slope δ emanating from x_n . We note that (3.144) is exact in this case because $\psi^*(z) = 1/(1-z) = \sum_{j=0}^{\infty} z^j$ for $|z| < 1$, so that $\psi_j^* = 1$ for all j .

The w_t are Gaussian, so estimation is straightforward because the differenced data, say $y_t = \nabla x_t$, are independent and identically distributed normal variates with mean δ and variance σ_w^2 . Consequently, optimal estimates of δ and σ_w^2 are the sample mean and variance of the y_t , respectively.

Example 3.37 IMA(1, 1) and EWMA

The ARIMA(0,1,1), or IMA(1,1) model is of interest because many economic time series can be successfully modeled this way. In addition, the model leads to a frequently used, and abused, forecasting method called exponentially weighted moving averages (EWMA). We will write the model as

$$x_t = x_{t-1} + w_t - \lambda w_{t-1}, \quad (3.147)$$

with $|\lambda| < 1$, for $t = 1, 2, \dots$, and $x_0 = 0$, because this model formulation is easier to work with here, and it leads to the standard representation for EWMA. We could

have included a drift term in (3.147), as was done in the previous example, but for the sake of simplicity, we leave it out of the discussion. If we write

$$y_t = w_t - \lambda w_{t-1},$$

we may write (3.147) as $x_t = x_{t-1} + y_t$. Because $|\lambda| < 1$, y_t has an invertible representation, $y_t = \sum_{j=1}^{\infty} \lambda^j y_{t-j} + w_t$, and substituting $y_t = x_t - x_{t-1}$, we may write

$$x_t = \sum_{j=1}^{\infty} (1 - \lambda) \lambda^{j-1} x_{t-j} + w_t. \quad (3.148)$$

as an approximation for large t (put $x_t = 0$ for $t \leq 0$). Verification of (3.148) is left to the reader (Problem 3.28). Using the approximation (3.148), we have that the approximate one-step-ahead predictor, using the notation of §3.5, is

$$\begin{aligned} \tilde{x}_{n+1} &= \sum_{j=1}^{\infty} (1 - \lambda) \lambda^{j-1} x_{n+1-j} \\ &= (1 - \lambda) x_n + \lambda \sum_{j=1}^{\infty} (1 - \lambda) \lambda^{j-1} x_{n-j} \\ &= (1 - \lambda) x_n + \lambda \tilde{x}_n. \end{aligned} \quad (3.149)$$

From (3.149), we see that the new forecast is a linear combination of the old forecast and the new observation. Based on (3.149) and the fact that we only observe $x_1, \dots, x_n$, and consequently $y_1, \dots, y_n$ (because $y_t = x_t - x_{t-1}$; $x_0 = 0$), the truncated forecasts are

$$\tilde{x}_{n+1}^n = (1 - \lambda) x_n + \lambda \tilde{x}_n^{n-1}, \quad n \geq 1, \quad (3.150)$$

with $\tilde{x}_1^0 = x_1$ as an initial value. The mean-square prediction error can be approximated using (3.144) by noting that $\psi^*(z) = (1 - \lambda z)/(1 - z) = 1 + (1 - \lambda) \sum_{j=1}^{\infty} z^j$ for $|z| < 1$; consequently, for large n , (3.144) leads to

$$P_{n+m}^n \approx \sigma_w^2 [1 + (m - 1)(1 - \lambda)^2].$$

In EWMA, the parameter $1 - \lambda$ is often called the smoothing parameter and is restricted to be between zero and one. Larger values of λ lead to smoother forecasts. This method of forecasting is popular because it is easy to use; we need only retain the previous forecast value and the current observation to forecast the next time period. Unfortunately, as previously suggested, the method is often abused because some forecasters do not verify that the observations follow an IMA(1, 1) process, and often arbitrarily pick values of λ . In the following, we show how to generate 100 observations from an IMA(1, 1) model with $\lambda = -\theta = .8$ and then calculate and display the fitted EWMA superimposed on the data. This is accomplished using the Holt-Winters command in R (see the help file [?HoltWinters](#) for details; no output is shown):

```

set.seed(666)
x = arima.sim(list(order = c(0,1,1), ma = -0.8), n = 100)
(x.ima = HoltWinters(x, beta=FALSE, gamma=FALSE)) #  $\alpha$  below is  $1-\lambda$ 
Smoothing parameter: alpha: 0.1663072
plot(x.ima)

```

3.8 Building ARIMA Models

There are a few basic steps to fitting ARIMA models to time series data. These steps involve plotting the data, possibly transforming the data, identifying the dependence orders of the model, parameter estimation, diagnostics, and model choice. First, as with any data analysis, we should construct a time plot of the data, and inspect the graph for any anomalies. If, for example, the variability in the data grows with time, it will be necessary to transform the data to stabilize the variance. In such cases, the Box–Cox class of power transformations, equation (2.34), could be employed. Also, the particular application might suggest an appropriate transformation. For example, suppose a process evolves as a fairly small and stable percent-change, such as an investment. For example, we might have

$$x_t = (1 + p_t)x_{t-1},$$

where x_t is the value of the investment at time t and p_t is the percentage-change from period $t - 1$ to t , which may be negative. Taking logs we have

$$\log(x_t) = \log(1 + p_t) + \log(x_{t-1}),$$

or

$$\nabla \log(x_t) = \log(1 + p_t).$$

If the percent change p_t stays relatively small in magnitude, then $\log(1 + p_t) \approx p_t$ ⁸ and, thus,

$$\nabla \log(x_t) \approx p_t,$$

will be a relatively stable process. Frequently, $\nabla \log(x_t)$ is called the return or growth rate. This general idea was used in [Example 3.32](#), and we will use it again in [Example 3.38](#).

After suitably transforming the data, the next step is to identify preliminary values of the autoregressive order, p , the order of differencing, d , and the moving average order, q . We have already addressed, in part, the problem of selecting d . A time plot of the data will typically suggest whether any differencing is needed. If differencing is called for, then difference the data once, $d = 1$, and inspect the time plot of ∇x_t . If additional differencing is necessary, then try differencing again and inspect a time plot of $\nabla^2 x_t$. Be careful not to overdifference because this may introduce dependence where none exists. For example, $x_t = w_t$ is serially uncorrelated, but $\nabla x_t = w_t - w_{t-1}$

⁸ $\log(1 + p) = p - \frac{p^2}{2} + \frac{p^3}{3} - \dots$ for $-1 < p \leq 1$. If p is a small percent-change, then the higher-order terms in the expansion are negligible.




Fig. 3.14. Quarterly U.S. GNP from 1947(1) to 2002(3).

is MA(1). In addition to time plots, the sample ACF can help in indicating whether differencing is needed. Because the polynomial $\phi(z)(1 - z)^d$ has a unit root, the sample ACF, $\hat{\rho}(h)$, will not decay to zero fast as h increases. Thus, a slow decay in $\hat{\rho}(h)$ is an indication that differencing may be needed.

When preliminary values of d have been settled, the next step is to look at the sample ACF and PACF of $\nabla^d x_t$ for whatever values of d have been chosen. Using Table 3.1 as a guide, preliminary values of p and q are chosen. Recall that, if $p = 0$ and $q > 0$, the ACF cuts off after lag q , and the PACF tails off. If $q = 0$ and $p > 0$, the PACF cuts off after lag p , and the ACF tails off. If $p > 0$ and $q > 0$, both the ACF and PACF will tail off. Because we are dealing with estimates, it will not always be clear whether the sample ACF or PACF is tailing off or cutting off. Also, two models that are seemingly different can actually be very similar. With this in mind, we should not worry about being so precise at this stage of the model fitting. At this stage, a few preliminary values of p , d , and q should be at hand, and we can start estimating the parameters.

Example 3.38 Analysis of GNP Data

In this example, we consider the analysis of quarterly U.S. GNP from 1947(1) to 2002(3), $n = 223$ observations. The data are real U.S. gross national product in billions of chained 1996 dollars and have been seasonally adjusted. The data were obtained from the Federal Reserve Bank of St. Louis (<http://research.stlouisfed.org/>). Figure 3.14 shows a plot of the data, say, y_t . Because strong trend hides any other effect, it is not clear from Figure 3.14 that the variance is increasing with time. For the purpose of demonstration, the sample ACF of the data is displayed in Figure 3.15. Figure 3.16 shows the first difference of the data, ∇y_t , and now that the trend has been removed we are able to notice that the variability in the second half of the data is larger than in the first half of the data. Also, it appears as though a trend is still present after differencing. The growth rate, say, $x_t = \nabla \log(y_t)$, is plotted in Figure 3.17, and, appears to be a stable process. Moreover, we may interpret the values of x_t as the percentage quarterly growth of U.S. GNP.




Fig. 3.15. Sample ACF of the GNP data. Lag is in terms of years.




Fig. 3.16. First difference of the U.S. GNP data.

The sample ACF and PACF of the quarterly growth rate are plotted in [Figure 3.18](#). Inspecting the sample ACF and PACF, we might feel that the ACF is cutting off at lag 2 and the PACF is tailing off. This would suggest the GNP growth rate follows an MA(2) process, or log GNP follows an ARIMA(0, 1, 2) model. Rather than focus on one model, we will also suggest that it appears that the ACF is tailing off and the PACF is cutting off at lag 1. This suggests an AR(1) model for the growth rate, or ARIMA(1, 1, 0) for log GNP. As a preliminary analysis, we will fit both models.

Using MLE to fit the MA(2) model for the growth rate, x_t , the estimated model is

$$x_t = .008_{(.001)} + .303_{(.065)}\widehat{w}_{t-1} + .204_{(.064)}\widehat{w}_{t-2} + \widehat{w}_t, \quad (3.151)$$




Fig. 3.17. U.S. GNP quarterly growth rate.




Fig. 3.18. Sample ACF and PACF of the GNP quarterly growth rate. Lag is in terms of years.

where $\hat{\sigma}_w = .0094$ is based on 219 degrees of freedom. The values in parentheses are the corresponding estimated standard errors. All of the regression coefficients are significant, including the constant. *We make a special note of this because, as a default, some computer packages do not fit a constant in a differenced model.* That is, these packages assume, by default, that there is no drift. In this example, not including a constant leads to the wrong conclusions about the nature of the U.S. economy. Not including a constant assumes the average quarterly growth rate is zero, whereas the U.S. GNP average quarterly growth rate is about 1% (which can be seen easily in [Figure 3.17](#)). We leave it to the reader to investigate what happens when the constant is not included.

The estimated AR(1) model is

$$x_t = .008_{(.001)} (1 - .347) + .347_{(.063)} x_{t-1} + \hat{w}_t, \quad (3.152)$$

where $\widehat{\sigma}_w = .0095$ on 220 degrees of freedom; note that the constant in (3.152) is $.008(1 - .347) = .005$.

We will discuss diagnostics next, but assuming both of these models fit well, how are we to reconcile the apparent differences of the estimated models (3.151) and (3.152)? In fact, the fitted models are nearly the same. To show this, consider an AR(1) model of the form in (3.152) without a constant term; that is,

$$x_t = .35x_{t-1} + w_t,$$

and write it in its causal form, $x_t = \sum_{j=0}^{\infty} \psi_j w_{t-j}$, where we recall $\psi_j = .35^j$. Thus, $\psi_0 = 1, \psi_1 = .350, \psi_2 = .123, \psi_3 = .043, \psi_4 = .015, \psi_5 = .005, \psi_6 = .002, \psi_7 = .001, \psi_8 = 0, \psi_9 = 0, \psi_{10} = 0$, and so forth. Thus,

$$x_t \approx .35w_{t-1} + .12w_{t-2} + w_t,$$

which is similar to the fitted MA(2) model in (3.152).

The analysis can be performed in R as follows.

```
plot(gnp)
acf2(gnp, 50)
gnpgr = diff(log(gnp)) # growth rate
plot(gnpgr)
acf2(gnpgr, 24)
sarima(gnpgr, 1, 0, 0) # AR(1)
sarima(gnpgr, 0, 0, 2) # MA(2)
ARMAtoMA(ar=.35, ma=0, 10) # prints psi-weights
```

The next step in model fitting is diagnostics. This investigation includes the analysis of the residuals as well as model comparisons. Again, the first step involves a time plot of the innovations (or residuals), $x_t - \widehat{x}_t^{t-1}$, or of the standardized innovations

$$e_t = (x_t - \widehat{x}_t^{t-1}) / \sqrt{\widehat{P}_t^{t-1}}, \quad (3.153)$$

where $\widehat{x}_t^{t-1}$ is the one-step-ahead prediction of x_t based on the fitted model and $\widehat{P}_t^{t-1}$ is the estimated one-step-ahead error variance. If the model fits well, the standardized residuals should behave as an iid sequence with mean zero and variance one. The time plot should be inspected for any obvious departures from this assumption. Unless the time series is Gaussian, it is not enough that the residuals are uncorrelated. For example, it is possible in the non-Gaussian case to have an uncorrelated process for which values contiguous in time are highly dependent. As an example, we mention the family of GARCH models that are discussed in Chapter 5.

Investigation of marginal normality can be accomplished visually by looking at a histogram of the residuals. In addition to this, a normal probability plot or a Q-Q plot can help in identifying departures from normality. See Johnson and Wichern (1992, Chapter 4) for details of this test as well as additional tests for multivariate normality.

There are several tests of randomness, for example the runs test, that could be applied to the residuals. We could also inspect the sample autocorrelations of the residuals, say, $\widehat{\rho}_e(h)$, for any patterns or large values. Recall that, for a white noise sequence, the sample autocorrelations are approximately independently and normally

distributed with zero means and variances $1/n$. Hence, a good check on the correlation structure of the residuals is to plot $\hat{\rho}_e(h)$ versus h along with the error bounds of $\pm 2/\sqrt{n}$. The residuals from a model fit, however, will not quite have the properties of a white noise sequence and the variance of $\hat{\rho}_e(h)$ can be much less than $1/n$. Details can be found in Box and Pierce (1970) and McLeod (1978). This part of the diagnostics can be viewed as a visual inspection of $\hat{\rho}_e(h)$ with the main concern being the detection of obvious departures from the independence assumption.

In addition to plotting $\hat{\rho}_e(h)$, we can perform a general test that takes into consideration the magnitudes of $\hat{\rho}_e(h)$ as a group. For example, it may be the case that, individually, each $\hat{\rho}_e(h)$ is small in magnitude, say, each one is just slightly less than $2/\sqrt{n}$ in magnitude, but, collectively, the values are large. The Ljung–Box–Pierce Q-statistic given by

$$Q = n(n+2) \sum_{h=1}^H \frac{\hat{\rho}_e^2(h)}{n-h} \quad (3.154)$$

can be used to perform such a test. The value H in (3.154) is chosen somewhat arbitrarily, typically, $H = 20$. Under the null hypothesis of model adequacy, asymptotically ($n \rightarrow \infty$), $Q \sim \chi_{H-p-q}^2$. Thus, we would reject the null hypothesis at level α if the value of Q exceeds the $(1 - \alpha)$ -quantile of the χ_{H-p-q}^2 distribution. Details can be found in Box and Pierce (1970), Ljung and Box (1978), and Davies et al. (1977). The basic idea is that if w_t is white noise, then by [Property 1.1](#), $n\hat{\rho}_w^2(h)$, for $h = 1, \dots, H$, are asymptotically independent χ_1^2 random variables. This means that $n \sum_{h=1}^H \hat{\rho}_w^2(h)$ is approximately a χ_H^2 random variable. Because the test involves the ACF of residuals from a model fit, there is a loss of $p + q$ degrees of freedom; the other values in (3.154) are used to adjust the statistic to better match the asymptotic chi-squared distribution.

Example 3.39 Diagnostics for GNP Growth Rate Example

We will focus on the MA(2) fit from [Example 3.38](#); the analysis of the AR(1) residuals is similar. [Figure 3.19](#) displays a plot of the standardized residuals, the ACF of the residuals, a boxplot of the standardized residuals, and the p-values associated with the Q-statistic, (3.154), at lags $H = 3$ through $H = 20$ (with corresponding degrees of freedom $H - 2$).

Inspection of the time plot of the standardized residuals in [Figure 3.19](#) shows no obvious patterns. Notice that there are outliers, however, with a few values exceeding 3 standard deviations in magnitude. The ACF of the standardized residuals shows no apparent departure from the model assumptions, and the Q-statistic is never significant at the lags shown. The normal Q-Q plot of the residuals shows departure from normality at the tails due to the outliers that occurred primarily in the 1950s and the early 1980s.

The model appears to fit well except for the fact that a distribution with heavier tails than the normal distribution should be employed. We discuss some possibilities




Fig. 3.19. Diagnostics of the residuals from MA(2) fit on GNP growth rate.

in Chapters 5 and 6. The diagnostics shown in [Figure 3.19](#) are a by-product of the `sarima` command from the previous example.⁹

Example 3.40 Diagnostics for the Glacial Varve Series

In [Example 3.32](#), we fit an ARIMA(0, 1, 1) model to the logarithms of the glacial varve data and there appears to be a small amount of autocorrelation left in the residuals and the Q-tests are all significant; see [Figure 3.20](#).

To adjust for this problem, we fit an ARIMA(1, 1, 1) to the logged varve data and obtained the estimates

$$\hat{\phi} = .23_{(.05)}, \hat{\theta} = -.89_{(.03)}, \hat{\sigma}_w^2 = .23.$$

Hence the AR term is significant. The Q-statistic p-values for this model are also displayed in [Figure 3.20](#), and it appears this model fits the data well.

As previously stated, the diagnostics are byproducts of the individual `sarima` runs. We note that we did not fit a constant in either model because there is no apparent drift in the differenced, logged varve series. This fact can be verified by noting the constant is not significant when the command `no.constant=TRUE` is removed in the code:

⁹ The script `tsdiag` is available in R to run diagnostics for an ARIMA object, however, the script has errors and we do not recommend using it.




Fig. 3.20. Q-statistic p-values for the ARIMA(0, 1, 1) fit (top) and the ARIMA(1, 1, 1) fit (bottom) to the logged varve data.

```
sarima(log(varve), 0, 1, 1, no.constant=TRUE) # ARIMA(0, 1, 1)
sarima(log(varve), 1, 1, 1, no.constant=TRUE) # ARIMA(1, 1, 1)
```

In [Example 3.38](#), we have two competing models, an AR(1) and an MA(2) on the GNP growth rate, that each appear to fit the data well. In addition, we might also consider that an AR(2) or an MA(3) might do better for forecasting. Perhaps combining both models, that is, fitting an ARMA(1, 2) to the GNP growth rate, would be the best. As previously mentioned, we have to be concerned with overfitting the model; it is not always the case that more is better. Overfitting leads to less-precise estimators, and adding more parameters may fit the data better but may also lead to bad forecasts. This result is illustrated in the following example.

Example 3.41 A Problem with Overfitting

[Figure 3.21](#) shows the U.S. population by official census, every ten years from 1910 to 1990, as points. If we use these nine observations to predict the future population, we can use an eight-degree polynomial so the fit to the nine observations is perfect. The model in this case is

$$x_t = \beta_0 + \beta_1 t + \beta_2 t^2 + \cdots + \beta_8 t^8 + w_t.$$

The fitted line, which is plotted in the figure, passes through the nine observations. The model predicts that the population of the United States will be close to zero in the year 2000, and will cross zero sometime in the year 2002!

The final step of model fitting is model choice or model selection. That is, we must decide which model we will retain for forecasting. The most popular techniques, AIC, AICc, and BIC, were described in §2.2 in the context of regression models.




Fig. 3.21. A perfect fit and a terrible forecast.

Example 3.42 Model Choice for the U.S. GNP Series

Returning to the analysis of the U.S. GNP data presented in [Example 3.38](#) and [Example 3.39](#), recall that two models, an AR(1) and an MA(2), fit the GNP growth rate well. To choose the final model, we compare the AIC, the AICc, and the BIC for both models. These values are a byproduct of the `sarima` runs displayed at the end of [Example 3.38](#), but for convenience, we display them again here (recall the growth rate data are in `gnpgr`):

```
sarima(gnpgr, 1, 0, 0) # AR(1)
  $AIC: -8.294403  $AICc: -8.284898  $BIC: -9.263748
sarima(gnpgr, 0, 0, 2) # MA(2)
  $AIC: -8.297693  $AICc: -8.287854  $BIC: -9.251711
```

The AIC and AICc both prefer the MA(2) fit, whereas the BIC prefers the simpler AR(1) model. It is often the case that the BIC will select a model of smaller order than the AIC or AICc. It would not be unreasonable in this case to retain the AR(1) because pure autoregressive models are easier to work with.

3.9 Regression with Autocorrelated Errors

In [§2.1](#), we covered the classical regression model with uncorrelated errors w_t . In this section, we discuss the modifications that might be considered when the errors are correlated. That is, consider the regression model

$$y_t = \sum_{j=1}^r \beta_j z_{tj} + x_t \quad (3.155)$$

where x_t is a process with some covariance function $\gamma_x(s, t)$. In ordinary least squares, the assumption is that x_t is white Gaussian noise, in which case $\gamma_x(s, t) = 0$ for $s \neq t$ and $\gamma_x(t, t) = \sigma^2$, independent of t . If this is not the case, then weighted least squares should be used.

Write the model in vector notation as in (??), $\mathbf{y} = \mathbf{Z}\boldsymbol{\beta} + \mathbf{x}$, and let $\Gamma = \{\gamma_x(s, t)\}$. Then $\Gamma^{-1/2}\mathbf{y} = \Gamma^{-1/2}\mathbf{Z}\boldsymbol{\beta} + \Gamma^{-1/2}\mathbf{x}$, so that we can write the model as $\mathbf{y}^* = \mathbf{Z}^*\boldsymbol{\beta} + \boldsymbol{\delta}$, where the covariance matrix of $\boldsymbol{\delta}$ is the identity. It follows that the weighted estimate of $\boldsymbol{\beta}$ is $\widehat{\boldsymbol{\beta}}_w = (\mathbf{Z}'\Gamma^{-1}\mathbf{Z})^{-1}\mathbf{Z}'\Gamma^{-1}\mathbf{y}$, and the variance-covariance matrix of the estimator is $\text{var}(\widehat{\boldsymbol{\beta}}_w) = (\mathbf{Z}'\Gamma^{-1}\mathbf{Z})^{-1}$. If x_t is white noise, then $\Gamma = \sigma^2\mathbf{I}$ and these results reduce to the usual least squares results.

In the time series case, it is often possible to assume a stationary covariance structure for the error process x_t that corresponds to a linear process and try to find an ARMA representation for x_t . For example, if we have a pure AR(p) error, then

$$\phi(B)x_t = w_t,$$

and $\phi(B) = 1 - \phi_1 B - \dots - \phi_p B^p$ is the linear transformation that, when applied to the error process, produces the white noise w_t . Multiplying the regression equation through by the transformation $\phi(B)$ yields,

$$\underbrace{\phi(B)y_t}_{y_t^*} = \sum_{j=1}^r \beta_j \underbrace{\phi(B)z_{tj}}_{z_{tj}^*} + \underbrace{\phi(B)x_t}_{w_t},$$

and we are back to the linear regression model where the observations have been transformed so that $y_t^* = \phi(B)y_t$ is the dependent variable, $z_{tj}^* = \phi(B)z_{tj}$ for $j = 1, \dots, r$, are the independent variables, but the β s are the same as in the original model. For example, if $p = 1$, then $y_t^* = y_t - \phi y_{t-1}$ and $z_{tj}^* = z_{tj} - \phi z_{t-1,j}$.

In the AR case, we may set up the least squares problem as minimizing the error sum of squares

$$S(\boldsymbol{\phi}, \boldsymbol{\beta}) = \sum_{t=1}^n w_t^2 = \sum_{t=1}^n \left[\phi(B)y_t - \sum_{j=1}^r \beta_j \phi(B)z_{tj} \right]^2$$

with respect to all the parameters, $\boldsymbol{\phi} = \{\phi_1, \dots, \phi_p\}$ and $\boldsymbol{\beta} = \{\beta_1, \dots, \beta_r\}$. Of course, this is done using numerical methods.

If the error process is ARMA(p, q), i.e., $\phi(B)x_t = \theta(B)w_t$, then in the above discussion, we transform by $\pi(B)x_t = w_t$, where $\pi(B) = \theta(B)^{-1}\phi(B)$. In this case the error sum of squares also depends on $\boldsymbol{\theta} = \{\theta_1, \dots, \theta_q\}$:

$$S(\boldsymbol{\phi}, \boldsymbol{\theta}, \boldsymbol{\beta}) = \sum_{t=1}^n w_t^2 = \sum_{t=1}^n \left[\pi(B)y_t - \sum_{j=1}^r \beta_j \pi(B)z_{tj} \right]^2$$

At this point, the main problem is that we do not typically know the behavior of the noise x_t prior to the analysis. An easy way to tackle this problem was first presented in Cochrane and Orcutt (1949), and with the advent of cheap computing is modernized below:

- (i) First, run an ordinary regression of y_t on $z_{t1}, \dots, z_{tr}$ (acting as if the errors are uncorrelated). Retain the residuals, $\widehat{x}_t = y_t - \sum_{j=1}^r \widehat{\beta}_j z_{tj}$.




Fig. 3.22. Sample ACF and PACF of the mortality residuals indicating an AR(2) process.

- (ii) Identify ARMA model(s) for the residuals $\hat{x}_t$.
- (iii) Run weighted least squares (or MLE) on the regression model with autocorrelated errors using the model specified in step (ii).
- (iv) Inspect the residuals $\hat{w}_t$ for whiteness, and adjust the model if necessary.

Example 3.43 Mortality, Temperature and Pollution

We consider the analyses presented in [Example 2.2](#), relating mean adjusted temperature T_t , and particulate levels P_t to cardiovascular mortality M_t . We consider the regression model

$$M_t = \beta_1 + \beta_2 t + \beta_3 T_t + \beta_4 T_t^2 + \beta_5 P_t + x_t, \quad (3.156)$$

where, for now, we assume that x_t is white noise. The sample ACF and PACF of the residuals from the ordinary least squares fit of (3.156) are shown in [Figure 3.22](#), and the results suggest an AR(2) model for the residuals.

Our next step is to fit the correlated error model (3.156), but where x_t is AR(2),

$$x_t = \phi_1 x_{t-1} + \phi_2 x_{t-2} + w_t$$

and w_t is white noise. The model can be fit using the `sarima` function as follows (partial output shown).

```
trend = time(cmort); temp = tempr - mean(tempr); temp2 = tempr^2
fit = lm(cmort~trend + temp + temp2 + part, na.action=NULL)
acf2(resid(fit), 52) # implies AR2
sarima(cmort, 2,0,0, xreg=cbind(trend,temp,temp2,part))
```



```

Coefficients:
      ar1      ar2 intercept      trend      temp      temp2      part
    0.3848  0.4326    80.2116   -1.5165   -0.0190   0.0154   0.1545
s.e.    0.0436  0.0400     1.8072    0.4226    0.0495   0.0020   0.0272
sigma^2 estimated as 26.01: loglikelihood = -1549.04, aic = 3114.07

```

The residual analysis output from [sarima](#) (not shown) shows no obvious departure of the residuals from whiteness.

3.10 Multiplicative Seasonal ARIMA Models

In this section, we introduce several modifications made to the ARIMA model to account for seasonal and nonstationary behavior. Often, the dependence on the past tends to occur most strongly at multiples of some underlying seasonal lag s . For example, with monthly economic data, there is a strong yearly component occurring at lags that are multiples of $s = 12$, because of the strong connections of all activity to the calendar year. Data taken quarterly will exhibit the yearly repetitive period at $s = 4$ quarters. Natural phenomena such as temperature also have strong components corresponding to seasons. Hence, the natural variability of many physical, biological, and economic processes tends to match with seasonal fluctuations. Because of this, it is appropriate to introduce autoregressive and moving average polynomials that identify with the seasonal lags. The resulting pure seasonal autoregressive moving average model, say, $\text{ARMA}(P, Q)_s$, then takes the form

$$\Phi_P(B^s)x_t = \Theta_Q(B^s)w_t, \quad (3.157)$$

where the operators

$$\Phi_P(B^s) = 1 - \Phi_1 B^s - \Phi_2 B^{2s} - \dots - \Phi_P B^{Ps} \quad (3.158)$$

and

$$\Theta_Q(B^s) = 1 + \Theta_1 B^s + \Theta_2 B^{2s} + \dots + \Theta_Q B^{Qs} \quad (3.159)$$

are the **seasonal autoregressive operator** and the **seasonal moving average operator** of orders P and Q , respectively, with seasonal period s .

Analogous to the properties of nonseasonal ARMA models, the pure seasonal $\text{ARMA}(P, Q)_s$ is causal only when the roots of $\Phi_P(z^s)$ lie outside the unit circle, and it is invertible only when the roots of $\Theta_Q(z^s)$ lie outside the unit circle.

Example 3.44 A Seasonal AR Series

A first-order seasonal autoregressive series that might run over months could be written as

$$(1 - \Phi B^{12})x_t = w_t$$

or

$$x_t = \Phi x_{t-12} + w_t.$$

This model exhibits the series x_t in terms of past lags at the multiple of the yearly seasonal period $s = 12$ months. It is clear from the above form that estimation and




Fig. 3.23. Data generated from a seasonal ($s = 12$) AR(1), and the true ACF and PACF of the model $x_t = .9x_{t-12} + w_t$.

forecasting for such a process involves only straightforward modifications of the unit lag case already treated. In particular, the causal condition requires $|\Phi| < 1$.

We simulated 3 years of data from the model with $\Phi = .9$, and exhibit the *theoretical* ACF and PACF of the model. See [Figure 3.23](#).

```
set.seed(666)
phi = c(rep(0,11),.9)
sAR = arima.sim(list(order=c(12,0,0), ar=phi), n=37)
sAR = ts(sAR, freq=12)
layout(matrix(c(1,2, 1,3), nc=2))
par(mar=c(3,3,2,1), mgp=c(1.6,.6,0))
plot(sAR, axes=FALSE, main='seasonal AR(1)', xlab="year", type='c')
Months = c("J","F","M","A","M","J","J","A","S","O","N","D")
points(sAR, pch=Months, cex=1.25, font=4, col=1:4)
axis(1, 1:4); abline(v=1:4, lty=2, col='#cccccc')
axis(2); box()
ACF = ARMAacf(ar=phi, ma=0, 100)
PACF = ARMAacf(ar=phi, ma=0, 100, pacf=TRUE)
plot(ACF,type="h", xlab="lag", ylim=c(-.1,1)); abline(h=0)
plot(PACF, type="h", xlab="lag", ylim=c(-.1,1)); abline(h=0)
```

For the first-order seasonal ($s = 12$) MA model, $x_t = w_t + \Theta w_{t-12}$, it is easy to verify that

$$\begin{aligned}\gamma(0) &= (1 + \Theta^2)\sigma^2 \\ \gamma(\pm 12) &= \Theta\sigma^2 \\ \gamma(h) &= 0, \quad \text{otherwise.}\end{aligned}$$

Table 3.3. Behavior of the ACF and PACF for Pure SARMA Models

	$AR(P)_s$	$MA(Q)_s$	$ARMA(P, Q)_s$
ACF*	Tails off at lags ks , $k = 1, 2, \dots$,	Cuts off after lag Qs	Tails off at lags ks
PACF*	Cuts off after lag Ps	Tails off at lags ks $k = 1, 2, \dots$,	Tails off at lags ks

*The values at nonseasonal lags $h \neq ks$, for $k = 1, 2, \dots$, are zero.

Thus, the only nonzero correlation, aside from lag zero, is

$$\rho(\pm 12) = \theta / (1 + \theta^2).$$

For the first-order seasonal ($s = 12$) AR model, using the techniques of the nonseasonal AR(1), we have

$$\begin{aligned} \gamma(0) &= \sigma^2 / (1 - \Phi^2) \\ \gamma(\pm 12k) &= \sigma^2 \Phi^k / (1 - \Phi^2) \quad k = 1, 2, \dots \\ \gamma(h) &= 0, \quad \text{otherwise.} \end{aligned}$$

In this case, the only non-zero correlations are

$$\rho(\pm 12k) = \Phi^k, \quad k = 0, 1, 2, \dots$$

These results can be verified using the general result that $\gamma(h) = \Phi \gamma(h - 12)$, for $h \geq 1$. For example, when $h = 1$, $\gamma(1) = \Phi \gamma(11)$, but when $h = 11$, we have $\gamma(11) = \Phi \gamma(1)$, which implies that $\gamma(1) = \gamma(11) = 0$. In addition to these results, the PACF have the analogous extensions from nonseasonal to seasonal models. These results are demonstrated in [Figure 3.23](#).

As an initial diagnostic criterion, we can use the properties for the pure seasonal autoregressive and moving average series listed in [Table 3.3](#). These properties may be considered as generalizations of the properties for nonseasonal models that were presented in [Table 3.1](#).

In general, we can combine the seasonal and nonseasonal operators into a multiplicative seasonal autoregressive moving average model, denoted by $ARMA(p, q) \times (P, Q)_s$, and write

$$\Phi_P(B^s) \phi(B) x_t = \Theta_Q(B^s) \theta(B) w_t \quad (3.160)$$

as the overall model. Although the diagnostic properties in [Table 3.3](#) are not strictly true for the overall mixed model, the behavior of the ACF and PACF tends to show rough patterns of the indicated form. In fact, for mixed models, we tend to see a mixture of the facts listed in [Table 3.1](#) and [Table 3.3](#). In fitting such models, focusing on the seasonal autoregressive and moving average components first generally leads to more satisfactory results.

Example 3.45 A Mixed Seasonal Model

Consider an $\text{ARMA}(0, 1) \times (1, 0)_{12}$ model

$$x_t = \Phi x_{t-12} + w_t + \theta w_{t-1},$$

where $|\Phi| < 1$ and $|\theta| < 1$. Then, because x_{t-12} , w_t , and w_{t-1} are uncorrelated, and x_t is stationary, $\gamma(0) = \Phi^2 \gamma(0) + \sigma_w^2 + \theta^2 \sigma_w^2$, or

$$\gamma(0) = \frac{1 + \theta^2}{1 - \Phi^2} \sigma_w^2.$$

In addition, multiplying the model by x_{t-h} , $h > 0$, and taking expectations, we have $\gamma(1) = \Phi \gamma(11) + \theta \sigma_w^2$, and $\gamma(h) = \Phi \gamma(h - 12)$, for $h \geq 2$. Thus, the ACF for this model is

$$\begin{aligned} \rho(12h) &= \Phi^h \quad h = 1, 2, \dots \\ \rho(12h - 1) &= \rho(12h + 1) = \frac{\theta}{1 + \theta^2} \Phi^h \quad h = 0, 1, 2, \dots, \\ \rho(h) &= 0, \quad \text{otherwise.} \end{aligned}$$

The ACF and PACF for this model, with $\Phi = .8$ and $\theta = -.5$, are shown in **Figure 3.24**. These type of correlation relationships, although idealized here, are typically seen with seasonal data.

To reproduce **Figure 3.24** in R, use the following commands:

```
phi = c(rep(0,11),.8)
ACF = ARMAacf(ar=phi, ma=-.5, 50)[-1]      # [-1] removes 0 lag
PACF = ARMAacf(ar=phi, ma=-.5, 50, pacf=TRUE)
par(mfrow=c(1,2))
plot(ACF, type="h", xlab="lag", ylim=c(-.4,.8)); abline(h=0)
plot(PACF, type="h", xlab="lag", ylim=c(-.4,.8)); abline(h=0)
```

Seasonal persistence occurs when the process is nearly periodic in the season. For example, with average monthly temperatures over the years, each January would be approximately the same, each February would be approximately the same, and so on. In this case, we might think of average monthly temperature x_t as being modeled as

$$x_t = S_t + w_t,$$

where S_t is a seasonal component that varies a little from one year to the next, according to a random walk,

$$S_t = S_{t-12} + v_t.$$

In this model, w_t and v_t are uncorrelated white noise processes. The tendency of data to follow this type of model will be exhibited in a sample ACF that is large and decays very slowly at lags $h = 12k$, for $k = 1, 2, \dots$. If we subtract the effect of successive years from each other, we find that

$$(1 - B^{12})x_t = x_t - x_{t-12} = v_t + w_t - w_{t-12}.$$




Fig. 3.24. ACF and PACF of the mixed seasonal ARMA model $x_t = .8x_{t-12} + w_t - .5w_{t-1}$.

This model is a stationary $MA(1)_{12}$, and its ACF will have a peak only at lag 12. In general, seasonal differencing can be indicated when the ACF decays slowly at multiples of some season s , but is negligible between the periods. Then, a seasonal difference of order D is defined as

$$\nabla_s^D x_t = (1 - B^s)^D x_t, \quad (3.161)$$

where $D = 1, 2, \dots$, takes positive integer values. Typically, $D = 1$ is sufficient to obtain seasonal stationarity. Incorporating these ideas into a general model leads to the following definition.

Definition 3.12 *The multiplicative seasonal autoregressive integrated moving average model, or SARIMA model is given by*

$$\Phi_P(B^s)\phi(B)\nabla_s^D\nabla^d x_t = \delta + \Theta_Q(B^s)\theta(B)w_t, \quad (3.162)$$

where w_t is the usual Gaussian white noise process. The general model is denoted as $ARIMA(p, d, q) \times (P, D, Q)_s$. The ordinary autoregressive and moving average components are represented by polynomials $\phi(B)$ and $\theta(B)$ of orders p and q , respectively, and the seasonal autoregressive and moving average components by $\Phi_P(B^s)$ and $\Theta_Q(B^s)$ of orders P and Q and ordinary and seasonal difference components by $\nabla^d = (1 - B)^d$ and $\nabla_s^D = (1 - B^s)^D$.

Example 3.46 An SARIMA Model

Consider the following model, which often provides a reasonable representation for seasonal, nonstationary, economic time series. We exhibit the equations for the model, denoted by $ARIMA(0, 1, 1) \times (0, 1, 1)_{12}$ in the notation given above, where the seasonal fluctuations occur every 12 months. Then, with $\delta = 0$, the model (3.162) becomes

$$\nabla_{12}\nabla x_t = \Theta(B^{12})\theta(B)w_t$$

or

$$(1 - B^{12})(1 - B)x_t = (1 + \Theta B^{12})(1 + \theta B)w_t. \quad (3.163)$$

Expanding both sides of (3.163) leads to the representation

$$(1 - B - B^{12} + B^{13})x_t = (1 + \theta B + \Theta B^{12} + \Theta\theta B^{13})w_t,$$

or in difference equation form

$$x_t = x_{t-1} + x_{t-12} - x_{t-13} + w_t + \theta w_{t-1} + \Theta w_{t-12} + \Theta\theta w_{t-13}.$$

Note that the multiplicative nature of the model implies that the coefficient of w_{t-13} is the product of the coefficients of w_{t-1} and w_{t-12} rather than a free parameter. The multiplicative model assumption seems to work well with many seasonal time series data sets while reducing the number of parameters that must be estimated.

Selecting the appropriate model for a given set of data from all of those represented by the general form (3.162) is a daunting task, and we usually think first in terms of finding difference operators that produce a roughly stationary series and then in terms of finding a set of simple autoregressive moving average or multiplicative seasonal ARMA to fit the resulting residual series. Differencing operations are applied first, and then the residuals are constructed from a series of reduced length. Next, the ACF and the PACF of these residuals are evaluated. Peaks that appear in these functions can often be eliminated by fitting an autoregressive or moving average component in accordance with the general properties of Table 3.1 and Table 3.3. In considering whether the model is satisfactory, the diagnostic techniques discussed in §3.8 still apply.

Example 3.47 Air Passengers

We consider the R data set `AirPassengers`, which are the monthly totals of international airline passengers, 1949 to 1960, taken from Box & Jenkins (1970). Various plots of the data and transformed data are shown in Figure 3.25 and were obtained as follows:

```
x = AirPassengers
lx = log(x); dlx = diff(lx); ddlx = diff(dlx, 12)
plot.ts(cbind(x, lx, dlx, ddlx), main="")
# below of interest for showing seasonal RW (not shown here):
par(mfrow=c(2,1))
monthplot(dlx); monthplot(ddlx)
```

Note that `x` is the original series, which shows trend plus increasing variance. The logged data are in `lx`, and the transformation stabilizes the variance. The logged data are then differenced to remove trend, and are stored in `dlx`. It is clear there is still persistence in the seasons (i.e., $dlx_t \approx dlx_{t-12}$), so that a twelfth-order difference is applied and stored in `ddlx`. The transformed data appears to be stationary and we are now ready to fit a model.

The sample ACF and PACF of `ddlx` ($\nabla_{12}\nabla \log x_t$) are shown in Figure 3.26. The R code is:




Fig. 3.25. R data set [AirPassengers](#), which are the monthly totals of international airline passengers x , and the transformed data: $lx = \log x_t$, $dlx = \nabla \log x_t$, and $ddlx = \nabla_{12} \nabla \log x_t$.

```
acf2(ddlx, 50)
```

Between Seasons: It appears that at the seasons, the ACF is cutting off a lag $1s$ ($s = 12$), whereas the PACF is tailing off at lags $1s, 2s, 3s, 4s, \dots$. These results implies an $SMA(1)$, $P = 0$, $Q = 1$, in the season ($s = 12$).

Within Seasons: Inspecting the sample ACF and PACF at the lower lags, it appears as though both are tailing off. This suggests an $ARMA(1, 1)$ within the seasons, $p = q = 1$.

Thus, we first try an $ARIMA(1, 1, 1) \times (0, 1, 1)_{12}$ on the logged data:

```
sarima(lx, 1,1,1, 0,1,1,12)
Coefficients:
      ar1      ma1      sma1
    0.1960 -0.5784 -0.5643
s.e.  0.2475  0.2132  0.0747
sigma^2 estimated as 0.001341
$AIC -5.5726 $AICc -5.556713 $BIC -6.510729
```

However, the AR parameter is not significant, so we should try dropping one parameter from the within seasons part. In this case, we try both an $ARIMA(0, 1, 1) \times (0, 1, 1)_{12}$ and an $ARIMA(1, 1, 0) \times (0, 1, 1)_{12}$ model:




Fig. 3.26. Sample ACF and PACF of `ddlx` ($\nabla_{12} \nabla \log x_t$).

```
sarima(lx, 0,1,1, 0,1,1,12)
Coefficients:
      mal      sma1
    -0.4018  -0.5569
s.e.    0.0896   0.0731
sigma^2 estimated as 0.001348
$AIC -5.58133  $AICc -5.56625  $BIC -6.540082

sarima(lx, 1,1,0, 0,1,1,12)
Coefficients:
      ar1      sma1
    -0.3395  -0.5619
s.e.    0.0822   0.0748
sigma^2 estimated as 0.001367
$AIC -5.567081  $AICc -5.552002  $BIC -6.525834
```

All information criteria prefer the $\text{ARIMA}(0, 1, 1) \times (0, 1, 1)_{12}$ model, which is the model displayed in (3.163). The residual diagnostics are shown in Figure 3.27, and except for one or two outliers, the model seems to fit well.

Finally, we forecast the logged data out twelve months, and the results are shown in Figure 3.28.

```
sarima.for(lx, 12, 0,1,1, 0,1,1,12)
```

Example 3.48 Don't use `auto.arima`

The `auto.arima` function from the package `forecast` tries to fit a model in an automated fashion. It may be tempting, but automated analysis can lead to strange




Fig. 3.27. Residual analysis for the $\text{ARIMA}(0, 1, 1) \times (0, 1, 1)_{12}$ fit to the logged air passengers data set.

results and is no substitute for understanding what you're doing. Here is what happens when `auto.arima` is used for the previous example (assuming you know enough to log the data first).

```
require(forecast) # remember you have to download it first
auto.arima(log(AirPassengers))
Series: log(AirPassengers) ARIMA(0,1,1)(2,1,2)[12]
Coefficients:
      ma1      sar1      sar2      sma1      sma2
    -0.3632 -0.4852 -0.0933 -0.0204 -0.1803
s.e.  0.0949  0.1561  0.1107  0.1700  0.1117
---
sigma^2 estimated as 0.001258: log likelihood=251.52
AIC=-409.06 AICc=-408.39 BIC=-391.81
Warning message: Unable to fit final model using maximum likelihood. AIC
value approximated
```

Three of the five regression parameters are not significant.




Fig. 3.28. Twelve month forecast using the $\text{ARIMA}(0, 1, 1) \times (0, 1, 1)_{12}$ model on the logged air passenger data set.

Example 3.49 Regression with Lagged Variables (cont)

In [Example 2.8](#) we fit the model

$$R_t = \beta_1 + \beta_2 S_{t-6} + \beta_3 D_{t-6} + \beta_4 D_{t-6} S_{t-6} + w_t,$$

where R_t is Recruitment, S_t is SOI, and D_t is a dummy variable that is 0 if $S_t < 0$ and 1 otherwise. However, residual analysis indicates that the residuals are not white noise. The sample (P)ACF of the residuals indicates that an AR(2) model might be appropriate, which is similar to the results of [Example 3.43](#). Additional analysis suggests that a seasonal model might be needed, and after some investigation, we settled on an $\text{SARIMA}(2, 0, 0) \times (2, 0, 0)_{12}$ model for the residuals. Both AIC and BIC confirm that this model is the best. We only display the results of the final model for this example, but the code to carry out the complete analysis is listed below.

```
dummy = ifelse(soi<0,0,1)
fish = ts.intersect(rec, soiL6=lag(soi,-6), dL6=lag(dummy,-6), dframe=TRUE)
summary(fit<-lm(rec~soiL6*dL6, data=fish, na.action=NULL))
attach(fish)
plot(resid(fit))
acf2(resid(fit)) # indicates AR(2)
(fit1 = sarima(rec,2,0,0, xreg=cbind(soiL6,dL6,I(soiL6*dL6))))
acf2(resid(fit1)) # indicates seasonal AR
# try seasonal AR order 1 and then 2, which appears to be best
(fit2 = sarima(rec,2,0,0,1,0,0,12,xreg=cbind(soiL6,dL6,I(soiL6*dL6))))
(fit3 = sarima(rec,2,0,0,2,0,0,12,xreg=cbind(soiL6,dL6,I(soiL6*dL6))))
Coefficients: (rounded)
      ar1      ar2      sar1      sar2      intrcpt      soiL6      dL6      I(soiL6*dL6)
      1.349    -0.444    0.114    0.15         64         8.4     -2.25         -8.3
s.e.    0.044    0.045    0.047    0.05         6         2.2     0.95         2.9
sigma^2 estimated as 83.75: log likelihood = -1625.53, aic = 3269.07
$AIC          $AICc          $BIC
[1] 5.463671    [1] 5.469067    [1] 4.537095
```

Problems

Section 3.2

3.1 For an MA(1), $x_t = w_t + \theta w_{t-1}$, show that $|\rho_x(1)| \leq 1/2$ for any number θ . For which values of θ does $\rho_x(1)$ attain its maximum and minimum?

3.2 Let $\{w_t; t = 0, 1, \dots\}$ be a white noise process with variance σ_w^2 and let $|\phi| < 1$ be a constant. Consider the process $x_0 = w_0$, and

$$x_t = \phi x_{t-1} + w_t, \quad t = 1, 2, \dots$$

We might use this method to simulate an AR(1) process from simulated white noise.

- (a) Show that $x_t = \sum_{j=0}^t \phi^j w_{t-j}$ for any $t = 0, 1, \dots$
- (b) Find the $E(x_t)$.
- (c) Show that, for $t = 0, 1, \dots$,

$$\text{var}(x_t) = \frac{\sigma_w^2}{1 - \phi^2} (1 - \phi^{2(t+1)})$$

- (d) Show that, for $h \geq 0$,

$$\text{cov}(x_{t+h}, x_t) = \phi^h \text{var}(x_t)$$

- (e) Is x_t stationary?
- (f) Argue that, as $t \rightarrow \infty$, the process becomes stationary, so in a sense, x_t is “asymptotically stationary.”
- (g) Comment on how you could use these results to simulate n observations of a stationary Gaussian AR(1) model from simulated iid $N(0, 1)$ values.
- (h) Now suppose $x_0 = w_0/\sqrt{1 - \phi^2}$. Is this process stationary? *Hint:* Show $\text{var}(x_t)$ is constant.

3.3 Verify the calculations made in **Example 3.3** as follows.

- (a) Let $x_t = \phi x_{t-1} + w_t$ where $|\phi| > 1$ and $w_t \sim \text{iid } N(0, \sigma_w^2)$. Show $E(x_t) = 0$ and $\gamma_x(h) = \sigma_w^2 \phi^{-2} \phi^{-h} / (1 - \phi^{-2})$ for $h \geq 0$.
- (b) Let $y_t = \phi^{-1} y_{t-1} + v_t$ where $v_t \sim \text{iid } N(0, \sigma_w^2 \phi^{-2})$ and ϕ and σ_w are as in part (a). Argue that y_t is causal with the same mean function and autocovariance function as x_t .

3.4 Identify the following models as ARMA(p, q) models (watch out for parameter redundancy), and determine whether they are causal and/or invertible:

- (a) $x_t = .80x_{t-1} - .15x_{t-2} + w_t - .30w_{t-1}$.
- (b) $x_t = x_{t-1} + .50x_{t-2} + w_t - w_{t-1}$.

3.5 Verify the causal conditions for an AR(2) model given in (3.28). That is, show that an AR(2) is causal if and only if (3.28) holds.

Section 3.3

3.6 For the AR(2) model given by $x_t = -.9x_{t-2} + w_t$, find the roots of the autoregressive polynomial, and then sketch the ACF, $\rho(h)$.

3.7 For the AR(2) series shown below, use the results of [Example 3.9](#) to determine a set of difference equations that can be used to find the ACF $\rho(h)$, $h = 0, 1, \dots$; solve for the constants in the ACF using the initial conditions. Then plot the ACF values to lag 10 (use [ARMAacf](#) as a check on your answers).

(a) $x_t + 1.6x_{t-1} + .64x_{t-2} = w_t$.

(b) $x_t - .40x_{t-1} - .45x_{t-2} = w_t$.

(c) $x_t - 1.2x_{t-1} + .85x_{t-2} = w_t$.

Section 3.4

3.8 Verify the calculations for the autocorrelation function of an ARMA(1, 1) process given in [Example 3.13](#). Compare the form with that of the ACF for the ARMA(1, 0) and the ARMA(0, 1) series. Plot (or sketch) the ACFs of the three series on the same graph for $\phi = .6$, $\theta = .9$, and comment on the diagnostic capabilities of the ACF in this case.

3.9 Generate $n = 100$ observations from each of the three models discussed in [Problem 3.8](#). Compute the sample ACF for each model and compare it to the theoretical values. Compute the sample PACF for each of the generated series and compare the sample ACFs and PACFs with the general results given in [Table 3.1](#).

Section 3.5

3.10 Let x_t represent the cardiovascular mortality series ([cmort](#)) discussed in Chapter 2, [Example 2.2](#).

(a) Fit an AR(2) to x_t using linear regression as in [Example 3.17](#).

(b) Assuming the fitted model in (a) is the true model, find the forecasts over a four-week horizon, x_{n+m}^n , for $m = 1, 2, 3, 4$, and the corresponding 95% prediction intervals.

3.11 Consider the MA(1) series

$$x_t = w_t + \theta w_{t-1},$$

where w_t is white noise with variance σ_w^2 .

(a) Derive the minimum mean-square error one-step forecast based on the infinite past, and determine the mean-square error of this forecast.

(b) Let $\tilde{x}_{n+1}^n$ be the truncated one-step-ahead forecast as given in (3.92). Show that

$$E \left[(x_{n+1} - \tilde{x}_{n+1}^n)^2 \right] = \sigma^2 (1 + \theta^{2+2n}).$$

Compare the result with (a), and indicate how well the finite approximation works in this case.

3.12 In the context of equation (3.63), show that, if $\gamma(0) > 0$ and $\gamma(h) \rightarrow 0$ as $h \rightarrow \infty$, then Γ_n is positive definite.

3.13 Suppose x_t is stationary with zero mean and recall the definition of the PACF given by (3.55) and (3.56). That is, let

$$\epsilon_t = x_t - \sum_{i=1}^{h-1} a_i x_{t-i}$$

and

$$\delta_{t-h} = x_{t-h} - \sum_{j=1}^{h-1} b_j x_{t-j}$$

be the two residuals where $\{a_1, \dots, a_{h-1}\}$ and $\{b_1, \dots, b_{h-1}\}$ are chosen so that they minimize the mean-squared errors

$$E[\epsilon_t^2] \quad \text{and} \quad E[\delta_{t-h}^2].$$

The PACF at lag h was defined as the cross-correlation between ϵ_t and δ_{t-h} ; that is,

$$\phi_{hh} = \frac{E(\epsilon_t \delta_{t-h})}{\sqrt{E(\epsilon_t^2)E(\delta_{t-h}^2)}}.$$

Let R_h be the $h \times h$ matrix with elements $\rho(i-j)$, $i, j = 1, \dots, h$, and let $\boldsymbol{\rho}_h = (\rho(1), \rho(2), \dots, \rho(h))'$ be the vector of lagged autocorrelations, $\rho(h) = \text{corr}(x_{t+h}, x_t)$. Let $\tilde{\boldsymbol{\rho}}_h = (\rho(h), \rho(h-1), \dots, \rho(1))'$ be the reversed vector. In addition, let x_t^h denote the BLP of x_t given $\{x_{t-1}, \dots, x_{t-h}\}$:

$$x_t^h = \alpha_{h1}x_{t-1} + \dots + \alpha_{hh}x_{t-h},$$

as described in Property 3.3. Prove

$$\phi_{hh} = \frac{\rho(h) - \tilde{\boldsymbol{\rho}}_{h-1}' R_{h-1}^{-1} \boldsymbol{\rho}_h}{1 - \tilde{\boldsymbol{\rho}}_{h-1}' R_{h-1}^{-1} \tilde{\boldsymbol{\rho}}_{h-1}} = \alpha_{hh}.$$

In particular, this result proves Property 3.4.

Hint: Divide the prediction equations [see (3.63)] by $\gamma(0)$ and write the matrix equation in the partitioned form as

$$\begin{pmatrix} R_{h-1} & \tilde{\boldsymbol{\rho}}_{h-1} \\ \tilde{\boldsymbol{\rho}}_{h-1}' & \rho(0) \end{pmatrix} \begin{pmatrix} \boldsymbol{\alpha}_1 \\ \alpha_{hh} \end{pmatrix} = \begin{pmatrix} \boldsymbol{\rho}_{h-1} \\ \rho(h) \end{pmatrix},$$

where the $h \times 1$ vector of coefficients $\boldsymbol{\alpha} = (\alpha_{h1}, \dots, \alpha_{hh})'$ is partitioned as $\boldsymbol{\alpha} = (\boldsymbol{\alpha}'_1, \alpha_{hh})'$.

3.14 Suppose we wish to find a prediction function $g(x)$ that minimizes

$$MSE = E[(y - g(x))^2],$$

where x and y are jointly distributed random variables with density function $f(x, y)$.

(a) Show that MSE is minimized by the choice

$$g(x) = E(y \mid x).$$

Hint:

$$MSE = EE[(y - g(x))^2 \mid x].$$

(b) Apply the above result to the model

$$y = x^2 + z,$$

where x and z are independent zero-mean normal variables with variance one. Show that $MSE = 1$.

(c) Suppose we restrict our choices for the function $g(x)$ to linear functions of the form

$$g(x) = a + bx$$

and determine a and b to minimize MSE . Show that $a = 1$ and

$$b = \frac{E(xy)}{E(x^2)} = 0$$

and $MSE = 3$. What do you interpret this to mean?

3.15 For an AR(1) model, determine the general form of the m -step-ahead forecast x_{t+m}^t and show

$$E[(x_{t+m} - x_{t+m}^t)^2] = \sigma_w^2 \frac{1 - \phi^{2m}}{1 - \phi^2}.$$

3.16 Consider the ARMA(1,1) model discussed in [Example 3.7](#), equation (3.27); that is, $x_t = .9x_{t-1} + .5w_{t-1} + w_t$. Show that truncated prediction as defined in (3.91) is equivalent to truncated prediction using the recursive formula (3.92).

3.17 Verify statement (3.87), that for a fixed sample size, the ARMA prediction errors are correlated.

Section 3.6

3.18 Fit an AR(2) model to the cardiovascular mortality series ([cmort](#)) discussed in Chapter 2, [Example 2.2](#). using linear regression and using Yule–Walker.

(a) Compare the parameter estimates obtained by the two methods.

- (b) Compare the estimated standard errors of the coefficients obtained by linear regression with their corresponding asymptotic approximations, as given in **Property 3.10**.

3.19 Suppose $x_1, \dots, x_n$ are observations from an AR(1) process with $\mu = 0$.

- (a) Show the backcasts can be written as $x_t^n = \phi^{1-t} x_1$, for $t \leq 1$.
 (b) In turn, show, for $t \leq 1$, the backcasted errors are

$$\widehat{w}_t(\phi) = x_t^n - \phi x_{t-1}^n = \phi^{1-t} (1 - \phi^2) x_1.$$

- (c) Use the result of (b) to show $\sum_{t=-\infty}^1 \widehat{w}_t^2(\phi) = (1 - \phi^2) x_1^2$.
 (d) Use the result of (c) to verify the unconditional sum of squares, $S(\phi)$, can be written as $\sum_{t=-\infty}^n \widehat{w}_t^2(\phi)$.
 (e) Find x_t^{t-1} and r_t for $1 \leq t \leq n$, and show that

$$S(\phi) = \sum_{t=1}^n (x_t - x_t^{t-1})^2 / r_t.$$

3.20 Repeat the following numerical exercise three times. Generate $n = 500$ observations from the ARMA model given by

$$x_t = .9x_{t-1} + w_t - .9w_{t-1},$$

with $w_t \sim \text{iid } N(0, 1)$. Plot the simulated data, compute the sample ACF and PACF of the simulated data, and fit an ARMA(1, 1) model to the data. What happened and how do you explain the results?

3.21 Generate 10 realizations of length $n = 200$ each of an ARMA(1,1) process with $\phi = .9$, $\theta = .5$ and $\sigma^2 = 1$. Find the MLEs of the three parameters in each case and compare the estimators to the true values.

3.22 Generate $n = 50$ observations from a Gaussian AR(1) model with $\phi = .99$ and $\sigma_w = 1$. Using an estimation technique of your choice, compare the approximate asymptotic distribution of your estimate (the one you would use for inference) with the results of a bootstrap experiment (use $B = 200$).

3.23 Using **Example 3.31** as your guide, find the Gauss–Newton procedure for estimating the autoregressive parameter, ϕ , from the AR(1) model, $x_t = \phi x_{t-1} + w_t$, given data $x_1, \dots, x_n$. Does this procedure produce the unconditional or the conditional estimator? *Hint:* Write the model as $w_t(\phi) = x_t - \phi x_{t-1}$; your solution should work out to be a non-recursive procedure.

3.24 Consider the stationary series generated by

$$x_t = \alpha + \phi x_{t-1} + w_t + \theta w_{t-1},$$

where $E(x_t) = \mu$, $|\theta| < 1$, $|\phi| < 1$ and the w_t are iid random variables with zero mean and variance σ_w^2 .

- (a) Determine the mean as a function of α for the above model. Find the autocovariance and ACF of the process x_t , and show that the process is weakly stationary. Is the process strictly stationary?
- (b) Prove the limiting distribution as $n \rightarrow \infty$ of the sample mean,

$$\bar{x} = n^{-1} \sum_{t=1}^n x_t,$$

is normal, and find its limiting mean and variance in terms of α , ϕ , θ , and σ_w^2 . (Note: This part uses results from [Appendix A](#).)

3.25 A problem of interest in the analysis of geophysical time series involves a simple model for observed data containing a signal and a reflected version of the signal with unknown amplification factor a and unknown time delay δ . For example, the depth of an earthquake is proportional to the time delay δ for the P wave and its reflected form pP on a seismic record. Assume the signal, say s_t , is white and Gaussian with variance σ_s^2 , and consider the generating model

$$x_t = s_t + a s_{t-\delta}.$$

- (a) Prove the process x_t is stationary. If $|a| < 1$, show that

$$s_t = \sum_{j=0}^{\infty} (-a)^j x_{t-\delta j}$$

is a mean square convergent representation for the signal s_t , for $t = 1, \pm 1, \pm 2, \dots$

- (b) If the time delay δ is assumed to be known, suggest an approximate computational method for estimating the parameters a and σ_s^2 using maximum likelihood and the Gauss–Newton method.
- (c) If the time delay δ is an unknown integer, specify how we could estimate the parameters including δ . Generate a $n = 500$ point series with $a = .9$, $\sigma_w^2 = 1$ and $\delta = 5$. Estimate the integer time delay δ by searching over $\delta = 3, 4, \dots, 7$.

3.26 Forecasting with estimated parameters: Let $x_1, x_2, \dots, x_n$ be a sample of size n from a causal AR(1) process, $x_t = \phi x_{t-1} + w_t$. Let $\hat{\phi}$ be the Yule–Walker estimator of ϕ .

- (a) Show $\hat{\phi} - \phi = O_p(n^{-1/2})$. See [Appendix A](#) for the definition of $O_p(\cdot)$.
- (b) Let x_{n+1}^n be the one-step-ahead forecast of x_{n+1} given the data $x_1, \dots, x_n$, based on the known parameter, ϕ , and let $\hat{x}_{n+1}^n$ be the one-step-ahead forecast when the parameter is replaced by $\hat{\phi}$. Show $x_{n+1}^n - \hat{x}_{n+1}^n = O_p(n^{-1/2})$.

Section 3.7

3.27 Suppose

$$y_t = \beta_0 + \beta_1 t + \cdots + \beta_q t^q + x_t, \quad \beta_q \neq 0,$$

where x_t is stationary. First, show that $\nabla^k x_t$ is stationary for any $k = 1, 2, \dots$, and then show that $\nabla^k y_t$ is not stationary for $k < q$, but is stationary for $k \geq q$.

3.28 Verify that the IMA(1,1) model given in (3.147) can be inverted and written as (3.148).**3.29** For the ARIMA(1, 1, 0) model with drift, $(1 - \phi B)(1 - B)x_t = \delta + w_t$, let $y_t = (1 - B)x_t = \nabla x_t$.

(a) Noting that y_t is AR(1), show that, for $j \geq 1$,

$$y_{n+j}^n = \delta [1 + \phi + \cdots + \phi^{j-1}] + \phi^j y_n.$$

(b) Use part (a) to show that, for $m = 1, 2, \dots$,

$$x_{n+m}^n = x_n + \frac{\delta}{1 - \phi} \left[m - \frac{\phi(1 - \phi^m)}{(1 - \phi)} \right] + (x_n - x_{n-1}) \frac{\phi(1 - \phi^m)}{(1 - \phi)}.$$

Hint: From (a), $x_{n+j}^n - x_{n+j-1}^n = \delta \frac{1 - \phi^j}{1 - \phi} + \phi^j (x_n - x_{n-1})$. Now sum both sides over j from 1 to m .

(c) Use (3.144) to find P_{n+m}^n by first showing that $\psi_0^* = 1$, $\psi_1^* = (1 + \phi)$, and $\psi_j^* - (1 + \phi)\psi_{j-1}^* + \phi\psi_{j-2}^* = 0$ for $j \geq 2$, in which case $\psi_j^* = \frac{1 - \phi^{j+1}}{1 - \phi}$, for $j \geq 1$. Note that, as in Example 3.36, equation (3.144) is exact here.

3.30 For the logarithm of the glacial varve data, say, x_t , presented in Example 3.32, use the first 100 observations and calculate the EWMA, $\tilde{x}_{t+1}^t$, given in (3.150) for $t = 1, \dots, 100$, using $\lambda = .25, .50$, and $.75$, and plot the EWMA's and the data superimposed on each other. Comment on the results.

Section 3.8

3.31 In Example 3.39, we presented the diagnostics for the MA(2) fit to the GNP growth rate series. Using that example as a guide, complete the diagnostics for the AR(1) fit.**3.32** Crude oil prices in dollars per barrel are in oil; see Appendix R for more details. Fit an ARIMA(p, d, q) model to the growth rate performing all necessary diagnostics. Comment.**3.33** Fit an ARIMA(p, d, q) model to the global temperature data gtemp performing all of the necessary diagnostics. After deciding on an appropriate model, forecast (with limits) the next 10 years. Comment.

3.34 One of the series collected along with particulates, temperature, and mortality described in [Example 2.2](#) is the sulfur dioxide series, [so2](#). Fit an $\text{ARIMA}(p, d, q)$ model to the data, performing all of the necessary diagnostics. After deciding on an appropriate model, forecast the data into the future four time periods ahead (about one month) and calculate 95% prediction intervals for each of the four forecasts. Comment.

Section 3.9

3.35 Let S_t represent the monthly sales data in [sales](#) ($n = 150$), and let L_t be the leading indicator in [lead](#).

- Fit an ARIMA model to S_t , the monthly sales data. Discuss your model fitting in a step-by-step fashion, presenting your (A) initial examination of the data, (B) transformations, if necessary, (C) initial identification of the dependence orders and degree of differencing, (D) parameter estimation, (E) residual diagnostics and model choice.
- Use the CCF and lag plots between ∇S_t and ∇L_t to argue that a regression of ∇S_t on ∇L_{t-3} is reasonable. [Note that in `lag2.plot()`, the first named series is the one that gets lagged.]
- Fit the regression model $\nabla S_t = \beta_0 + \beta_1 \nabla L_{t-3} + x_t$, where x_t is an ARMA process (explain how you decided on your model for x_t). Discuss your results. [See [Example 3.49](#) for help on coding this problem.]

3.36 Redo [Problem 2.2](#) without assuming the error term is white noise.

Section 3.10

3.37 Consider the ARIMA model

$$x_t = w_t + \Theta w_{t-2}.$$

- Identify the model using the notation $\text{ARIMA}(p, d, q) \times (P, D, Q)_s$.
- Show that the series is invertible for $|\Theta| < 1$, and find the coefficients in the representation

$$w_t = \sum_{k=0}^{\infty} \pi_k x_{t-k}.$$

- Develop equations for the m -step ahead forecast, $\tilde{x}_{n+m}$, and its variance based on the infinite past, $x_n, x_{n-1}, \dots$.

3.38 Plot (or sketch) the ACF of the seasonal $\text{ARIMA}(0, 1) \times (1, 0)_{12}$ model with $\Phi = .8$ and $\theta = .5$.

3.39 Fit a seasonal ARIMA model of your choice to the unemployment data in [unemp](#). Use the estimated model to forecast the next 12 months.

3.40 Fit a seasonal ARIMA model of your choice to the U.S. Live Birth Series ([birth](#)). Use the estimated model to forecast the next 12 months.

3.41 Fit an appropriate seasonal ARIMA model to the log-transformed Johnson and Johnson earnings series ([jj](#)) of [Example 1.1](#). Use the estimated model to forecast the next 4 quarters.

The following problems require supplemental material given in [Appendix B](#).

3.42 Suppose $x_t = \sum_{j=1}^p \phi_j x_{t-j} + w_t$, where $\phi_p \neq 0$ and w_t is white noise such that w_t is uncorrelated with $\{x_k; k < t\}$. Use the Projection Theorem, [Theorem B.1](#), to show that, for $n > p$, the BLP of x_{n+1} on $\overline{\text{sp}}\{x_k, k \leq n\}$ is

$$\hat{x}_{n+1} = \sum_{j=1}^p \phi_j x_{n+1-j}.$$

3.43 Use the Projection Theorem to derive the Innovations Algorithm, [Property 3.6](#), equations (3.77)-(3.79). Then, use [Theorem B.2](#) to derive the m -step-ahead forecast results given in (3.80) and (3.81).

3.44 Consider the series $x_t = w_t - w_{t-1}$, where w_t is a white noise process with mean zero and variance σ_w^2 . Suppose we consider the problem of predicting x_{n+1} , based on only $x_1, \dots, x_n$. Use the Projection Theorem to answer the questions below.

(a) Show the best linear predictor is

$$x_{n+1}^n = -\frac{1}{n+1} \sum_{k=1}^n k x_k.$$

(b) Prove the mean square error is

$$E(x_{n+1} - x_{n+1}^n)^2 = \frac{n+2}{n+1} \sigma_w^2.$$

3.45 Use [Theorem B.2](#) and [Theorem B.3](#) to verify (3.116).

3.46 Prove [Theorem B.2](#).

3.47 Prove [Property 3.2](#).

Chapter 4

Spectral Analysis and Filtering

4.1 Introduction

The notion that a time series exhibits repetitive or regular behavior over time is of fundamental importance because it distinguishes time series analysis from classical statistics, which assumes complete independence over time. We have seen how dependence over time can be introduced through models that describe in detail the way certain empirical data behaves, even to the extent of producing forecasts based on the models. It is natural that models based on predicting the present as a regression on the past, such as are provided by the celebrated ARIMA or state-space forms, will be attractive to statisticians, who are trained to view nature in terms of linear models. In fact, the difference equations used to represent these kinds of models are simply the discrete versions of linear differential equations that may, in some instances, provide the ideal physical model for a certain phenomenon. An alternate version of the way nature behaves exists, however, and is based on a decomposition of an empirical series into its regular components.

In this chapter, we argue that the concept of regularity of a series can best be expressed in terms of periodic variations of the underlying phenomenon that produced the series, expressed as Fourier frequencies being driven by sines and cosines. Such a possibility was discussed in Chapters 1 and 2. From a regression point of view, we may imagine a system responding to various driving frequencies by producing linear combinations of sine and cosine functions. Expressed in these terms, the time domain approach may be thought of as regression of the present on the past, whereas the frequency domain approach may be considered as regression of the present on periodic sines and cosines.

Frequency domain approaches are the focus of this chapter. To illustrate the two methods for generating series with a single primary periodic component, consider [Figure 1.9](#), which was generated from a simple second-order autoregressive model, and the middle and bottom panels of [Figure 1.11](#), which were generated by adding a cosine wave with a period of 50 points to white noise. Both series exhibit strong periodic fluctuations, illustrating that both models can generate time series with regular behavior. As discussed in [Example 2.9](#), a fundamental objective of spectral

analysis is to identify the dominant frequencies in a series and to find an explanation of the system from which the measurements were derived.

Of course, the primary justification for any alternate model must lie in its potential for explaining the behavior of some empirical phenomenon. In this sense, an explanation involving only a few kinds of primary oscillations becomes simpler and more physically meaningful than a collection of parameters estimated for some selected difference equation. It is the tendency of observed data to show periodic kinds of fluctuations that justifies the use of frequency domain methods. Many of the examples in §1.2 are time series representing real phenomena that are driven by periodic components. The speech recording of the syllable *aa...hh* in [Figure 1.3](#) contains a complicated mixture of frequencies related to the opening and closing of the glottis. [Figure 1.5](#) shows the monthly SOI, which we later explain as a combination of two kinds of periodicities, a seasonal periodic component of 12 months and an El Niño component of about three to five years. Of fundamental interest is the return period of the El Niño phenomenon, which can have profound effects on local climate. Also of interest is whether the different periodic components of the new fish population depend on corresponding seasonal and El Niño-type oscillations. We introduce the coherence as a tool for relating the common periodic behavior of two series. Seasonal periodic components are often pervasive in economic time series; this phenomenon can be seen in the quarterly earnings series shown in [Figure 1.1](#). In [Figure 1.6](#), we see the extent to which various parts of the brain will respond to a periodic stimulus generated by having the subject do alternate left and right finger tapping. [Figure 1.7](#) shows series from an earthquake and a nuclear explosion. The relative amounts of energy at various frequencies for the two phases can produce statistics, useful for discriminating between earthquakes and explosions.

In this chapter, we summarize an approach to handling correlation generated in stationary time series that begins by transforming the series to the frequency domain. This simple linear transformation essentially matches sines and cosines of various frequencies against the underlying data and serves two purposes as discussed in [Example 2.9](#) and [Example 2.10](#). The periodogram that was introduced in [Example 2.10](#) has its population counterpart called the power spectrum, and its estimation is a main goal of spectral analysis. Another purpose of exploring this topic is statistical convenience resulting from the periodic components being nearly uncorrelated. This property facilitates writing likelihoods based on classical statistical methods.

An important part of analyzing data in the frequency domain, as well as the time domain, is the investigation and exploitation of the properties of the time-invariant linear filter. This special linear transformation is used similarly to linear regression in conventional statistics, and we use many of the same terms in the time series context. We have previously mentioned the coherence as a measure of the relation between two series at a given frequency, and we show later that this coherence also measures the performance of the best linear filter relating the two series. Linear filtering can also be an important step in isolating a signal embedded in noise. For example, the lower panels of [Figure 1.11](#) contain a signal contaminated with an additive noise, whereas the upper panel contains the pure signal. It might also be appropriate to ask whether a linear filter transformation exists that could be applied to the lower panel to produce

a series closer to the signal in the upper panel. The use of filtering for reducing noise will also be a part of the presentation in this chapter. We emphasize, throughout, the analogy between filtering techniques and conventional linear regression.

Many frequency scales will often coexist, depending on the nature of the problem. For example, in the Johnson & Johnson data set in [Figure 1.1](#), the predominant frequency of oscillation is one cycle per year (4 quarters), or .25 cycles per observation. The predominant frequency in the SOI and fish populations series in [Figure 1.5](#) is also one cycle per year, but this corresponds to 1 cycle every 12 months, or .083 cycles per observation. For simplicity, we measure frequency, ω , at cycles per time point and discuss the implications of certain frequencies in terms of the problem context. Of descriptive interest is the *period* of a time series, defined as the number of points in a cycle, i.e., $1/\omega$. Hence, the predominant period of the Johnson & Johnson series is $1/.25$ or 4 quarters per cycle, whereas the predominant period of the SOI series is 12 months per cycle.

4.2 Cyclical Behavior and Periodicity

As previously mentioned, we have already encountered the notion of periodicity in numerous examples in Chapters 1, 2 and 3. The general notion of periodicity can be made more precise by introducing some terminology. In order to define the rate at which a series oscillates, we first define a cycle as one complete period of a sine or cosine function defined over a unit time interval. As in [\(1.5\)](#), we consider the periodic process

$$x_t = A \cos(2\pi\omega t + \phi) \quad (4.1)$$

for $t = 0, \pm 1, \pm 2, \dots$, where ω is a frequency index, defined in cycles per unit time with A determining the height or *amplitude* of the function and ϕ , called the *phase*, determining the start point of the cosine function. We can introduce random variation in this time series by allowing the amplitude and phase to vary randomly.

As discussed in [Example 2.9](#), for purposes of data analysis, it is easier to use a trigonometric identity¹ and write [\(4.1\)](#) as

$$x_t = U_1 \cos(2\pi\omega t) + U_2 \sin(2\pi\omega t), \quad (4.2)$$

where $U_1 = A \cos \phi$ and $U_2 = -A \sin \phi$ are often taken to be normally distributed random variables. In this case, the amplitude is $A = \sqrt{U_1^2 + U_2^2}$ and the phase is $\phi = \tan^{-1}(-U_2/U_1)$. From these facts we can show that if, and only if, in [\(4.1\)](#), A and ϕ are independent random variables, where A^2 is chi-squared with 2 degrees of freedom, and ϕ is uniformly distributed on $(-\pi, \pi)$, then U_1 and U_2 are independent, standard normal random variables (see [Problem 4.2](#)).

The above random process is also a function of its frequency, defined by the parameter ω . The frequency is measured in cycles per unit time, or in cycles per point

¹ $\cos(\alpha \pm \beta) = \cos(\alpha) \cos(\beta) \mp \sin(\alpha) \sin(\beta)$.

in the above illustration. For $\omega = 1$, the series makes one cycle per time unit; for $\omega = .50$, the series makes a cycle every two time units; for $\omega = .25$, every four units, and so on. In general, for data that occur at discrete time points will need at least two points to determine a cycle, so the highest frequency of interest is .5 cycles per point. This frequency is called the *folding frequency* and defines the highest frequency that can be seen in discrete sampling. Higher frequencies sampled this way will appear at lower frequencies, called *aliases*; an example is the way a camera samples a rotating wheel on a moving automobile in a movie, in which the wheel appears to be rotating at a different rate. For example, movies are recorded at 24 frames per second. If the camera is filming a wheel that is rotating at the rate of 24 cycles per second (or 24 Hertz), the wheel will appear to stand still (that's about 110 miles per hour in case you were wondering).

Consider a generalization of (4.2) that allows mixtures of periodic series with multiple frequencies and amplitudes,

$$x_t = \sum_{k=1}^q [U_{k1} \cos(2\pi\omega_k t) + U_{k2} \sin(2\pi\omega_k t)], \quad (4.3)$$

where U_{k1}, U_{k2} , for $k = 1, 2, \dots, q$, are independent zero-mean random variables with variances σ_k^2 , and the ω_k are distinct frequencies. Notice that (4.3) exhibits the process as a sum of independent components, with variance σ_k^2 for frequency ω_k . Using the independence of the U s and the trig identity in footnote 1, it is easy to show² (Problem 4.3) that the autocovariance function of the process is

$$\gamma(h) = \sum_{k=1}^q \sigma_k^2 \cos(2\pi\omega_k h), \quad (4.4)$$

and we note the autocovariance function is the sum of periodic components with weights proportional to the variances σ_k^2 . Hence, x_t is a mean-zero stationary processes with variance

$$\gamma(0) = E(x_t^2) = \sum_{k=1}^q \sigma_k^2, \quad (4.5)$$

which exhibits the overall variance as a sum of variances of each of the component parts.

Example 4.1 A Periodic Series

Figure 4.1 shows an example of the mixture (4.3) with $q = 3$ constructed in the following way. First, for $t = 1, \dots, 100$, we generated three series

$$\begin{aligned} x_{t1} &= 2 \cos(2\pi t \cdot 6/100) + 3 \sin(2\pi t \cdot 6/100) \\ x_{t2} &= 4 \cos(2\pi t \cdot 10/100) + 5 \sin(2\pi t \cdot 10/100) \\ x_{t3} &= 6 \cos(2\pi t \cdot 40/100) + 7 \sin(2\pi t \cdot 40/100) \end{aligned}$$

² For example, for x_t in (4.2) we have $\text{cov}(x_{t+h}, x_t) = \sigma^2 \{\cos(2\pi\omega[t+h])\cos(2\pi\omega t) + \sin(2\pi\omega[t+h])\sin(2\pi\omega t)\} = \sigma^2 \cos(2\pi\omega h)$, noting that $\text{cov}(U_1, U_2) = 0$.




Fig. 4.1. Periodic components and their sum as described in [Example 4.1](#).

These three series are displayed in [Figure 4.1](#) along with the corresponding frequencies and squared amplitudes. For example, the squared amplitude of x_{t1} is $A^2 = 2^2 + 3^2 = 13$. Hence, the maximum and minimum values that x_{t1} will attain are $\pm\sqrt{13} = \pm 3.61$.

Finally, we constructed

$$x_t = x_{t1} + x_{t2} + x_{t3}$$

and this series is also displayed in [Figure 4.1](#). We note that x_t appears to behave as some of the periodic series we saw in Chapters 1 and 2. The systematic sorting out of the essential frequency components in a time series, including their relative contributions, constitutes one of the main objectives of spectral analysis. The R code to reproduce [Figure 4.1](#) is

```
x1 = 2*cos(2*pi*1:100*6/100) + 3*sin(2*pi*1:100*6/100)
x2 = 4*cos(2*pi*1:100*10/100) + 5*sin(2*pi*1:100*10/100)
x3 = 6*cos(2*pi*1:100*40/100) + 7*sin(2*pi*1:100*40/100)
x = x1 + x2 + x3
par(mfrow=c(2,2))
plot.ts(x1, ylim=c(-10,10), main=expression(omega==6/100~~~A^2==13))
plot.ts(x2, ylim=c(-10,10), main=expression(omega==10/100~~~A^2==41))
plot.ts(x3, ylim=c(-10,10), main=expression(omega==40/100~~~A^2==85))
plot.ts(x, ylim=c(-16,16), main="sum")
```

Example 4.2 The Scaled Periodogram for [Example 4.1](#)

In §2.3, [Example 2.10](#), we introduced the periodogram as a way to discover the periodic components of a time series. Recall that the scaled periodogram is given




Fig. 4.2. Periodogram of the data generated in **Example 4.1**.

by

$$P(j/n) = \left(\frac{2}{n} \sum_{t=1}^n x_t \cos(2\pi t j/n) \right)^2 + \left(\frac{2}{n} \sum_{t=1}^n x_t \sin(2\pi t j/n) \right)^2, \quad (4.6)$$

and it may be regarded as a measure of the squared correlation of the data with sinusoids oscillating at a frequency of $\omega_j = j/n$, or j cycles in n time points. Recall that we are basically computing the regression of the data on the sinusoids varying at the fundamental frequencies, j/n . As discussed in **Example 2.10**, the periodogram may be computed quickly using the fast Fourier transform (FFT), and there is no need to run repeated regressions.

The scaled periodogram of the data, x_t , simulated in **Example 4.1** is shown in **Figure 4.2**, and it clearly identifies the three components x_{t1} , x_{t2} , and x_{t3} of x_t . Note that

$$P(j/n) = P(1 - j/n), \quad j = 0, 1, \dots, n-1,$$

so there is a mirroring effect at the folding frequency of $1/2$; consequently, the periodogram is typically not plotted for frequencies higher than the folding frequency. In addition, note that the heights of the scaled periodogram shown in the figure are

$$P(6/100) = 13, \quad P(10/100) = 41, \quad P(40/100) = 85,$$

$P(j/n) = P(1 - j/n)$ and $P(j/n) = 0$ otherwise. These are exactly the values of the squared amplitudes of the components generated in **Example 4.1**. This outcome suggests that the periodogram may provide some insight into the variance components, (4.5), of a real set of data.

Assuming the simulated data, **x**, were retained from the previous example, the R code to reproduce **Figure 4.2** is

```
P = Mod(2*fft(x)/100)^2
Fr = 0:99/100
plot(Fr, P, type="o", xlab="frequency", ylab="periodogram")
```

If we consider the data x_t in [Example 4.1](#) as a color (waveform) made up of primary colors x_{t1}, x_{t2}, x_{t3} at various strengths (amplitudes), then we might consider the periodogram as a prism that decomposes the color x_t into its primary colors (spectrum). Hence the term *spectral analysis*.

Another fact that may be of use in understanding the periodogram is that for any time series sample $x_1, \dots, x_n$, where n is odd, we may write, *exactly*

$$x_t = a_0 + \sum_{j=1}^{(n-1)/2} [a_j \cos(2\pi t j/n) + b_j \sin(2\pi t j/n)], \quad (4.7)$$

for $t = 1, \dots, n$ and suitably chosen coefficients. If n is even, the representation (4.7) can be modified by summing to $(n/2 - 1)$ and adding an additional component given by $a_{n/2} \cos(2\pi t 1/2) = a_{n/2}(-1)^t$. The crucial point here is that (4.7) is exact for any sample. Hence (4.3) may be thought of as an approximation to (4.7), the idea being that many of the coefficients in (4.7) may be close to zero. Recall from [Example 2.10](#) that

$$P(j/n) = a_j^2 + b_j^2, \quad (4.8)$$

so the scaled periodogram indicates which components in (4.7) are large in magnitude and which components are small. We also saw (4.8) in [Example 4.2](#).

The following is an example using actual data.

Example 4.3 Star Magnitude

The data in [Figure 4.3](#) are the magnitude of a star taken at midnight for 600 consecutive days. The data are taken from the classic text, *The Calculus of Observations, a Treatise on Numerical Mathematics*, by E.T. Whittaker and G. Robinson, (1923, Blackie & Son, Ltd.).

The periodogram for frequencies less than .08 is also displayed in the figure; the periodogram for frequencies higher than .08 are essentially zero. Note that the 29 day cycle and the 25 day cycle are the most prominent periodic components of the data.

We can interpret this result as follows. The data are what is referred to as an amplitude modulated signal; e.g., consider $x_t = \cos(2\pi\omega t) \cos(2\pi\delta t)$, where δ is very small. In this case, the process will oscillate at frequency ω , but the amplitude will be modulated by $\cos(2\pi\delta t)$. Since $2 \cos(a) \cos(b) = \cos(a + b) + \cos(a - b)$, the periodogram of data generated as x_t will have two peaks close to each other. For example, try this on your own:

```
t = 1:200
plot.ts(2*cos(2*pi*.2*t)*cos(2*pi*.01*t)) # not shown
lines(cos(2*pi*.19*t)+cos(2*pi*.21*t))   # the same
```

The R code to reproduce [Figure 4.3](#) is




Fig. 4.3. Star magnitudes and part of the corresponding periodogram.

```

n = length(star)
par(mfrow=c(2,1), mar=c(3,3,1,1), mgp=c(1.6,.6,0))
plot(star, ylab="star magnitude", xlab="day")
Per = Mod(fft(star-mean(star)))^2/n
Freq = (1:n -1)/n
plot(Freq[1:50], Per[1:50], type='h', lwd=3, ylab="Periodogram",
     xlab="Frequency")
u = which.max(Per[1:50])           # 22  freq=21/600=.035 cycles/day
uu = which.max(Per[1:50][-u])     # 25  freq=24/600=.04 cycles/day
1/Freq[22]; 1/Freq[25]             # period = days/cycle
text(.05, 7000, "25 day cycle"); text(.027, 9000, "29 day cycle")

```

The periodogram, which was introduced in Schuster (1898) and used in Schuster (1906) for studying the periodicities in the sunspot series (shown in [Figure 4.22](#) in the Problems section) is a sample based statistic. In [Example 4.2](#) and [Example 4.3](#), we discussed the fact that the periodogram may be giving us an idea of the variance components associated with each frequency, as presented in (4.5), of a time series. These variance components, however, are population parameters. The concepts of population parameters and sample statistics, as they relate to spectral analysis of time series can be generalized to cover stationary time series and that is the topic of the next section.

4.3 The Spectral Density

The idea that a time series is composed of periodic components, appearing in proportion to their underlying variances, is fundamental in the spectral representation, given in **Theorem C.2** of Appendix C. The result is quite technical because it involves stochastic integration; that is, integration with respect to a stochastic process. The essence of the result is that (4.3) is approximately true for any stationary time series. In other words, we have the following.

Property 4.1 Spectral Representation of a Stationary Process

In nontechnical terms, any stationary time series may be thought of, approximately, as the random superposition of sines and cosines oscillating at various frequencies.

Given that (4.3) is approximately true for all stationary time series, the next question is whether a meaningful representation for its autocovariance function, like the one displayed in (4.4), also exists. The answer is yes. The following example will help explain the result.

Example 4.4 A Periodic Stationary Process

Consider a periodic stationary random process given by (4.2), with a fixed frequency ω_0 , say,

$$x_t = U_1 \cos(2\pi\omega_0 t) + U_2 \sin(2\pi\omega_0 t),$$

where U_1 and U_2 are independent zero-mean random variables with equal variance σ^2 . The number of time periods needed for the above series to complete one cycle is exactly $1/\omega_0$, and the process makes exactly ω_0 cycles per point for $t = 0, \pm 1, \pm 2, \dots$. It is easily shown that³

$$\begin{aligned} \gamma(h) &= \sigma^2 \cos(2\pi\omega_0 h) = \frac{\sigma^2}{2} e^{-2\pi i \omega_0 h} + \frac{\sigma^2}{2} e^{2\pi i \omega_0 h} \\ &= \int_{-1/2}^{1/2} e^{2\pi i \omega h} dF(\omega) \end{aligned}$$

using a Riemann–Stieltjes integration, where $F(\omega)$ is the function defined by

$$F(\omega) = \begin{cases} 0 & \omega < -\omega_0, \\ \sigma^2/2 & -\omega_0 \leq \omega < \omega_0, \\ \sigma^2 & \omega \geq \omega_0. \end{cases}$$

The function $F(\omega)$ behaves like a cumulative distribution function for a discrete random variable, except that $F(\infty) = \sigma^2 = \text{var}(x_t)$ instead of one. In fact, $F(\omega)$ is a cumulative distribution function, not of probabilities, but rather of variances associated with the frequency ω_0 in an analysis of variance, with $F(\infty)$ being the total variance of the process x_t . Hence, we term $F(\omega)$ the *spectral distribution function*.

³ Some identities may be helpful here: $e^{i\alpha} = \cos(\alpha) + i \sin(\alpha)$ and consequently, $\cos(\alpha) = (e^{i\alpha} + e^{-i\alpha})/2$ and $\sin(\alpha) = (e^{i\alpha} - e^{-i\alpha})/2i$.

A representation such as the one given in **Example 4.4** always exists for a stationary process. In particular, if x_t is stationary with autocovariance $\gamma(h) = E[(x_{t+h} - \mu)(x_t - \mu)]$, then there exists a unique monotonically increasing function $F(\omega)$, called the spectral distribution function, that is bounded, with $F(-\infty) = F(-1/2) = 0$, and $F(\infty) = F(1/2) = \gamma(0)$ such that

$$\gamma(h) = \int_{-1/2}^{1/2} e^{2\pi i \omega h} dF(\omega). \quad (4.9)$$

A more important situation we use repeatedly is the case when the autocovariance function is absolutely summable, in which case the spectral distribution function is absolutely continuous with $dF(\omega) = f(\omega) d\omega$, and the representation (4.9) becomes the motivation for the property given below.

Property 4.2 The Spectral Density

If the autocovariance function, $\gamma(h)$, of a stationary process satisfies

$$\sum_{h=-\infty}^{\infty} |\gamma(h)| < \infty, \quad (4.10)$$

then it has the representation

$$\gamma(h) = \int_{-1/2}^{1/2} e^{2\pi i \omega h} f(\omega) d\omega \quad h = 0, \pm 1, \pm 2, \dots \quad (4.11)$$

as the inverse transform of the spectral density, which has the representation

$$f(\omega) = \sum_{h=-\infty}^{\infty} \gamma(h) e^{-2\pi i \omega h} \quad -1/2 \leq \omega \leq 1/2. \quad (4.12)$$

This spectral density is the analogue of the probability density function; the fact that $\gamma(h)$ is non-negative definite ensures

$$f(\omega) \geq 0$$

for all ω . It follows immediately from (4.12) that

$$f(\omega) = f(-\omega)$$

verifying the spectral density is an even function. Because of the evenness, we will typically only plot $f(\omega)$ for $\omega \geq 0$. In addition, putting $h = 0$ in (4.11) yields

$$\gamma(0) = \text{var}(x_t) = \int_{-1/2}^{1/2} f(\omega) d\omega,$$

which expresses the total variance as the integrated spectral density over all of the frequencies. We show later on, that a linear filter can isolate the variance in certain frequency intervals or bands.

Analogous to probability theory, $\gamma(h)$ in (4.11) is the characteristic function⁴ of the spectral density $f(\omega)$ in (4.12). These facts should make it clear that, when the conditions of **Property 4.2** are satisfied, *the autocovariance function, $\gamma(h)$, and the spectral density function, $f(\omega)$, contain the same information*. That information, however, is expressed in different ways. The autocovariance function expresses information in terms of lags, whereas the spectral density expresses the same information in terms of cycles. Some problems are easier to work with when considering lagged information and we would tend to handle those problems in the time domain. Nevertheless, other problems are easier to work with when considering periodic information and we would tend to handle those problems in the spectral domain.

We note that the autocovariance function, $\gamma(h)$, in (4.11) and the spectral density, $f(\omega)$, in (4.12) are Fourier transform pairs. In particular, this means that if $f(\omega)$ and $g(\omega)$ are two spectral densities for which

$$\gamma_f(h) = \int_{-1/2}^{1/2} f(\omega) e^{2\pi i \omega h} d\omega = \int_{-1/2}^{1/2} g(\omega) e^{2\pi i \omega h} d\omega = \gamma_g(h) \quad (4.13)$$

for all $h = 0, \pm 1, \pm 2, \dots$, then

$$f(\omega) = g(\omega). \quad (4.14)$$

We also mention, at this point, that we have been focusing on the frequency ω , expressed in cycles per point rather than the more common (in statistics) alternative $\lambda = 2\pi\omega$ that would give radians per point. Finally, the absolute summability condition, (4.10), is not satisfied by (4.4), the example that we have used to introduce the idea of a spectral representation. The condition, however, is satisfied for ARMA models.

It is illuminating to examine the spectral density for the series that we have looked at in earlier discussions.

Example 4.5 White Noise Series

As a simple example, consider the theoretical power spectrum of a sequence of uncorrelated random variables, w_t , with variance σ_w^2 . A simulated set of data is displayed in the top of **Figure 1.8**. Because the autocovariance function was computed in **Example 1.16** as $\gamma_w(h) = \sigma_w^2$ for $h = 0$, and zero, otherwise, it follows from (4.12), that

$$f_w(\omega) = \sigma_w^2$$

for $-1/2 \leq \omega \leq 1/2$. Hence the process contains equal power at all frequencies. This property is seen in the realization, which seems to contain all different frequencies in a roughly equal mix. In fact, the name white noise comes from the analogy to white light, which contains all frequencies in the color spectrum at the same level of intensity. **Figure 4.4** shows a plot of the white noise spectrum for $\sigma_w^2 = 1$.

Since the linear process is an essential tool, it is worthwhile investigating the spectrum of such a process. In general, a linear filter uses a set of specified coefficients,

⁴ If $M_X(\lambda) = E(e^{\lambda X})$ for $\lambda \in \mathbb{R}$ is the moment generating function of random variable X , then $\varphi_X(\lambda) = M_X(i\lambda)$ is the characteristic function.

say a_j , for $j = 0, \pm 1, \pm 2, \dots$, to transform an input series, y_t , producing an output series, x_t , of the form

$$y_t = \sum_{j=-\infty}^{\infty} a_j x_{t-j}, \quad \sum_{j=-\infty}^{\infty} |a_j| < \infty. \quad (4.15)$$

The form (4.15) is also called a convolution in some statistical contexts. The coefficients, collectively called the *impulse response function*, and the Fourier transform

$$A(\omega) = \sum_{j=-\infty}^{\infty} a_j e^{-2\pi i \omega j}, \quad (4.16)$$

is called the *frequency response function*. If we further assume that y_t has spectral density $f_y(\omega)$ and x_t has spectral density $f_x(\omega)$, then we have the following result.

Property 4.3 Output Spectrum of a Filtered Stationary Series

Under the established conditions, the spectrum of the filtered output y_t in (4.15) is related to the spectrum of the input x_t by

$$f_y(\omega) = |A(\omega)|^2 f_x(\omega), \quad (4.17)$$

where the frequency response function $A(\omega)$ is defined in (4.16).

Proof: The autocovariance function of the filtered output y_t in (4.15) is

$$\begin{aligned} \gamma_y(h) &= \text{cov}(x_{t+h}, x_t) \\ &= \text{cov}\left(\sum_r a_r x_{t+h-r}, \sum_s a_s x_{t-s}\right) \\ &= \sum_r \sum_s a_r \gamma_x(h-r+s) a_s \\ &= \sum_r \sum_s a_r \left[\int_{-1/2}^{1/2} e^{2\pi i \omega (h-r+s)} f_x(\omega) d\omega \right] a_s \\ &= \int_{-1/2}^{1/2} \left(\sum_r a_r e^{-2\pi i \omega r} \right) \left(\sum_s a_s e^{2\pi i \omega s} \right) e^{2\pi i \omega h} f_x(\omega) d\omega \\ &= \int_{-1/2}^{1/2} e^{2\pi i \omega h} \underbrace{|A(\omega)|^2 f_x(\omega)}_{f_y(\omega)} d\omega, \end{aligned}$$

where we have first replaced $\gamma_x(\cdot)$ by its representation (4.11) and then substituted $A(\omega)$ from (4.16). The result holds by exploiting the uniqueness of the Fourier transform. $\square$

The use of **Property 4.3** is explored further in §4.8. If x_t is ARMA, its spectral density can be obtained explicitly using the fact that it is a linear process, i.e., $x_t = \sum_{j=0}^{\infty} \psi_j w_{t-j}$, where $\sum_{j=0}^{\infty} |\psi_j| < \infty$. The following property is a direct consequence of the more general result, **Property 4.3**, by using the additional facts that the spectral density of white noise is $f_w(\omega) = \sigma_w^2$, and by **Property 3.1**, $\psi(z) = \theta(z)/\phi(z)$.

Property 4.4 The Spectral Density of ARMA

If x_t is ARMA(p, q), $\phi(B)x_t = \theta(B)w_t$, its spectral density is given by

$$f_x(\omega) = \sigma_w^2 \frac{|\theta(e^{-2\pi i\omega})|^2}{|\phi(e^{-2\pi i\omega})|^2} \quad (4.18)$$

where $\phi(z) = 1 - \sum_{k=1}^p \phi_k z^k$ and $\theta(z) = 1 + \sum_{k=1}^q \theta_k z^k$.

Example 4.6 Moving Average

As an example of a series that does not have an equal mix of frequencies, we consider a moving average model. Specifically, consider the MA(1) model given by

$$x_t = w_t + .5w_{t-1}.$$

A sample realization is shown in the top of [Figure 3.2](#) and we note that the series has less of the higher or faster frequencies. The spectral density will verify this observation.

The autocovariance function is displayed in [Example 3.4](#), and for this particular example, we have

$$\gamma(0) = (1 + .5^2)\sigma_w^2 = 1.25\sigma_w^2; \quad \gamma(\pm 1) = .5\sigma_w^2; \quad \gamma(\pm h) = 0 \text{ for } h > 1.$$

Substituting this directly into the definition given in (4.12), we have

$$\begin{aligned} f(\omega) &= \sum_{h=-\infty}^{\infty} \gamma(h) e^{-2\pi i\omega h} = \sigma_w^2 \left[1.25 + .5(e^{-2\pi i\omega} + e^{2\pi i\omega}) \right] \\ &= \sigma_w^2 [1.25 + \cos(2\pi\omega)]. \end{aligned} \quad (4.19)$$

We can also compute the spectral density using [Property 4.4](#), which states that for an MA, $f(\omega) = \sigma_w^2 |\theta(e^{-2\pi i\omega})|^2$. Because $\theta(z) = 1 + .5z$, we have

$$\begin{aligned} |\theta(e^{-2\pi i\omega})|^2 &= |1 + .5e^{-2\pi i\omega}|^2 = (1 + .5e^{-2\pi i\omega})(1 + .5e^{2\pi i\omega}) \\ &= 1.25 + .5(e^{-2\pi i\omega} + e^{2\pi i\omega}) \end{aligned}$$

which leads to agreement with (4.19).

Plotting the spectrum for $\sigma_w^2 = 1$, as in the middle of [Figure 4.4](#), shows the lower or slower frequencies have greater power than the higher or faster frequencies.

Example 4.7 A Second-Order Autoregressive Series

We now consider the spectrum of an AR(2) series of the form

$$x_t - \phi_1 x_{t-1} - \phi_2 x_{t-2} = w_t,$$

for the special case $\phi_1 = 1$ and $\phi_2 = -.9$. [Figure 1.9](#) on page 13 shows a sample realization of such a process for $\sigma_w = 1$. We note the data exhibit a strong periodic component that makes a cycle about every six points.




Fig. 4.4. Theoretical spectra of white noise (top), a first-order moving average (middle), and a second-order autoregressive process (bottom).

To use [Property 4.4](#), note that $\theta(z) = 1$, $\phi(z) = 1 - z + .9z^2$ and

$$\begin{aligned}
 |\phi(e^{-2\pi i\omega})|^2 &= (1 - e^{-2\pi i\omega} + .9e^{-4\pi i\omega})(1 - e^{2\pi i\omega} + .9e^{4\pi i\omega}) \\
 &= 2.81 - 1.9(e^{2\pi i\omega} + e^{-2\pi i\omega}) + .9(e^{4\pi i\omega} + e^{-4\pi i\omega}) \\
 &= 2.81 - 3.8 \cos(2\pi\omega) + 1.8 \cos(4\pi\omega).
 \end{aligned}$$

Using this result in [\(4.18\)](#), we have that the spectral density of x_t is

$$f_x(\omega) = \frac{\sigma_w^2}{2.81 - 3.8 \cos(2\pi\omega) + 1.8 \cos(4\pi\omega)}.$$

Setting $\sigma_w = 1$, the bottom of [Figure 4.4](#) displays $f_x(\omega)$ and shows a strong power component at about $\omega = .16$ cycles per point or a period between six and seven cycles per point and very little power at other frequencies. In this case, modifying the white noise series by applying the second-order AR operator has concentrated the power or variance of the resulting series in a very narrow frequency band.

The spectral density can also be obtained from first principles, without having to use [Property 4.4](#). Because $w_t = x_t - x_{t-1} + .9x_{t-2}$ in this example, we have

$$\begin{aligned}
\gamma_w(h) &= \text{cov}(w_{t+h}, w_t) \\
&= \text{cov}(x_{t+h} - x_{t+h-1} + .9x_{t+h-2}, x_t - x_{t-1} + .9x_{t-2}) \\
&= 2.81\gamma_x(h) - 1.9[\gamma_x(h+1) + \gamma_x(h-1)] + .9[\gamma_x(h+2) + \gamma_x(h-2)]
\end{aligned}$$

Now, substituting the spectral representation (4.11) for $\gamma_x(h)$ in the above equation yields

$$\begin{aligned}
\gamma_w(h) &= \int_{-1/2}^{1/2} [2.81 - 1.9(e^{2\pi i\omega} + e^{-2\pi i\omega}) + .9(e^{4\pi i\omega} + e^{-4\pi i\omega})] e^{2\pi i\omega h} f_x(\omega) d\omega \\
&= \int_{-1/2}^{1/2} [2.81 - 3.8 \cos(2\pi\omega) + 1.8 \cos(4\pi\omega)] e^{2\pi i\omega h} f_x(\omega) d\omega.
\end{aligned}$$

If the spectrum of the white noise process, w_t , is $g_w(\omega)$, the uniqueness of the Fourier transform allows us to identify

$$g_w(\omega) = [2.81 - 3.8 \cos(2\pi\omega) + 1.8 \cos(4\pi\omega)] f_x(\omega).$$

But, as we have already seen, $g_w(\omega) = \sigma_w^2$, from which we deduce that

$$f_x(\omega) = \frac{\sigma_w^2}{2.81 - 3.8 \cos(2\pi\omega) + 1.8 \cos(4\pi\omega)}$$

is the spectrum of the autoregressive series.

To reproduce [Figure 4.4](#), use the `spec.arma`:

```

par(mfrow=c(3,1))
spec.arma(log="no", main="White Noise")
spec.arma(ma=.5, log="no", main="Moving Average")
spec.arma(ar=c(1,-.9), log="no", main="Autoregression")

```

The above examples motivate the use of the power spectrum for describing the theoretical variance fluctuations of a stationary time series. Indeed, the interpretation of the spectral density function as the variance of the time series over a given frequency band gives us the intuitive explanation for its physical meaning. The plot of the function $f(\omega)$ over the frequency argument ω can even be thought of as an analysis of variance, in which the columns or block effects are the frequencies, indexed by ω .

Example 4.8 Every Explosion has a Cause (cont)

In [Example 3.3](#), we discussed the fact that explosive models have causal counterparts. In that example, we also indicated that it was easier to show this result in general in the spectral domain. In this example, we give the details for an AR(1) model, but the techniques used here will indicate how to generalize the result.

As in [Example 3.3](#), we suppose that $x_t = 2x_{t-1} + w_t$, where $w_t \sim \text{iid } N(0, \sigma_w^2)$. Then, the spectral density of x_t is

$$f_x(\omega) = \sigma_w^2 |1 - 2e^{-2\pi i\omega}|^{-2}. \quad (4.20)$$

But,

$$|1 - 2e^{-2\pi i\omega}| = |1 - 2e^{2\pi i\omega}| = |(2e^{2\pi i\omega})(\tfrac{1}{2}e^{-2\pi i\omega} - 1)| = 2|1 - \tfrac{1}{2}e^{-2\pi i\omega}|.$$

Thus, (4.20) can be written as

$$f_x(\omega) = \tfrac{1}{4}\sigma_w^2 |1 - \tfrac{1}{2}e^{-2\pi i\omega}|^{-2},$$

which implies that $x_t = \tfrac{1}{2}x_{t-1} + v_t$, with $v_t \sim \text{iid } N(0, \tfrac{1}{4}\sigma_w^2)$ is an equivalent form of the model.

4.4 Periodogram and Discrete Fourier Transform

We are now ready to tie together the periodogram, which is the sample-based concept presented in §4.2, with the spectral density, which is the population-based concept of §4.3.

Definition 4.1 Given data $x_1, \dots, x_n$, we define the **discrete Fourier transform (DFT)** to be

$$d(\omega_j) = n^{-1/2} \sum_{t=1}^n x_t e^{-2\pi i\omega_j t} \quad (4.21)$$

for $j = 0, 1, \dots, n-1$, where the frequencies $\omega_j = j/n$ are called the **Fourier or fundamental frequencies**.

If n is a highly composite integer (i.e., it has many factors), the DFT can be computed by the fast Fourier transform (FFT) introduced in Cooley and Tukey (1965). Also, different packages scale the FFT differently, so it is a good idea to consult the documentation. R computes the DFT defined in (4.21) without the factor $n^{-1/2}$, but with an additional factor of $e^{2\pi i\omega_j}$ that can be ignored because we will be interested in the squared modulus of the DFT. Sometimes it is helpful to exploit the inversion result for DFTs which shows the linear transformation is one-to-one. For the inverse DFT we have,

$$x_t = n^{-1/2} \sum_{j=0}^{n-1} d(\omega_j) e^{2\pi i\omega_j t} \quad (4.22)$$

for $t = 1, \dots, n$. The following example shows how to calculate the DFT and its inverse in R for the data set $\{1, 2, 3, 4\}$; note that R writes a complex number $z = a + ib$ as `a+bi`.

```
(dft = fft(1:4)/sqrt(4))
[1] 5+0i -1+1i -1+0i -1-1i
(idft = fft(dft, inverse=TRUE)/sqrt(4))
[1] 1+0i 2+0i 3+0i 4+0i
(Re(idft)) # keep it real
[1] 1 2 3 4
```

We now define the periodogram as the squared modulus⁵ of the DFT.

⁵ Recall that if $z = a + ib$, then $\bar{z} = a - ib$, and $|z|^2 = z\bar{z} = a^2 + b^2$.

Definition 4.2 Given data $x_1, \dots, x_n$, we define the **periodogram** to be

$$I(\omega_j) = |d(\omega_j)|^2 \quad (4.23)$$

for $j = 0, 1, 2, \dots, n-1$.

Note that $I(0) = n\bar{x}^2$, where $\bar{x}$ is the sample mean. In addition, because $\sum_{t=1}^n \exp(-2\pi i t \frac{j}{n}) = 0$ for $j \neq 0$,⁶ we can write the DFT as

$$d(\omega_j) = n^{-1/2} \sum_{t=1}^n (x_t - \bar{x}) e^{-2\pi i \omega_j t} \quad (4.24)$$

for $j \neq 0$. Thus, for $j \neq 0$,

$$\begin{aligned} I(\omega_j) &= |d(\omega_j)|^2 = n^{-1} \sum_{t=1}^n \sum_{s=1}^n (x_t - \bar{x})(x_s - \bar{x}) e^{-2\pi i \omega_j (t-s)} \\ &= n^{-1} \sum_{h=-(n-1)}^{n-1} \sum_{t=1}^{n-|h|} (x_{t+|h|} - \bar{x})(x_t - \bar{x}) e^{-2\pi i \omega_j h} \\ &= \sum_{h=-(n-1)}^{n-1} \hat{\gamma}(h) e^{-2\pi i \omega_j h} \end{aligned} \quad (4.25)$$

where we have put $h = t - s$, with $\hat{\gamma}(h)$ as given in (1.34).⁷

Recall, $P(\omega_j) = (4/n)I(\omega_j)$ where $P(\omega_j)$ is the scaled periodogram defined in (4.6). Henceforth we will work with $I(\omega_j)$ instead of $P(\omega_j)$. In view of (4.25), the periodogram, $I(\omega_j)$, is the sample version of $f(\omega_j)$ given in (4.12). That is, we may think of the periodogram as the “sample spectral density” of x_t .

It is sometimes useful to work with the real and imaginary parts of the DFT individually. To this end, we define the following transforms.

Definition 4.3 Given data $x_1, \dots, x_n$, we define the **cosine transform**

$$d_c(\omega_j) = n^{-1/2} \sum_{t=1}^n x_t \cos(2\pi \omega_j t) \quad (4.26)$$

and the **sine transform**

$$d_s(\omega_j) = n^{-1/2} \sum_{t=1}^n x_t \sin(2\pi \omega_j t) \quad (4.27)$$

where $\omega_j = j/n$ for $j = 0, 1, \dots, n-1$.

We note that $d(\omega_j) = d_c(\omega_j) - i d_s(\omega_j)$ and hence

$$I(\omega_j) = d_c^2(\omega_j) + d_s^2(\omega_j). \quad (4.28)$$

We have also discussed the fact that spectral analysis can be thought of as an analysis of variance. The next example examines this notion.

⁶ $\sum_{t=1}^n z^t = z \frac{1-z^n}{1-z}$ for $z \neq 1$.

⁷ Note that (4.25) can be used to obtain $\hat{\gamma}(h)$ by taking the inverse DFT of $I(\omega_j)$. This approach was used in Example 1.29 to obtain a two-dimensional ACF.

Example 4.9 Spectral ANOVA

Let $x_1, \dots, x_n$ be a sample of size n , where for ease, n is odd. Then, recalling [Example 2.10](#) and the discussion around (4.7) and (4.8),

$$x_t = a_0 + \sum_{j=1}^m \left[a_j \cos(2\pi\omega_j t) + b_j \sin(2\pi\omega_j t) \right], \quad (4.29)$$

where $m = (n-1)/2$, is exact for $t = 1, \dots, n$. In particular, using multiple regression formulas, we have $a_0 = \bar{x}$,

$$a_j = \frac{2}{n} \sum_{t=1}^n x_t \cos(2\pi\omega_j t) = \frac{2}{\sqrt{n}} d_c(\omega_j)$$

$$b_j = \frac{2}{n} \sum_{t=1}^n x_t \sin(2\pi\omega_j t) = \frac{2}{\sqrt{n}} d_s(\omega_j).$$

Hence, we may write

$$(x_t - \bar{x}) = \frac{2}{\sqrt{n}} \sum_{j=1}^m \left[d_c(\omega_j) \cos(2\pi\omega_j t) + d_s(\omega_j) \sin(2\pi\omega_j t) \right]$$

for $t = 1, \dots, n$. Squaring both sides and summing we obtain

$$\sum_{t=1}^n (x_t - \bar{x})^2 = 2 \sum_{j=1}^m \left[d_c^2(\omega_j) + d_s^2(\omega_j) \right] = 2 \sum_{j=1}^m I(\omega_j)$$

using the results of [Problem 2.10\(d\)](#). Thus, we have partitioned the sum of squares into harmonic components represented by frequency ω_j with the periodogram, $I(\omega_j)$, being the mean square regression. This leads to the ANOVA table for n odd:

Source	df	SS	MS
ω_1	2	$2I(\omega_1)$	$I(\omega_1)$
ω_2	2	$2I(\omega_2)$	$I(\omega_2)$
$\vdots$	$\vdots$	$\vdots$	$\vdots$
ω_m	2	$2I(\omega_m)$	$I(\omega_m)$
Total	$n - 1$	$\sum_{t=1}^n (x_t - \bar{x})^2$	

This decomposition means that if the data contain some strong periodic components, the periodogram values corresponding to those frequencies (or near those frequencies) will be large. On the other hand, the corresponding values of the periodogram will be small for periodic components not present in the data.

The following is an R example to help explain this concept. We consider $n = 5$ observations given by $x_1 = 1, x_2 = 2, x_3 = 3, x_4 = 2, x_5 = 1$. Note that the data complete one cycle, but not in a sinusoidal way. Thus, we should expect the $\omega_1 = 1/5$ component to be relatively large but not exhaustive, and the $\omega_2 = 2/5$ component to be small.

```

x = c(1, 2, 3, 2, 1)
c1 = cos(2*pi*1:5*1/5); s1 = sin(2*pi*1:5*1/5)
c2 = cos(2*pi*1:5*2/5); s2 = sin(2*pi*1:5*2/5)
omega1 = cbind(c1, s1); omega2 = cbind(c2, s2)
anova(lm(x~omega1+omega2)) # ANOVA Table
      Df Sum Sq Mean Sq
omega1  2  2.74164  1.37082
omega2  2  .05836  .02918
Residuals 0  .00000
Mod(fft(x))^2/5 # the periodogram (as a check)
[1] 16.2  1.37082 .029179 .029179  1.37082
# I(0) I(1/5) I(2/5) I(3/5) I(4/5)

```

Note that $\bar{x} = 1.8$, and $I(0) = 16.2 = 5 \times 1.8^2 (= n\bar{x}^2)$. Also, note that

$$I(1/5) = 1.37082 = \text{Mean Sq}(\omega_1) \quad \text{and} \quad I(2/5) = .02918 = \text{Mean Sq}(\omega_2)$$

and $I(j/5) = I(1 - j/5)$, for $j = 3, 4$. Finally, we note that the sum of squares associated with the residuals (SSE) is zero, indicating an exact fit.

We are now ready to present some large sample properties of the periodogram. First, let μ be the mean of a stationary process x_t with absolutely summable autocovariance function $\gamma(h)$ and spectral density $f(\omega)$. We can use the same argument as in (4.25), replacing $\bar{x}$ by μ in (4.24), to write

$$I(\omega_j) = n^{-1} \sum_{h=-(n-1)}^{n-1} \sum_{t=1}^{n-|h|} (x_{t+|h|} - \mu)(x_t - \mu) e^{-2\pi i \omega_j h} \quad (4.30)$$

where ω_j is a non-zero fundamental frequency. Taking expectation in (4.30) we obtain

$$E[I(\omega_j)] = \sum_{h=-(n-1)}^{n-1} \left(\frac{n-|h|}{n} \right) \gamma(h) e^{-2\pi i \omega_j h}. \quad (4.31)$$

For any given $\omega \neq 0$, choose a sequence of fundamental frequencies $\omega_{j:n} \rightarrow \omega$ ⁸ from which it follows by (4.31) that, as $n \rightarrow \infty$ ⁹

$$E[I(\omega_{j:n})] \rightarrow f(\omega) = \sum_{h=-\infty}^{\infty} \gamma(h) e^{-2\pi i h \omega}. \quad (4.32)$$

In other words, under absolute summability of $\gamma(h)$, the spectral density is the long-term average of the periodogram.

To examine the asymptotic distribution of the periodogram, we note that if x_t is a normal time series, the sine and cosine transforms will also be jointly normal, because they are linear combinations of the jointly normal random variables $x_1, x_2, \dots, x_n$. In that case, the assumption that the covariance function satisfies the condition

⁸ By this we mean $\omega_{j:n} = j_n/n$, where $\{j_n\}$ is a sequence of integers chosen so that j_n/n is the closest Fourier frequency to ω ; consequently, $|j_n/n - \omega| \leq \frac{1}{2n}$.

⁹ From Definition 4.2 we have $I(0) = n\bar{x}^2$, so the analogous result of (4.32) for the case $\omega = 0$ is $E[I(0)] - n\mu^2 = n \text{var}(\bar{x}) \rightarrow f(0)$ as $n \rightarrow \infty$.

$$\theta = \sum_{h=-\infty}^{\infty} |h| |\gamma(h)| < \infty \quad (4.33)$$

is enough to obtain simple large sample approximations for the variances and covariances. Using the same argument used to develop (4.31) we have

$$\text{cov}[d_c(\omega_j), d_c(\omega_k)] = n^{-1} \sum_{s=1}^n \sum_{t=1}^n \gamma(s-t) \cos(2\pi\omega_j s) \cos(2\pi\omega_k t), \quad (4.34)$$

$$\text{cov}[d_c(\omega_j), d_s(\omega_k)] = n^{-1} \sum_{s=1}^n \sum_{t=1}^n \gamma(s-t) \cos(2\pi\omega_j s) \sin(2\pi\omega_k t), \quad (4.35)$$

and

$$\text{cov}[d_s(\omega_j), d_s(\omega_k)] = n^{-1} \sum_{s=1}^n \sum_{t=1}^n \gamma(s-t) \sin(2\pi\omega_j s) \sin(2\pi\omega_k t), \quad (4.36)$$

where the variance terms are obtained by setting $\omega_j = \omega_k$ in (4.34) and (4.36). In [Appendix C, §C.2](#), we show the terms in (4.34)-(4.36) have interesting properties under assumption (4.33), namely, for $\omega_j, \omega_k \neq 0$ or $1/2$,

$$\text{cov}[d_c(\omega_j), d_c(\omega_k)] = \begin{cases} f(\omega_j)/2 + \varepsilon_n & \omega_j = \omega_k, \\ \varepsilon_n & \omega_j \neq \omega_k, \end{cases} \quad (4.37)$$

$$\text{cov}[d_s(\omega_j), d_s(\omega_k)] = \begin{cases} f(\omega_j)/2 + \varepsilon_n & \omega_j = \omega_k, \\ \varepsilon_n & \omega_j \neq \omega_k, \end{cases} \quad (4.38)$$

and

$$\text{cov}[d_c(\omega_j), d_s(\omega_k)] = \varepsilon_n, \quad (4.39)$$

where the error term ε_n in the approximations can be bounded,

$$|\varepsilon_n| \leq \theta/n, \quad (4.40)$$

and θ is given by (4.33). If $\omega_j = \omega_k = 0$ or $1/2$ in (4.37), the multiplier $1/2$ disappears; note that $d_s(0) = d_s(1/2) = 0$, so (4.38) does not apply.

Example 4.10 Covariance of Sine and Cosine Transforms

For the three-point moving average series of [Example 1.9](#) and $n = 256$ observations, the theoretical covariance matrix of the vector $\mathbf{d} = (d_c(\omega_{26}), d_s(\omega_{26}), d_c(\omega_{27}), d_s(\omega_{27}))'$ is

$$\text{cov}(\mathbf{d}) = \begin{pmatrix} .3752 & -.0009 & -.0022 & -.0010 \\ -.0009 & .3777 & -.0009 & .0003 \\ -.0022 & -.0009 & .3667 & -.0010 \\ -.0010 & .0003 & -.0010 & .3692 \end{pmatrix}.$$

The diagonal elements can be compared with half the theoretical spectral values of $\frac{1}{2}f(\omega_{26}) = .3774$ for the spectrum at frequency $\omega_{26} = 26/256$, and of $\frac{1}{2}f(\omega_{27}) = .3689$ for the spectrum at $\omega_{27} = 27/256$. Hence, the cosine and sine transforms produce nearly uncorrelated variables with variances approximately equal to one half of the theoretical spectrum. For this particular case, the uniform bound is determined from $\theta = 8/9$, yielding $|\varepsilon_{256}| \leq .0035$ for the bound on the approximation error.

If $x_t \sim \text{iid}(0, \sigma^2)$, then it follows from (4.33)-(4.39), Problem 2.10(d), and a central limit theorem¹⁰ that

$$d_c(\omega_{j:n}) \sim \text{AN}(0, \sigma^2/2) \quad \text{and} \quad d_s(\omega_{j:n}) \sim \text{AN}(0, \sigma^2/2) \quad (4.41)$$

jointly and independently, and independent of $d_c(\omega_{k:n})$ and $d_s(\omega_{k:n})$ provided $\omega_{j:n} \rightarrow \omega_1$ and $\omega_{k:n} \rightarrow \omega_2$ where $0 < \omega_1 \neq \omega_2 < 1/2$. We note that in this case, $f_x(\omega) = \sigma^2$. In view of (4.41), it follows immediately that as $n \rightarrow \infty$,

$$\frac{2I(\omega_{j:n})}{\sigma^2} \xrightarrow{d} \chi_2^2 \quad \text{and} \quad \frac{2I(\omega_{k:n})}{\sigma^2} \xrightarrow{d} \chi_2^2 \quad (4.42)$$

with $I(\omega_{j:n})$ and $I(\omega_{k:n})$ being asymptotically independent, where χ_ν^2 denotes a chi-squared random variable with ν degrees of freedom.

Using the central limit theory of §C.2, it is fairly easy to extend the results of the iid case to the case of a linear process.

Property 4.5 Distribution of the Periodogram Ordinates

If

$$x_t = \sum_{j=-\infty}^{\infty} \psi_j w_{t-j}, \quad \sum_{j=-\infty}^{\infty} |\psi_j| < \infty \quad (4.43)$$

where $w_t \sim \text{iid}(0, \sigma_w^2)$, and (4.33) holds, then for any collection of m distinct frequencies $\omega_j \in (0, 1/2)$ with $\omega_{j:n} \rightarrow \omega_j$

$$\frac{2I(\omega_{j:n})}{f(\omega_j)} \xrightarrow{d} \text{iid } \chi_2^2 \quad (4.44)$$

provided $f(\omega_j) > 0$, for $j = 1, \dots, m$.

This result is stated more precisely in Theorem C.7. Other approaches to large sample normality of the periodogram ordinates are in terms of cumulants, as in Brillinger (1981), or in terms of mixing conditions, such as in Rosenblatt (1956a). Here, we adopt the approach used by Hannan (1970), Fuller (1996), and Brockwell and Davis (1991).

¹⁰ If $Y_j \sim \text{iid}(0, \sigma^2)$ and $\{a_j\}$ are constants for which $\sum_{j=1}^n a_j^2 / \max_{1 \leq j \leq n} a_j^2 \rightarrow \infty$ as $n \rightarrow \infty$, then $\sum_{j=1}^n a_j Y_j \sim \text{AN}(0, \sigma^2 \sum_{j=1}^n a_j^2)$. AN is read *asymptotically normal*; see Definition A.5.

The distributional result (4.44) can be used to derive an approximate confidence interval for the spectrum in the usual way. Let $\chi^2_\nu(\alpha)$ denote the lower α probability tail for the chi-squared distribution with ν degrees of freedom; that is,

$$\Pr\{\chi^2_\nu \leq \chi^2_\nu(\alpha)\} = \alpha. \quad (4.45)$$

Then, an approximate $100(1 - \alpha)\%$ confidence interval for the spectral density function would be of the form

$$\frac{2 I(\omega_{j:n})}{\chi^2_2(1 - \alpha/2)} \leq f(\omega) \leq \frac{2 I(\omega_{j:n})}{\chi^2_2(\alpha/2)}. \quad (4.46)$$

Often, nonstationary trends are present that should be eliminated before computing the periodogram. Trends introduce extremely low frequency components in the periodogram that tend to obscure the appearance at higher frequencies. For this reason, it is usually conventional to center the data prior to a spectral analysis using either mean-adjusted data of the form $x_t - \bar{x}$ to eliminate the zero or d-c component or to use detrended data of the form $x_t - \hat{\beta}_1 - \hat{\beta}_2 t$ to eliminate the term that will be considered a half cycle by the spectral analysis. Note that higher order polynomial regressions in t or nonparametric smoothing (linear filtering) could be used in cases where the trend is nonlinear.

As previously indicated, it is often convenient to calculate the DFTs, and hence the periodogram, using the fast Fourier transform algorithm. The FFT utilizes a number of redundancies in the calculation of the DFT when n is highly composite; that is, an integer with many factors of 2, 3, or 5, the best case being when $n = 2^p$ is a factor of 2. Details may be found in Cooley and Tukey (1965). To accommodate this property, we can pad the centered (or detrended) data of length n to the next highly composite integer n' by adding zeros, i.e., setting $x^c_{n+1} = x^c_{n+2} = \dots = x^c_{n'} = 0$, where x^c_t denotes the centered data. This means that the fundamental frequency ordinates will be $\omega_j = j/n'$ instead of j/n . We illustrate by considering the periodogram of the SOI and Recruitment series, as has been given in Figure 1.5 of Chapter 1. Recall that they are monthly series and $n = 453$ months. To find n' in R, use the command `nextn(453)` to see that $n' = 480$ will be used in the spectral analyses by default.

Example 4.11 Periodogram of SOI and Recruitment Series

Figure 4.5 shows the periodograms of each series, where the frequency axis is labeled in multiples of $\Delta = 1/12$. As previously indicated, the centered data have been padded to a series of length 480. We notice a narrow-band peak at the obvious yearly (12 month) cycle, $\omega = 1\Delta = 1/12$. In addition, there is considerable power in a wide band at the lower frequencies that is centered around the four-year (48 month) cycle $\omega = \frac{1}{4}\Delta = 1/48$ representing a possible El Niño effect. This wide band activity suggests that the possible El Niño cycle is irregular, but tends to be around four years on average. We will continue to address this problem as we move to more sophisticated analyses.

Noting $\chi^2_2(.025) = .05$ and $\chi^2_2(.975) = 7.38$, we can obtain approximate 95% confidence intervals for the frequencies of interest. For example, the periodogram




Fig. 4.5. Periodogram of SOI and Recruitment, $n = 453$ ($n' = 480$), where the frequency axis is labeled in multiples of $\Delta = 1/12$. Note the common peaks at $\omega = 1\Delta = 1/12$, or one cycle per year (12 months), and $\omega = \frac{1}{4}\Delta = 1/48$, or one cycle every four years (48 months).

of the SOI series is $I_S(1/12) = .97$ at the yearly cycle. An approximate 95% confidence interval for the spectrum $f_S(1/12)$ is then

$$[2(.97)/7.38, 2(.97)/.05] = [.26, 38.4],$$

which is too wide to be of much use. We do notice, however, that the lower value of .26 is higher than any other periodogram ordinate, so it is safe to say that this value is significant. On the other hand, an approximate 95% confidence interval for the spectrum at the four-year cycle, $f_S(1/48)$, is

$$[2(.05)/7.38, 2(.05)/.05] = [.01, 2.12],$$

which again is extremely wide, and with which we are unable to establish significance of the peak.

We now give the R commands that can be used to reproduce **Figure 4.5**. To calculate and graph the periodogram, we used the `mvspec` command in available from `astsa`. We note that the value of Δ is the reciprocal of the value of `frequency`

for the data of a time series object. If the data are not a time series object, `frequency` is set to 1. Also, we set `log="no"` because the periodogram is plotted on a $\log_{10}$ scale by default. Figure 4.5 displays a `bandwidth`. We will discuss bandwidth in the next section, so ignore this concepts for the time being.

```
par(mfrow=c(2,1))
soi.per = mvspec(soi, log="no")
abline(v=1/4, lty="dotted")
rec.per = mvspec(rec, log="no")
abline(v=1/4, lty="dotted")
```

The confidence intervals for the SOI series at the yearly cycle, $\omega = 1/12 = 40/480$, and the possible El Niño cycle of four years $\omega = 1/48 = 10/480$ can be computed in R as follows:

```
soi.per$spec[40] # 0.97223; soi pgram at freq 1/12 = 40/480
soi.per$spec[10] # 0.05372; soi pgram at freq 1/48 = 10/480
# conf intervals - returned value:
U = qchisq(.025,2) # 0.05063
L = qchisq(.975,2) # 7.37775
2*soi.per$spec[10]/L # 0.01456
2*soi.per$spec[10]/U # 2.12220
2*soi.per$spec[40]/L # 0.26355
2*soi.per$spec[40]/U # 38.40108
```

The example above makes it clear that the periodogram as an estimator is susceptible to large uncertainties, and we need to find a way to reduce the variance. Not surprisingly, this result follows if we think about the periodogram, $I(\omega_j)$ as an estimator of the spectral density $f(\omega)$ and realize that it is the sum of squares of only two random variables for any sample size. The solution to this dilemma is suggested by the analogy with classical statistics where we look for independent random variables with the same variance and average the squares of these common variance observations. Independence and equality of variance do not hold in the time series case, but the covariance structure of the two adjacent estimators given in Example 4.10 suggests that for neighboring frequencies, these assumptions are approximately true.

4.5 Nonparametric Spectral Estimation

To continue the discussion that ended the previous section, we introduce a frequency band, $\mathcal{B}$, of $L \ll n$ contiguous fundamental frequencies, centered around frequency $\omega_j = j/n$, which is chosen close to a frequency of interest, ω . For frequencies of the form $\omega^* = \omega_j + k/n$, let

$$\mathcal{B} = \left\{ \omega^*: \omega_j - \frac{m}{n} \leq \omega^* \leq \omega_j + \frac{m}{n} \right\}, \quad (4.47)$$

where

$$L = 2m + 1 \quad (4.48)$$

is an odd number, chosen such that the spectral values in the interval $\mathcal{B}$,

$$f(\omega_j + k/n), \quad k = -m, \dots, 0, \dots, m$$

are approximately equal to $f(\omega)$.

We now define an averaged (or smoothed) periodogram as the average of the periodogram values, say,

$$\bar{f}(\omega) = \frac{1}{L} \sum_{k=-m}^m I(\omega_j + k/n), \quad (4.49)$$

over the band $\mathcal{B}$. Under the assumption that the spectral density is fairly constant in the band $\mathcal{B}$, and in view of (4.44) we can show that under appropriate conditions,¹¹ for large n , the periodograms in (4.49) are approximately distributed as independent $f(\omega)\chi_2^2/2$ random variables, for $0 < \omega < 1/2$, as long as we keep L fairly small relative to n . This result is discussed formally in §C.2. Thus, under these conditions, $L\bar{f}(\omega)$ is the sum of L approximately independent $f(\omega)\chi_2^2/2$ random variables. It follows that, for large n ,

$$\frac{2L\bar{f}(\omega)}{f(\omega)} \sim \chi_{2L}^2 \quad (4.50)$$

where $\sim$ means *is approximately distributed as*.

In this scenario, where we smooth the periodogram by simple averaging, it seems reasonable to call the width of the frequency interval defined by (4.47),

$$B = \frac{L}{n}, \quad (4.51)$$

the bandwidth. The concept of the bandwidth, however, becomes more complicated with the introduction of spectral estimators that smooth with unequal weights. Note (4.51) implies the degrees of freedom can be expressed as

$$2L = 2Bn, \quad (4.52)$$

or twice the time-bandwidth product. The result (4.50) can be rearranged to obtain an approximate $100(1 - \alpha)\%$ confidence interval of the form

$$\frac{2L\bar{f}(\omega)}{\chi_{2L}^2(1 - \alpha/2)} \leq f(\omega) \leq \frac{2L\bar{f}(\omega)}{\chi_{2L}^2(\alpha/2)} \quad (4.53)$$

for the true spectrum, $f(\omega)$.

Many times, the visual impact of a spectral density plot will be improved by plotting the logarithm of the spectrum instead of the spectrum (the log transformation is the variance stabilizing transformation in this situation). This phenomenon can occur when regions of the spectrum exist with peaks of interest much smaller than some of the main power components. For the log spectrum, we obtain an interval of the form

¹¹ The conditions, which are sufficient, are that x_t is a linear process, as described in Property 4.5, with $\sum_j \sqrt{|j|} |\psi_j| < \infty$, and w_t has a finite fourth moment.

$$[\log \bar{f}(\omega) + \log 2L - \log \chi_{2L}^2(1 - \alpha/2), \log \bar{f}(\omega) + \log 2L - \log \chi_{2L}^2(\alpha/2)]. \quad (4.54)$$

We can also test hypotheses relating to the equality of spectra using the fact that the distributional result (4.50) implies that the ratio of spectra based on roughly independent samples will have an approximate $F_{2L, 2L}$ distribution. The independent estimators can either be from different frequency bands or from different series.

If zeros are appended before computing the spectral estimators, we need to adjust the degrees of freedom and an approximation is to replace $2L$ by $2Ln/n'$. Hence, we define the adjusted degrees of freedom as

$$df = \frac{2Ln}{n'} \quad (4.55)$$

and use it instead of $2L$ in the confidence intervals (4.53) and (4.54). For example, (4.53) becomes

$$\frac{df \bar{f}(\omega)}{\chi_{df}^2(1 - \alpha/2)} \leq f(\omega) \leq \frac{df \bar{f}(\omega)}{\chi_{df}^2(\alpha/2)}. \quad (4.56)$$

A number of assumptions are made in computing the approximate confidence intervals given above, which may not hold in practice. In such cases, it may be reasonable to employ resampling techniques such as one of the parametric bootstraps proposed by Hurvich and Zeger (1987) or a nonparametric local bootstrap proposed by Paparoditis and Politis (1999). To develop the bootstrap distributions, we assume that the contiguous DFTs in a frequency band of the form (4.47) all came from a time series with identical spectrum $f(\omega)$. This, in fact, is exactly the same assumption made in deriving the large-sample theory. We may then simply resample the L DFTs in the band, with replacement, calculating a spectral estimate from each bootstrap sample. The sampling distribution of the bootstrap estimators approximates the distribution of the nonparametric spectral estimator. For further details, including the theoretical properties of such estimators, see Paparoditis and Politis (1999).

Before proceeding further, we pause to consider computing the average periodograms for the SOI and Recruitment series, as shown in Figure 4.6.

Example 4.12 Averaged Periodogram for SOI and Recruitment

Generally, it is a good idea to try several bandwidths that seem to be compatible with the general overall shape of the spectrum, as suggested by the periodogram. The SOI and Recruitment series periodograms, previously computed in Figure 4.5, suggest the power in the lower El Niño frequency needs smoothing to identify the predominant overall period. Trying values of L leads to the choice $L = 9$ as a reasonable value, and the result is displayed in Figure 4.6.

The smoothed spectra shown in Figure 4.6 provide a sensible compromise between the noisy version, shown in Figure 4.5, and a more heavily smoothed spectrum, which might lose some of the peaks. An undesirable effect of averaging can be noticed at the yearly cycle, $\omega = 1A$, where the narrow band peaks that appeared in the periodograms in Figure 4.5 have been flattened and spread out to




Fig. 4.6. The averaged periodogram of the SOI and Recruitment series $n = 453$, $n' = 480$, $L = 9$, $df = 17$, showing common peaks at the four year period, $\omega = \frac{1}{4}\Delta = 1/48$ cycles/month, the yearly period, $\omega = 1\Delta = 1/12$ cycles/month and some of its harmonics $\omega = k\Delta$ for $k = 2, 3$.

nearby frequencies. We also notice, and have marked, the appearance of harmonics of the yearly cycle, that is, frequencies of the form $\omega = k\Delta$ for $k = 1, 2, \dots$. Harmonics typically occur when a periodic component is present, but not in a sinusoidal fashion; see [Example 4.13](#).

[Figure 4.6](#) can be reproduced in R using the following commands. To compute averaged periodograms, use the Daniell kernel, and specify m , where $L = 2m + 1$ ($L = 9$ and $m = 4$ in this example). We will explain the kernel concept later in this section, specifically just prior to [Example 4.14](#).

```
par(mfrow=c(2,1))
k = kernel("daniell", 4)
soi.ave = mvspec(soi, k, log="no")
abline(v=c(.25,1,2,3), lty=2)
# Repeat above lines using rec in place of soi on line 3
soi.ave$bandwidth      # = 0.225
```

The displayed bandwidth (.225) is adjusted for the fact that the frequency scale of the plot is in terms of cycles per year instead of cycles per month (the original unit




Fig. 4.7. Figure 4.6 with the average periodogram ordinates plotted on a $\log_{10}$ scale. The display in the upper right-hand corner represents a generic 95% confidence interval.

of the data). Using (4.51), the bandwidth in terms of months is $9/480 = .01875$; the displayed value is simply converted to years, $.01875 \times 12 = .225$.

The adjusted degrees of freedom are $df = 2(9)(453)/480 \approx 17$. We can use this value for the 95% confidence intervals, with $\chi^2_{df}(.025) = 7.56$ and $\chi^2_{df}(.975) = 30.17$. Substituting into (4.56) gives the intervals in Table 4.1 for the two frequency bands identified as having the maximum power. To examine the two peak power possibilities, we may look at the 95% confidence intervals and see whether the lower limits are substantially larger than adjacent baseline spectral levels. For example, the El Niño frequency of 48 months has lower limits that exceed the values the spectrum would have if there were simply a smooth underlying spectral function without the peaks. The relative distribution of power over frequencies is different, with the SOI having less power at the lower frequency, relative to the seasonal periods, and the Recruitment series having more power at the lower or El Niño frequency.

The entries in Table 4.1 for SOI can be obtained in R as follows:

```
df = soi.ave$df          # df = 16.9875 (returned values)
U = qchisq(.025, df)     # U = 7.555916
L = qchisq(.975, df)     # L = 30.17425
soi.ave$spec[10]         # 0.0495202
```

Table 4.1. Confidence Intervals for the Spectra of the SOI and Recruitment Series

Series	ω	Period	Power	Lower	Upper
SOI	1/48	4 years	.05	.03	.11
	1/12	1 year	.12	.07	.27
Recruits $\times 10^2$	1/48	4 years	6.59	3.71	14.82
	1/12	1 year	2.19	1.24	4.93

```
soi.ave$spec[40]      # 0.1190800
# intervals
df*soi.ave$spec[10]/L # 0.0278789
df*soi.ave$spec[10]/U # 0.1113333
df*soi.ave$spec[40]/L # 0.0670396
df*soi.ave$spec[40]/U # 0.2677201
# repeat above commands with soi replaced by rec
```

Finally, [Figure 4.7](#) shows the averaged periodograms in [Figure 4.6](#) plotted on a $\log_{10}$ scale. This is the default can be obtained by removing the statement `log="no"`. Notice that the default plot also shows a generic confidence interval of the form (4.54) (with $\log$ replaced by $\log_{10}$) in the upper right-hand corner. To use it, imagine placing the tick mark on the averaged periodogram ordinate of interest; the resulting bar then constitutes an approximate 95% confidence interval for the spectrum at that frequency. We note that displaying the estimates on a log scale tends to emphasize the harmonic components.

Example 4.13 Harmonics

In the previous example, we saw that the spectra of the annual signals displayed minor peaks at the harmonics; that is, the signal spectra had a large peak at $\omega = 1\Delta = 1/12$ cycles/month (the one-year cycle) and minor peaks at its harmonics $\omega = k\Delta$ for $k = 2, 3, \dots$ (two-, three-, and so on, cycles per year). This will often be the case because most signals are not perfect sinusoids (or perfectly cyclic). In this case, the harmonics are needed to capture the non-sinusoidal behavior of the signal. As an example, consider the signal formed in [Figure 4.8](#) from a (fundamental) sinusoid oscillating at two cycles per unit time along with the second through sixth harmonics at decreasing amplitudes. In particular, the signal was formed as

$$\begin{aligned}
 x_t = & \sin(2\pi 2t) + .5 \sin(2\pi 4t) + .4 \sin(2\pi 6t) \\
 & + .3 \sin(2\pi 8t) + .2 \sin(2\pi 10t) + .1 \sin(2\pi 12t)
 \end{aligned}
 \tag{4.57}$$

for $0 \leq t \leq 1$. Notice that the signal is non-sinusoidal in appearance and rises quickly then falls slowly.

A figure similar to [Figure 4.8](#) can be generated in R as follows.

```
t = seq(0, 1, by=1/200)
amps = c(1, .5, .4, .3, .2, .1)
x = matrix(0, 201, 6)
for (j in 1:6){ x[,j] = amps[j]*sin(2*pi*t*2*j) }
```





Fig. 4.8. A signal (thick solid line) formed by a fundamental sinusoid (thin solid line) oscillating at two cycles per unit time and its harmonics as specified in (4.57).

```
x = ts(cbind(x, rowSums(x)), start=0, deltat=1/200)
ts.plot(x, lty=c(1:6, 1), lwd=c(rep(1,6), 2), ylab="Sinusoids")
names = c("Fundamental", "2nd Harmonic", "3rd Harmonic", "4th Harmonic", "5th
Harmonic", "6th Harmonic", "Formed Signal")
legend("topright", names, lty=c(1:6, 1), lwd=c(rep(1,6), 2))
```

Example 4.12 points out the necessity for having some relatively systematic procedure for deciding whether peaks are significant. The question of deciding whether a single peak is significant usually rests on establishing what we might think of as a baseline level for the spectrum, defined rather loosely as the shape that one would expect to see if no spectral peaks were present. This profile can usually be guessed by looking at the overall shape of the spectrum that includes the peaks; usually, a kind of baseline level will be apparent, with the peaks seeming to emerge from this baseline level. If the lower confidence limit for the spectral value is still greater than the baseline level at some predetermined level of significance, we may claim that frequency value as a statistically significant peak. To be consistent with our stated indifference to the upper limits, we might use a one-sided confidence interval.

An important aspect of interpreting the significance of confidence intervals and tests involving spectra is that typically, more than one frequency will be of interest, so that we will potentially be interested in simultaneous statements about a whole collection of frequencies. For example, it would be unfair to claim in Table 4.1 the two frequencies of interest as being statistically significant and all other potential candidates as nonsignificant at the overall level of $\alpha = .05$. In this case, we follow the usual statistical approach, noting that if K statements $S_1, S_2, \dots, S_K$ are made at significance level α , i.e., $P\{S_k\} = 1 - \alpha$, then the overall probability all statements are true satisfies the Bonferroni inequality

$$P\{\text{all } S_k \text{ true}\} \geq 1 - K\alpha. \quad (4.58)$$

For this reason, it is desirable to set the significance level for testing each frequency at α/K if there are K potential frequencies of interest. If, a priori, potentially $K = 10$ frequencies are of interest, setting $\alpha = .01$ would give an overall significance level of bound of .10.

The use of the confidence intervals and the necessity for smoothing requires that we make a decision about the bandwidth B_w over which the spectrum will be essentially constant. Taking too broad a band will tend to smooth out valid peaks in the data when the constant variance assumption is not met over the band. Taking too narrow a band will lead to confidence intervals so wide that peaks are no longer statistically significant. Thus, we note that there is a conflict here between variance properties or bandwidth stability, which can be improved by increasing B_w and resolution, which can be improved by decreasing B_w . A common approach is to try a number of different bandwidths and to look qualitatively at the spectral estimators for each case.

To address the problem of resolution, it should be evident that the flattening of the peaks in Figure 4.6 and Figure 4.7 was due to the fact that simple averaging was used in computing $\bar{f}(\omega)$ defined in (4.49). There is no particular reason to use simple averaging, and we might improve the estimator by employing a weighted average, say

$$\hat{f}(\omega) = \sum_{k=-m}^m h_k I(\omega_j + k/n), \quad (4.59)$$

using the same definitions as in (4.49) but where the weights $h_k > 0$ satisfy

$$\sum_{k=-m}^m h_k = 1.$$

In particular, it seems reasonable that the resolution of the estimator will improve if we use weights that decrease as distance from the center weight h_0 increases; we will return to this idea shortly. To obtain the averaged periodogram, $\bar{f}(\omega)$, in (4.59), set $h_k = L^{-1}$, for all k , where $L = 2m + 1$. The asymptotic theory established for $\bar{f}(\omega)$ still holds for $\hat{f}(\omega)$ provided that the weights satisfy the additional condition that if $m \rightarrow \infty$ as $n \rightarrow \infty$ but $m/n \rightarrow 0$, then

$$\sum_{k=-m}^m h_k^2 \rightarrow 0.$$

Under these conditions, as $n \rightarrow \infty$,

- (i) $E(\hat{f}(\omega)) \rightarrow f(\omega)$
- (ii) $\left(\sum_{k=-m}^m h_k^2\right)^{-1} \text{cov}(\hat{f}(\omega), \hat{f}(\lambda)) \rightarrow f^2(\omega) \quad \text{for } \omega = \lambda \neq 0, 1/2.$

In (ii), replace $f^2(\omega)$ by 0 if $\omega \neq \lambda$ and by $2f^2(\omega)$ if $\omega = \lambda = 0$ or $1/2$.

We have already seen these results in the case of $\bar{f}(\omega)$, where the weights are constant, $h_k = L^{-1}$, in which case $\sum_{k=-m}^m h_k^2 = L^{-1}$. The distributional properties

of (4.59) are more difficult now because $\widehat{f}(\omega)$ is a weighted linear combination of asymptotically independent χ^2 random variables. An approximation that seems to work well is to replace L by $(\sum_{k=-m}^m h_k^2)^{-1}$. That is, define

$$L_h = \left(\sum_{k=-m}^m h_k^2 \right)^{-1} \quad (4.60)$$

and use the approximation¹²

$$\frac{2L_h \widehat{f}(\omega)}{f(\omega)} \sim \chi_{2L_h}^2. \quad (4.61)$$

In analogy to (4.51), we will define the bandwidth in this case to be

$$B = \frac{L_h}{n}. \quad (4.62)$$

Using the approximation (4.61) we obtain an approximate $100(1 - \alpha)\%$ confidence interval of the form

$$\frac{2L_h \widehat{f}(\omega)}{\chi_{2L_h}^2(1 - \alpha/2)} \leq f(\omega) \leq \frac{2L_h \widehat{f}(\omega)}{\chi_{2L_h}^2(\alpha/2)} \quad (4.63)$$

for the true spectrum, $f(\omega)$. If the data are padded to n' , then replace $2L_h$ in (4.63) with $df = 2L_h n/n'$ as in (4.55).

An easy way to generate the weights in R is by repeated use of the Daniell kernel. For example, with $m = 1$ and $L = 2m + 1 = 3$, the Daniell kernel has weights $\{h_k\} = \{\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\}$; applying this kernel to a sequence of numbers, $\{u_t\}$, produces

$$\widehat{u}_t = \frac{1}{3}u_{t-1} + \frac{1}{3}u_t + \frac{1}{3}u_{t+1}.$$

We can apply the same kernel again to the $\widehat{u}_t$,

$$\widehat{\widehat{u}}_t = \frac{1}{3}\widehat{u}_{t-1} + \frac{1}{3}\widehat{u}_t + \frac{1}{3}\widehat{u}_{t+1},$$

which simplifies to

$$\widehat{\widehat{u}}_t = \frac{1}{9}u_{t-2} + \frac{2}{9}u_{t-1} + \frac{3}{9}u_t + \frac{2}{9}u_{t+1} + \frac{1}{9}u_{t+2}.$$

The modified Daniell kernel puts half weights at the end points, so with $m = 1$ the weights are $\{h_k\} = \{\frac{1}{4}, \frac{2}{4}, \frac{1}{4}\}$ and

$$\widehat{u}_t = \frac{1}{4}u_{t-1} + \frac{1}{2}u_t + \frac{1}{4}u_{t+1}.$$

¹² The approximation proceeds as follows: If $\widehat{f} \sim c\chi_\nu^2$, where c is a constant, then $E\widehat{f} \approx c\nu$ and $\text{var}\widehat{f} \approx f^2 \sum_k h_k^2 \approx c^2 2\nu$. Solving, $c \approx f \sum_k h_k^2 / 2 = f/2L_h$ and $\nu \approx 2 \left(\sum_k h_k^2 \right)^{-1} = 2L_h$.




Fig. 4.9. Smoothed spectral estimates of the SOI and Recruitment series; see [Example 4.14](#) for details.

Applying the same kernel again to $\hat{u}_t$ yields

$$\hat{\hat{u}}_t = \frac{1}{16}u_{t-2} + \frac{4}{16}u_{t-1} + \frac{6}{16}u_t + \frac{4}{16}u_{t+1} + \frac{1}{16}u_{t+2}.$$

These coefficients can be obtained in R by issuing the `kernel` command. For example, `kernel("modified.daniell", c(1,1))` would produce the coefficients of the last example. It is also possible to use different values of m , e.g., you can try `kernel("modified.daniell", c(1,2))` or `kernel("daniell", c(5,3))`. The other kernels that are currently available in R are the Dirichlet kernel and the Fejér kernel, which we will discuss shortly.

Example 4.14 Smoothed Periodogram for SOI and Recruitment

In this example, we estimate the spectra of the SOI and Recruitment series using the smoothed periodogram estimate in (4.59). We used a modified Daniell kernel twice, with $m = 3$ both times. This yields $L_h = 1 / \sum_{k=-m}^m h_k^2 = 9.232$, which is close to the value of $L = 9$ used in [Example 4.12](#). In this case, the bandwidth is $B = 9.232/480 = .019$ and the modified degrees of freedom is $df = 2L_h 453/480 = 17.43$. The weights, h_k , can be obtained and graphed in R as follows:

```
kernel("modified.daniell", c(3,3))
  coef[-6] = 0.006944 = coef[ 6]
  coef[-5] = 0.027778 = coef[ 5]
  coef[-4] = 0.055556 = coef[ 4]
  coef[-3] = 0.083333 = coef[ 3]
  coef[-2] = 0.111111 = coef[ 2]
  coef[-1] = 0.138889 = coef[ 1]
  coef[ 0] = 0.152778
plot(kernel("modified.daniell", c(3,3))) # not shown
```

The resulting spectral estimates can be viewed in Figure 4.9 and we notice that the estimates more appealing than those in Figure 4.6. Figure 4.9 was generated in R as follows; we also show how to obtain the associated bandwidth and degrees of freedom.

```
par(mfrow=c(2,1))
k = kernel("modified.daniell", c(3,3))
soi.smo = mvspec(soi, k, log="no")
abline(v=1, lty="dotted"); abline(v=1/4, lty="dotted")
## Repeat above lines with rec replacing soi in line 3
df = soi.smo$df # df = 17.42618
soi.smo$bandwidth # B = 0.2308103
```

Reissuing the `mvspec` commands with `log="no"` removed will result in a figure similar to Figure 4.7. Finally, we mention that the modified Daniell kernel is used by default. For example, an easier way to obtain `soi.smo` is to issue the command:

```
soi.smo = mvspec(soi, spans=c(7,7))
```

Notice that `spans` is a vector of odd integers, given in terms of $L = 2m + 1$ instead of m . These values give the widths of the modified Daniell smoother to be used to smooth the periodogram.

TAPERING

We are now ready to briefly introduce the concept of *tapering*; a more detailed discussion may be found in Bloomfield (2000, §9.5). Suppose x_t is a mean-zero, stationary process with spectral density $f_x(\omega)$. If we replace the original series by the tapered series

$$y_t = h_t x_t, \quad (4.64)$$

for $t = 1, 2, \dots, n$, use the modified DFT

$$d_y(\omega_j) = n^{-1/2} \sum_{t=1}^n h_t x_t e^{-2\pi i \omega_j t}, \quad (4.65)$$

and let $I_y(\omega_j) = |d_y(\omega_j)|^2$, we obtain (see Problem 4.15)

$$E[I_y(\omega_j)] = \int_{-1/2}^{1/2} W_n(\omega_j - \omega) f_x(\omega) d\omega \quad (4.66)$$

where

$$W_n(\omega) = |H_n(\omega)|^2 \quad (4.67)$$

and

$$H_n(\omega) = n^{-1/2} \sum_{t=1}^n h_t e^{-2\pi i \omega t}. \quad (4.68)$$

The value $W_n(\omega)$ is called a spectral window because, in view of (4.66), it is determining which part of the spectral density $f_x(\omega)$ is being “seen” by the estimator $I_y(\omega_j)$ on average. In the case that $h_t = 1$ for all t , $I_y(\omega_j) = I_x(\omega_j)$ is simply the periodogram of the data and the window is

$$W_n(\omega) = \frac{\sin^2(n\pi\omega)}{n \sin^2(\pi\omega)} \quad (4.69)$$

with $W_n(0) = n$, which is known as the Fejér or modified Bartlett kernel. If we consider the averaged periodogram in (4.49), namely

$$\bar{f}_x(\omega) = \frac{1}{L} \sum_{k=-m}^m I_x(\omega_j + k/n),$$

the window, $W_n(\omega)$, in (4.66) will take the form

$$W_n(\omega) = \frac{1}{nL} \sum_{k=-m}^m \frac{\sin^2[n\pi(\omega + k/n)]}{\sin^2[\pi(\omega + k/n)]}. \quad (4.70)$$

Tapers generally have a shape that enhances the center of the data relative to the extremities, such as a cosine bell of the form

$$h_t = .5 \left[1 + \cos\left(\frac{2\pi(t - \bar{t})}{n}\right) \right], \quad (4.71)$$

where $\bar{t} = (n + 1)/2$, favored by Blackman and Tukey (1959). In Figure 4.10, we have plotted the shapes of two windows, $W_n(\omega)$, for $n = 480$ and $L = 9$, when (i) $h_t \equiv 1$, in which case, (4.70) applies, and (ii) h_t is the cosine taper in (4.71). In both cases the predicted bandwidth should be $B_w = 9/480 = .01875$ cycles per point, which corresponds to the “width” of the windows shown in Figure 4.10. Both windows produce an integrated average spectrum over this band but the untapered window in the top panels shows considerable ripples over the band and outside the band. The ripples outside the band are called sidelobes and tend to introduce frequencies from outside the interval that may contaminate the desired spectral estimate within the band. For example, a large dynamic range for the values in the spectrum introduces spectra in contiguous frequency intervals several orders of magnitude greater than the value in the interval of interest. This effect is sometimes called leakage. Figure 4.10 emphasizes the suppression of the sidelobes in the Fejér kernel when a cosine taper is used.

Example 4.15 The Effect of Tapering the SOI Series

In this example, we examine the effect of tapering on the estimate of the spectrum of the SOI series. The results for the Recruitment series are similar. Figure 4.11 shows two spectral estimates plotted on a log scale. The degree of smoothing here




Fig. 4.10. Averaged Fejér window (top row) and the corresponding cosine taper window (bottom row) for $L = 9$, $n = 480$. The extra tic marks on the horizontal axis of the left-hand plots exhibit the predicted bandwidth, $B_w = 9/480 = .01875$.

is the same as in [Example 4.14](#). The dashed line in [Figure 4.11](#) shows the estimate without any tapering and hence it is the same as the estimated spectrum displayed in the top of [Figure 4.9](#). The solid line shows the result with full tapering. Notice that the tapered spectrum does a better job in separating the yearly cycle ($\omega = 1$) and the El Niño cycle ($\omega = 1/4$).

The following R session was used to generate [Figure 4.11](#). We note that, by default, `mvspec` does not taper. For full tapering, we use the argument `taper=.5` to instruct `mvspec` to taper 50% of each end of the data; any value between 0 and .5 is acceptable.

```
s0 = mvspec(soi, spans=c(7,7), plot=FALSE)           # no taper
s50 = mvspec(soi, spans=c(7,7), taper=.5, plot=FALSE) # full taper
plot(s0$freq, s0$spec, log="y", type="l", lty=2, ylab="spectrum",
     xlab="frequency")                                # dashed line
lines(s50$freq, s50$spec)                             # solid line
```

We close this section with a brief discussion of lag window estimators. First, consider the periodogram, $I(\omega_j)$, which was shown in (4.25) to be




Fig. 4.11. Smoothed spectral estimates of the SOI without tapering (dashed line) and with full tapering (solid line); see [Example 4.15](#) for details.

$$I(\omega_j) = \sum_{|h| < n} \hat{\gamma}(h) e^{-2\pi i \omega_j h}.$$

Thus, (4.59) can be written as

$$\begin{aligned} \hat{f}(\omega) &= \sum_{|k| \leq m} h_k I(\omega_j + k/n) \\ &= \sum_{|k| \leq m} h_k \sum_{|h| < n} \hat{\gamma}(h) e^{-2\pi i (\omega_j + k/n) h} \\ &= \sum_{|h| < n} g(h/n) \hat{\gamma}(h) e^{-2\pi i \omega_j h}. \end{aligned} \quad (4.72)$$

where $g(h/n) = \sum_{|k| \leq m} h_k \exp(-2\pi i k h/n)$. Equation (4.72) suggests estimators of the form

$$\tilde{f}(\omega) = \sum_{|h| \leq r} w(h/r) \hat{\gamma}(h) e^{-2\pi i \omega h} \quad (4.73)$$

where $w(\cdot)$ is a weight function, called the lag window, that satisfies

- (i) $w(0) = 1$
- (ii) $|w(x)| \leq 1$ and $w(x) = 0$ for $|x| > 1$,
- (iii) $w(x) = w(-x)$.

Note that if $w(x) = 1$ for $|x| < 1$ and $r = n$, then $\tilde{f}(\omega_j) = I(\omega_j)$, the periodogram. This result indicates the problem with the periodogram as an estimator of the spectral density is that it gives too much weight to the values of $\hat{\gamma}(h)$ when h is large, and hence is unreliable [e.g, there is only one pair of observations used in the estimate $\hat{\gamma}(n-1)$, and so on]. The smoothing window is defined to be

$$W(\omega) = \sum_{h=-r}^r w(h/r) e^{-2\pi i \omega h}, \quad (4.74)$$

and it determines which part of the periodogram will be used to form the estimate of $f(\omega)$. The asymptotic theory for $\hat{f}(\omega)$ holds for $\tilde{f}(\omega)$ under the same conditions and provided $r \rightarrow \infty$ as $n \rightarrow \infty$ but with $r/n \rightarrow 0$. We have

$$E\{\tilde{f}(\omega)\} \rightarrow f(\omega), \quad (4.75)$$

$$\frac{n}{r} \text{cov}(\tilde{f}(\omega), \tilde{f}(\lambda)) \rightarrow f^2(\omega) \int_{-1}^1 w^2(x) dx \quad \omega = \lambda \neq 0, 1/2. \quad (4.76)$$

In (4.76), replace $f^2(\omega)$ by 0 if $\omega \neq \lambda$ and by $2f^2(\omega)$ if $\omega = \lambda = 0$ or $1/2$.

Many authors have developed various windows and Brillinger (2001, Ch 3) and Brockwell and Davis (1991, Ch 10) are good sources of detailed information on this topic. We mention a few.

The rectangular lag window, which gives uniform weight in (4.73),

$$w(x) = 1, \quad |x| \leq 1,$$

corresponds to the Dirichlet smoothing window given by

$$W(\omega) = \frac{\sin(2\pi r + \pi)\omega}{\sin(\pi\omega)}. \quad (4.77)$$

This smoothing window takes on negative values, which may lead to estimates of the spectral density that are negative at various frequencies. Using (4.76) in this case, for large n we have

$$\text{var}\{\tilde{f}(\omega)\} \approx \frac{2r}{n} f^2(\omega).$$

The Parzen lag window is defined to be

$$w(x) = \begin{cases} 1 - 6x + 6|x|^3 & |x| < 1/2, \\ 2(1 - |x|)^3 & 1/2 \leq x \leq 1, \\ 0 & \text{otherwise.} \end{cases}$$

This leads to an approximate smoothing window of

$$W(\omega) = \frac{6}{\pi r^3} \frac{\sin^4(r\omega/4)}{\sin^4(\omega/2)}.$$

For large n , the variance of the estimator is approximately

$$\text{var}\{\tilde{f}(\omega)\} \approx .539 f^2(\omega)/n.$$

The Tukey-Hanning lag window has the form

$$w(x) = \frac{1}{2} (1 + \cos(\pi x)), \quad |x| \leq 1$$

which leads to the smoothing window

$$W(\omega) = \frac{1}{4}D_r(2\pi\omega - \pi/r) + \frac{1}{2}D_r(2\pi\omega) + \frac{1}{4}D_r(2\pi\omega + \pi/r)$$

where $D_r(\omega)$ is the Dirichlet kernel in (4.77). The approximate large sample variance of the estimator is

$$\text{var}\{\tilde{f}(\omega)\} \approx \frac{3r}{4n}f^2(\omega).$$

The triangular lag window, also known as the Bartlett or Fejér window, given by

$$w(x) = 1 - |x|, \quad |x| \leq 1$$

leads to the Fejér smoothing window:

$$W(\omega) = \frac{\sin^2(\pi r\omega)}{r \sin^2(\pi\omega)}.$$

In this case, (4.76) yields

$$\text{var}\{\tilde{f}(\omega)\} \approx \frac{2r}{3n}f^2(\omega).$$

The idealized rectangular smoothing window, also called the Daniell window, is given by

$$W(\omega) = \begin{cases} r & |\omega| \leq 1/2r, \\ 0 & \text{otherwise,} \end{cases}$$

and leads to the sinc lag window, namely

$$w(x) = \frac{\sin(\pi x)}{\pi x}, \quad |x| \leq 1.$$

From (4.76) we have

$$\text{var}\{\tilde{f}(\omega)\} \approx \frac{r}{n}f^2(\omega).$$

For lag window estimators, the width of the idealized rectangular window that leads to the same asymptotic variance as a given lag window estimator is sometimes called the equivalent bandwidth. For example, the bandwidth of the idealized rectangular window is $b_r = 1/r$ and the asymptotic variance is $\frac{1}{nb_r}f^2$. The asymptotic variance of the triangular window is $\frac{2r}{3n}f^2$, so setting $\frac{1}{nb_r}f^2 = \frac{2r}{3n}f^2$ and solving we get $b_r = 3/2r$ as the equivalent bandwidth.

4.6 Parametric Spectral Estimation

The methods of §4.5 lead to estimators generally referred to as nonparametric spectra because no assumption is made about the parametric form of the spectral density. In **Property 4.4**, we exhibited the spectrum of an ARMA process and we might consider basing a spectral estimator on this function, substituting the parameter estimates from

an ARMA(p, q) fit on the data into the formula for the spectral density $f_x(\omega)$ given in (4.18). Such an estimator is called a parametric spectral estimator. For convenience, a parametric spectral estimator is obtained by fitting an AR(p) to the data, where the order p is determined by one of the model selection criteria, such as AIC, AICc, and BIC, defined in (2.15)-(2.17). Parametric autoregressive spectral estimators will often have superior resolution in problems when several closely spaced narrow spectral peaks are present and are preferred by engineers for a broad variety of problems (see Kay, 1988). The development of autoregressive spectral estimators has been summarized by Parzen (1983).

If $\hat{\phi}_1, \hat{\phi}_2, \dots, \hat{\phi}_p$ and $\hat{\sigma}_w^2$ are the estimates from an AR(p) fit to x_t , then based on **Property 4.4**, a parametric spectral estimate of $f_x(\omega)$ is attained by substituting these estimates into (4.18), that is,

$$\hat{f}_x(\omega) = \frac{\hat{\sigma}_w^2}{|\hat{\phi}(e^{-2\pi i\omega})|^2}, \quad (4.78)$$

where

$$\hat{\phi}(z) = 1 - \hat{\phi}_1 z - \hat{\phi}_2 z^2 - \dots - \hat{\phi}_p z^p. \quad (4.79)$$

The asymptotic distribution of the autoregressive spectral estimator has been obtained by Berk (1974) under the conditions $p \rightarrow \infty$, $p^3/n \rightarrow 0$ as $p, n \rightarrow \infty$, which may be too severe for most applications. The limiting results imply a confidence interval of the form

$$\frac{\hat{f}_x(\omega)}{(1 + Cz_{\alpha/2})} \leq f_x(\omega) \leq \frac{\hat{f}_x(\omega)}{(1 - Cz_{\alpha/2})}, \quad (4.80)$$

where $C = \sqrt{2p/n}$ and $z_{\alpha/2}$ is the ordinate corresponding to the upper $\alpha/2$ probability of the standard normal distribution. If the sampling distribution is to be checked, we suggest applying the bootstrap estimator to get the sampling distribution of $\hat{f}_x(\omega)$ using a procedure similar to the one used for $p = 1$ in **Example 3.35**. An alternative for higher order autoregressive series is to put the AR(p) in state-space form and use the bootstrap procedure discussed in §6.7.

An interesting fact about rational spectra of the form (4.18) is that any spectral density can be approximated, arbitrarily close, by the spectrum of an AR process.

Property 4.6 AR Spectral Approximation

Let $g(\omega)$ be the spectral density of a stationary process. Then, given $\epsilon > 0$, there is a time series with the representation

$$x_t = \sum_{k=1}^p \phi_k x_{t-k} + w_t$$

where w_t is white noise with variance σ_w^2 , such that

$$|f_x(\omega) - g(\omega)| < \epsilon \quad \text{for all } \omega \in [-1/2, 1/2].$$

Moreover, p is finite and the roots of $\phi(z) = 1 - \sum_{k=1}^p \phi_k z^k$ are outside the unit circle.




Fig. 4.12. Model selection criteria AIC and BIC as a function of order p for autoregressive models fitted to the SOI series.

One drawback of the property is that it does not tell us how large p must be before the approximation is reasonable; in some situations p may be extremely large. **Property 4.6** also holds for MA and for ARMA processes in general, and a proof of the result may be found in Fuller (1996, Ch 4). We demonstrate the technique in the following example.

Example 4.16 Autoregressive Spectral Estimator for SOI

Consider obtaining results comparable to the nonparametric estimators shown in **Figure 4.6** for the SOI series. Fitting successively higher order AR(p) models for $p = 1, 2, \dots, 30$ yields a minimum BIC and a minimum AIC at $p = 15$, as shown in **Figure 4.12**. We can see from **Figure 4.12** that BIC is very definite about which model it chooses; that is, the minimum BIC is very distinct. On the other hand, it is not clear what is going to happen with AIC; that is, the minimum is not so clear, and there is some concern that AIC will start decreasing after $p = 30$. Minimum AICc selects the $p = 15$ model, but suffers from the same uncertainty as AIC. The spectrum is shown in **Figure 4.13**, and we note the strong peaks at 52 months and 12 months corresponding to the nonparametric estimators obtained in §4.5. In addition, the harmonics of the yearly period are evident in the estimated spectrum.

To perform a similar analysis in R, the command `spec.ar` can be used to fit the best model via AIC and plot the resulting spectrum. A quick way to obtain the AIC values is to run the `ar` command as follows.

```
spaic = spec.ar(soi, log="no") # min AIC spec
abline(v=frequency(soi)*1/52, lty="dotted") # El Nino Cycle
(soi.ar = ar(soi, order.max=30)) # estimates and AICs
dev.new()
plot(1:30, soi.ar$aic[-1], type="o") # plot AICs
```

R works only with the AIC in this case. To generate **Figure 4.12** we used the following code to obtain AIC, AICc, and BIC. Because AIC and AICc are nearly identical in this example, we only graphed AIC and BIC+1; we added 1 to the BIC to reduce white space in the graphic.




Fig. 4.13. Autoregressive spectral estimator for the SOI series using the AR(15) model selected by AIC, AICc, and BIC. The first peak (marked by a vertical dotted line) corresponds to the El Niño period of 52 months.

```
n = length(soi)
AIC = rep(0, 30) -> AICc -> BIC
for (k in 1:30){
  sigma2 = ar(soi, order=k, aic=FALSE)$var.pred
  BIC[k] = log(sigma2) + (k*log(n)/n)
  AICc[k] = log(sigma2) + ((n+k)/(n-k-2))
  AIC[k] = log(sigma2) + ((n+2*k)/n)
}
IC = cbind(AIC, BIC+1)
ts.plot(IC, type="o", xlab="p", ylab="AIC / BIC")
```

Finally, it should be mentioned that any parametric spectrum, say $f(\omega; \boldsymbol{\theta})$, depending on the vector parameter $\boldsymbol{\theta}$ can be estimated via the Whittle likelihood (Whittle, 1961), using the approximate properties of the discrete Fourier transform derived in Appendix C. We have that the DFTs, $d(\omega_j)$, are approximately complex normally distributed with mean zero and variance $f(\omega_j; \boldsymbol{\theta})$ and are approximately independent for $\omega_j \neq \omega_k$. This implies that an approximate log likelihood can be written in the form

$$\ln L(\mathbf{x}; \boldsymbol{\theta}) \approx - \sum_{0 < \omega_j < 1/2} \left(\ln f_x(\omega_j; \boldsymbol{\theta}) + \frac{|d(\omega_j)|^2}{f_x(\omega_j; \boldsymbol{\theta})} \right), \quad (4.81)$$

where the sum is sometimes expanded to include the frequencies $\omega_j = 0, 1/2$. If the form with the two additional frequencies is used, the multiplier of the sum will be

unity, except for the purely real points at $\omega_j = 0, 1/2$ for which the multiplier is $1/2$. For a discussion of applying the Whittle approximation to the problem of estimating parameters in an ARMA spectrum, see Anderson (1978). The Whittle likelihood is especially useful for fitting long memory models that will be discussed in Chapter 5.

4.7 Multiple Series and Cross-Spectra

The notion of analyzing frequency fluctuations using classical statistical ideas extends to the case in which there are several jointly stationary series, for example, x_t and y_t . In this case, we can introduce the idea of a correlation indexed by frequency, called the coherence. The results in Appendix C, §C.2, imply the covariance function

$$\gamma_{xy}(h) = E[(x_{t+h} - \mu_x)(y_t - \mu_y)]$$

has the representation

$$\gamma_{xy}(h) = \int_{-1/2}^{1/2} f_{xy}(\omega) e^{2\pi i \omega h} d\omega \quad h = 0, \pm 1, \pm 2, \dots, \quad (4.82)$$

where the cross-spectrum is defined as the Fourier transform

$$f_{xy}(\omega) = \sum_{h=-\infty}^{\infty} \gamma_{xy}(h) e^{-2\pi i \omega h} \quad -1/2 \leq \omega \leq 1/2, \quad (4.83)$$

assuming that the cross-covariance function is absolutely summable, as was the case for the autocovariance. The cross-spectrum is generally a complex-valued function, and it is often written as¹³

$$f_{xy}(\omega) = c_{xy}(\omega) - i q_{xy}(\omega), \quad (4.84)$$

where

$$c_{xy}(\omega) = \sum_{h=-\infty}^{\infty} \gamma_{xy}(h) \cos(2\pi \omega h) \quad (4.85)$$

and

$$q_{xy}(\omega) = \sum_{h=-\infty}^{\infty} \gamma_{xy}(h) \sin(2\pi \omega h) \quad (4.86)$$

are defined as the cospectrum and quadspectrum, respectively. Because of the relationship $\gamma_{yx}(h) = \gamma_{xy}(-h)$, it follows, by substituting into (4.83) and rearranging, that

$$f_{yx}(\omega) = \overline{f_{xy}(\omega)}. \quad (4.87)$$

This result, in turn, implies that the cospectrum and quadspectrum satisfy

¹³ For this section, it will be useful to recall the facts $e^{-i\alpha} = \cos(\alpha) - i \sin(\alpha)$ and if $z = a + ib$, then $\bar{z} = a - ib$.

$$c_{yx}(\omega) = c_{xy}(\omega) \quad (4.88)$$

and

$$q_{yx}(\omega) = -q_{xy}(\omega). \quad (4.89)$$

An important example of the application of the cross-spectrum is to the problem of predicting an output series y_t from some input series x_t through a linear filter relation such as the three-point moving average considered below. A measure of the strength of such a relation is the squared coherence function, defined as

$$\rho_{y \cdot x}^2(\omega) = \frac{|f_{yx}(\omega)|^2}{f_{xx}(\omega)f_{yy}(\omega)}, \quad (4.90)$$

where $f_{xx}(\omega)$ and $f_{yy}(\omega)$ are the individual spectra of the x_t and y_t series, respectively. Although we consider a more general form of this that applies to multiple inputs later, it is instructive to display the single input case as (4.90) to emphasize the analogy with conventional squared correlation, which takes the form

$$\rho_{yx}^2 = \frac{\sigma_{yx}^2}{\sigma_x^2 \sigma_y^2},$$

for random variables with variances σ_x^2 and σ_y^2 and covariance $\sigma_{yx} = \sigma_{xy}$. This motivates the interpretation of squared coherence and the squared correlation between two time series at frequency ω .

Example 4.17 Three-Point Moving Average

As a simple example, we compute the cross-spectrum between x_t and the three-point moving average $y_t = (x_{t-1} + x_t + x_{t+1})/3$, where x_t is a stationary input process with spectral density $f_{xx}(\omega)$. First,

$$\begin{aligned} \gamma_{xy}(h) &= \text{cov}(x_{t+h}, y_t) = \frac{1}{3} \text{cov}(x_{t+h}, x_{t-1} + x_t + x_{t+1}) \\ &= \frac{1}{3} (\gamma_{xx}(h+1) + \gamma_{xx}(h) + \gamma_{xx}(h-1)) \\ &= \frac{1}{3} \int_{-1/2}^{1/2} (e^{2\pi i \omega} + 1 + e^{-2\pi i \omega}) e^{2\pi i \omega h} f_{xx}(\omega) d\omega \\ &= \frac{1}{3} \int_{-1/2}^{1/2} [1 + 2 \cos(2\pi \omega)] f_{xx}(\omega) e^{2\pi i \omega h} d\omega, \end{aligned}$$

where we have used (4.11). Using the uniqueness of the Fourier transform, we argue from the spectral representation (4.82) that

$$f_{xy}(\omega) = \frac{1}{3} [1 + 2 \cos(2\pi \omega)] f_{xx}(\omega)$$

so that the cross-spectrum is real in this case. From Example 4.6, the spectral density of y_t is

$$\begin{aligned} f_{yy}(\omega) &= \frac{1}{9}[3 + 4 \cos(2\pi\omega) + 2 \cos(4\pi\omega)]f_{xx}(\omega) \\ &= \frac{1}{9}[1 + 2 \cos(2\pi\omega)]^2 f_{xx}(\omega), \end{aligned}$$

using the identity $\cos(2\alpha) = 2 \cos^2(\alpha) - 1$ in the last step. Substituting into (4.90) yields the squared coherence between x_t and y_t as unity over all frequencies. This is a characteristic inherited by more general linear filters, as will be shown in Problem 4.27. However, if some noise is added to the three-point moving average, the coherence is not unity; these kinds of models will be considered in detail later.

Property 4.7 Spectral Representation of a Vector Stationary Process

If the elements of the $p \times p$ autocovariance function matrix

$$\Gamma(h) = E[(\mathbf{x}_{t+h} - \boldsymbol{\mu})(\mathbf{x}_t - \boldsymbol{\mu})']$$

of a p -dimensional stationary time series, $\mathbf{x}_t = (x_{t1}, x_{t2}, \dots, x_{tp})'$, has elements satisfying

$$\sum_{h=-\infty}^{\infty} |\gamma_{jk}(h)| < \infty \quad (4.91)$$

for all $j, k = 1, \dots, p$, then $\Gamma(h)$ has the representation

$$\Gamma(h) = \int_{-1/2}^{1/2} e^{2\pi i \omega h} f(\omega) d\omega \quad h = 0, \pm 1, \pm 2, \dots, \quad (4.92)$$

as the inverse transform of the spectral density matrix, $f(\omega) = \{f_{jk}(\omega)\}$, for $j, k = 1, \dots, p$, with elements equal to the cross-spectral components. The matrix $f(\omega)$ has the representation

$$f(\omega) = \sum_{h=-\infty}^{\infty} \Gamma(h) e^{-2\pi i \omega h} \quad -1/2 \leq \omega \leq 1/2. \quad (4.93)$$

Example 4.18 Spectral Matrix of a Bivariate Process

Consider a jointly stationary bivariate process (x_t, y_t) . We arrange the autocovariances in the matrix

$$\Gamma(h) = \begin{pmatrix} \gamma_{xx}(h) & \gamma_{xy}(h) \\ \gamma_{yx}(h) & \gamma_{yy}(h) \end{pmatrix}.$$

The spectral matrix would be given by

$$f(\omega) = \begin{pmatrix} f_{xx}(\omega) & f_{xy}(\omega) \\ f_{yx}(\omega) & f_{yy}(\omega) \end{pmatrix},$$

where the Fourier transform (4.92) and (4.93) relate the autocovariance and spectral matrices.

The extension of spectral estimation to vector series is fairly obvious. For the vector series $\mathbf{x}_t = (x_{t1}, x_{t2}, \dots, x_{tp})'$, we may use the vector of DFTs, say $\mathbf{d}(\omega_j) = (d_1(\omega_j), d_2(\omega_j), \dots, d_p(\omega_j))'$, and estimate the spectral matrix by

$$\bar{f}(\omega) = L^{-1} \sum_{k=-m}^m I(\omega_j + k/n) \quad (4.94)$$

where now

$$I(\omega_j) = \mathbf{d}(\omega_j) \mathbf{d}^*(\omega_j) \quad (4.95)$$

is a $p \times p$ complex matrix.¹⁴

Again, the series may be tapered before the DFT is taken in (4.94) and we can use weighted estimation,

$$\hat{f}(\omega) = \sum_{k=-m}^m h_k I(\omega_j + k/n) \quad (4.96)$$

where $\{h_k\}$ are weights as defined in (4.59). The estimate of squared coherence between two series, y_t and x_t is

$$\hat{\rho}_{y \cdot x}^2(\omega) = \frac{|\hat{f}_{yx}(\omega)|^2}{\hat{f}_{xx}(\omega) \hat{f}_{yy}(\omega)}. \quad (4.97)$$

If the spectral estimates in (4.97) are obtained using equal weights, we will write $\bar{\rho}_{y \cdot x}^2(\omega)$ for the estimate.

Under general conditions, if $\rho_{y \cdot x}^2(\omega) > 0$ then

$$|\hat{\rho}_{y \cdot x}(\omega)| \sim AN(|\rho_{y \cdot x}(\omega)|, (1 - \rho_{y \cdot x}^2(\omega))^2 / 2L_h) \quad (4.98)$$

where L_h is defined in (4.60); the details of this result may be found in Brockwell and Davis (1991, Ch 11). We may use (4.98) to obtain approximate confidence intervals for the squared coherency $\rho_{y \cdot x}^2(\omega)$.

We can test the hypothesis that $\rho_{y \cdot x}^2(\omega) = 0$ if we use $\bar{\rho}_{y \cdot x}^2(\omega)$ for the estimate with $L > 1$,¹⁵ that is,

$$\bar{\rho}_{y \cdot x}^2(\omega) = \frac{|\bar{f}_{yx}(\omega)|^2}{\bar{f}_{xx}(\omega) \bar{f}_{yy}(\omega)}. \quad (4.99)$$

In this case, under the null hypothesis, the statistic

$$F = \frac{\bar{\rho}_{y \cdot x}^2(\omega)}{(1 - \bar{\rho}_{y \cdot x}^2(\omega))} (L - 1) \quad (4.100)$$

¹⁴ If Z is a complex matrix, then $Z^* = \bar{Z}'$ denotes the conjugate transpose operation. That is, Z^* is the result of replacing each element of Z by its complex conjugate and transposing the resulting matrix.

¹⁵ If $L = 1$ then $\bar{\rho}_{y \cdot x}^2(\omega) \equiv 1$.




Fig. 4.14. Squared coherency between the SOI and Recruitment series; $L = 19$, $n = 453$, $n' = 480$, and $\alpha = .001$. The horizontal line is $C_{.001}$.

has an approximate F -distribution with 2 and $2L - 2$ degrees of freedom. When the series have been extended to length n' , we replace $2L - 2$ by $df - 2$, where df is defined in (4.55). Solving (4.100) for a particular significance level α leads to

$$C_\alpha = \frac{F_{2,2L-2}(\alpha)}{L - 1 + F_{2,2L-2}(\alpha)} \quad (4.101)$$

as the approximate value that must be exceeded for the original squared coherence to be able to reject $\rho_{y \cdot x}^2(\omega) = 0$ at an a priori specified frequency.

Example 4.19 Coherence Between SOI and Recruitment

Figure 4.14 shows the squared coherence between the SOI and Recruitment series over a wider band than was used for the spectrum. In this case, we used $L = 19$, $df = 2(19)(453/480) \approx 36$ and $F_{2,df-2}(.001) \approx 8.53$ at the significance level $\alpha = .001$. Hence, we may reject the hypothesis of no coherence for values of $\bar{\rho}_{y \cdot x}^2(\omega)$ that exceed $C_{.001} = .32$. We emphasize that this method is crude because, in addition to the fact that the F -statistic is approximate, we are examining the squared coherence across all frequencies with the Bonferroni inequality, (4.58), in mind. Figure 4.14 also exhibits confidence bands as part of the R plotting routine. We emphasize that these bands are only valid for ω where $\rho_{y \cdot x}^2(\omega) > 0$.

In this case, the seasonal frequency and the El Niño frequencies ranging between about 3 and 7 year periods are strongly coherent. Other frequencies are also strongly coherent, although the strong coherence is less impressive because the underlying power spectrum at these higher frequencies is fairly small. Finally, we note that the coherence is persistent at the seasonal harmonic frequencies.

This example may be reproduced using the following R commands.

```
sr = mvspec(cbind(soi,rec), kernel("daniell",9), plot=FALSE)
sr$df          # df = 35.8625
f = qf(.999, 2, sr$df-2) # = 8.529792
C = f/(18+f)      # = 0.321517
plot(sr, plot.type = "coh", ci.lty = 2)
abline(h = C)
```

4.8 Linear Filters

Some of the examples of the previous sections have hinted at the possibility the distribution of power or variance in a time series can be modified by making a linear transformation. In this section, we explore that notion further by defining a linear filter and showing how it can be used to extract signals from a time series. The linear filter modifies the spectral characteristics of a time series in a predictable way, and the systematic development of methods for taking advantage of the special properties of linear filters is an important topic in time series analysis.

Recall **Property 4.3** that stated (with slightly different notation) if y_t and x_t have respective spectral densities $f_{yy}(\omega)$ and $f_{xx}(\omega)$, and if

$$y_t = \sum_{j=-\infty}^{\infty} a_j x_{t-j}, \quad \sum_{j=-\infty}^{\infty} |a_j| < \infty,$$

then

$$f_{yy}(\omega) = |A_{yx}(\omega)|^2 f_{xx}(\omega),$$

where

$$A_{yx}(\omega) = \sum_{j=-\infty}^{\infty} a_j e^{-2\pi i \omega j}$$

is the *frequency response function*. This result enables us to calculate the exact effect on the spectrum of any given filtering operation. This important property shows the spectrum of the input series is changed by filtering and the effect of the change can be characterized as a frequency-by-frequency multiplication by the squared magnitude of the frequency response function. An obvious analogy to a property of the variance in classical statistics holds, namely, if x is a random variable with variance σ_x^2 , then $y = ax$ will have variance $\sigma_y^2 = a^2 \sigma_x^2$.

Example 4.20 First Difference and Moving Average Filters

We illustrate the effect of filtering with two common examples, the first difference filter

$$y_t = \nabla x_t = x_t - x_{t-1}$$

and the symmetric moving average filter

$$y_t = \frac{1}{24}(x_{t-6} + x_{t+6}) + \frac{1}{12} \sum_{r=-5}^5 x_{t-r},$$




Fig. 4.15. SOI series (top) compared with the differenced SOI (middle) and a centered 12-month moving average (bottom).

which is a modified Daniell kernel with $m = 6$. The results of filtering the SOI series using the two filters are shown in the middle and bottom panels of **Figure 4.15**. Notice that the effect of differencing is to roughen the series because it tends to retain the higher or faster frequencies. The centered moving average smooths the series because it retains the lower frequencies and tends to attenuate the higher frequencies. In general, differencing is an example of a *high-pass filter* because it retains or passes the higher frequencies, whereas the moving average is a *low-pass filter* because it passes the lower or slower frequencies.

Notice that the slower periods are enhanced in the symmetric moving average and the seasonal or yearly frequencies are attenuated. The filtered series makes about 9 cycles in the length of the data (about one cycle every 52 months) and the moving average filter tends to enhance or extract the signal that is associated with El Niño. Moreover, by the low-pass filtering of the data, we get a better sense of the El Niño effect and its irregularity. **Figure 4.16** shows the results of a spectral analysis on the low-pass filtered SOI series. It is clear that all high frequency behavior has been removed and the El Niño cycle is accentuated; the dotted vertical line in the figure corresponds to the 52 months cycle.




Fig. 4.16. Spectral analysis of SOI after applying a 12-month moving average filter. The vertical line corresponds to the 52-month cycle.

Now, having done the filtering, it is essential to determine the exact way in which the filters change the input spectrum. We shall use (4.16) and (4.17) for this purpose. The first difference filter can be written in the form (4.15) by letting $a_0 = 1$, $a_1 = -1$, and $a_r = 0$ otherwise. This implies that

$$A_{yx}(\omega) = 1 - e^{-2\pi i\omega},$$

and the squared frequency response becomes

$$|A_{yx}(\omega)|^2 = (1 - e^{-2\pi i\omega})(1 - e^{2\pi i\omega}) = 2[1 - \cos(2\pi\omega)]. \quad (4.102)$$

The top panel of Figure 4.17 shows that the first difference filter will attenuate the lower frequencies and enhance the higher frequencies because the multiplier of the spectrum, $|A_{yx}(\omega)|^2$, is large for the higher frequencies and small for the lower frequencies. Generally, the slow rise of this kind of filter does not particularly recommend it as a procedure for retaining only the high frequencies.

For the centered 12-month moving average, we can take $a_{-6} = a_6 = 1/24$, $a_k = 1/12$ for $-5 \leq k \leq 5$ and $a_k = 0$ elsewhere. Substituting and recognizing the cosine terms gives

$$A_{yx}(\omega) = \frac{1}{12} \left[1 + \cos(12\pi\omega) + 2 \sum_{k=1}^5 \cos(2\pi\omega k) \right]. \quad (4.103)$$

Plotting the squared frequency response of this function as in Figure 4.17 shows that we can expect this filter to cut most of the frequency content above .05 cycles per point. This corresponds to eliminating periods shorter than $T = 1/.05 = 20$ points.




Fig. 4.17. Squared frequency response functions of the first difference and 12-month moving average filters.

In particular, this drives down the yearly components with periods of $T = 12$ months and enhances the El Niño frequency, which is somewhat lower. The filter is not completely efficient at attenuating high frequencies; some power contributions are left at higher frequencies, as shown in the function $|A_{yx}(\omega)|^2$ and in the spectrum of the moving average shown in [Figure 4.4](#).

The following R session shows how to filter the data, perform the spectral analysis of this example, and plot the squared frequency response curves of the difference and moving average filters.

```
par(mfrow=c(3,1), mar=c(3,3,1,1), mgp=c(1.6,.6,0))
plot(soi)                                # plot data
plot(diff(soi))                          # plot first difference
k = kernel("modified.daniell", 6)        # filter weights
plot(soif <- kernapply(soi, k))           # plot 12 month filter
dev.new()
spectrum(soif, spans=9, log="no")        # spectral analysis
abline(v=12/52, lty="dashed")
dev.new()
##-- frequency responses --##
par(mfrow=c(2,1), mar=c(3,3,1,1), mgp=c(1.6,.6,0))
w = seq(0, .5, by=.01)
FRdiff = abs(1-exp(2i*pi*w))^2
plot(w, FRdiff, type='l', xlab='frequency')
u = cos(2*pi*w)+cos(4*pi*w)+cos(6*pi*w)+cos(8*pi*w)+cos(10*pi*w)
FRma = ((1 + cos(12*pi*w) + 2*u)/12)^2
plot(w, FRma, type='l', xlab='frequency')
```

The two filters discussed in the previous example were different in that the frequency response function of the first difference was complex-valued, whereas the frequency response of the moving average was purely real. A short derivation similar to that used to verify (4.17) shows, when x_t and y_t are related by the linear filter relation (4.15), the cross-spectrum satisfies

$$f_{yx}(\omega) = A_{yx}(\omega)f_{xx}(\omega),$$

so the frequency response is of the form

$$A_{yx}(\omega) = \frac{f_{yx}(\omega)}{f_{xx}(\omega)} \quad (4.104)$$

$$= \frac{c_{yx}(\omega)}{f_{xx}(\omega)} - i \frac{q_{yx}(\omega)}{f_{xx}(\omega)}, \quad (4.105)$$

where we have used (4.84) to get the last form. Then, we may write (4.105) in polar coordinates as

$$A_{yx}(\omega) = |A_{yx}(\omega)| \exp\{-i \phi_{yx}(\omega)\}, \quad (4.106)$$

where the amplitude and phase of the filter are defined by

$$|A_{yx}(\omega)| = \frac{\sqrt{c_{yx}^2(\omega) + q_{yx}^2(\omega)}}{f_{xx}(\omega)} \quad (4.107)$$

and

$$\phi_{yx}(\omega) = \tan^{-1} \left(-\frac{q_{yx}(\omega)}{c_{yx}(\omega)} \right). \quad (4.108)$$

A simple interpretation of the phase of a linear filter is that it exhibits time delays as a function of frequency in the same way as the spectrum represents the variance as a function of frequency. Additional insight can be gained by considering the simple delaying filter

$$y_t = Ax_{t-D},$$

where the series gets replaced by a version, amplified by multiplying by A and delayed by D points. For this case,

$$f_{yx}(\omega) = Ae^{-2\pi i \omega D} f_{xx}(\omega),$$

and the amplitude is $|A|$, and the phase is

$$\phi_{yx}(\omega) = -2\pi\omega D,$$

or just a linear function of frequency ω . For this case, applying a simple time delay causes phase delays that depend on the frequency of the periodic component being delayed. Interpretation is further enhanced by setting

$$x_t = \cos(2\pi\omega t),$$

in which case

$$y_t = A \cos(2\pi\omega t - 2\pi\omega D).$$

Thus, the output series, y_t , has the same period as the input series, x_t , but the amplitude of the output has increased by a factor of $|A|$ and the phase has been changed by a factor of $-2\pi\omega D$.

Example 4.21 Difference and Moving Average Filters

We consider calculating the amplitude and phase of the two filters discussed in [Example 4.20](#). The case for the moving average is easy because $A_{yx}(\omega)$ given in [\(4.103\)](#) is purely real. So, the amplitude is just $|A_{yx}(\omega)|$ and the phase is $\phi_{yx}(\omega) = 0$. In general, symmetric ($a_j = a_{-j}$) filters have zero phase. The first difference, however, changes this, as we might expect from the example above involving the time delay filter. In this case, the squared amplitude is given in [\(4.102\)](#). To compute the phase, we write

$$\begin{aligned} A_{yx}(\omega) &= 1 - e^{-2\pi i\omega} = e^{-i\pi\omega} (e^{i\pi\omega} - e^{-i\pi\omega}) \\ &= 2ie^{-i\pi\omega} \sin(\pi\omega) = 2 \sin^2(\pi\omega) + 2i \cos(\pi\omega) \sin(\pi\omega) \\ &= \frac{c_{yx}(\omega)}{f_{xx}(\omega)} - i \frac{q_{yx}(\omega)}{f_{xx}(\omega)}, \end{aligned}$$

so

$$\phi_{yx}(\omega) = \tan^{-1} \left(-\frac{q_{yx}(\omega)}{c_{yx}(\omega)} \right) = \tan^{-1} \left(\frac{\cos(\pi\omega)}{\sin(\pi\omega)} \right).$$

Noting that

$$\cos(\pi\omega) = \sin(-\pi\omega + \pi/2)$$

and that

$$\sin(\pi\omega) = \cos(-\pi\omega + \pi/2),$$

we get

$$\phi_{yx}(\omega) = -\pi\omega + \pi/2,$$

and the phase is again a linear function of frequency.

The above tendency of the frequencies to arrive at different times in the filtered version of the series remains as one of two annoying features of the difference type filters. The other weakness is the gentle increase in the frequency response function. If low frequencies are really unimportant and high frequencies are to be preserved, we would like to have a somewhat sharper response than is obvious in [Figure 4.17](#). Similarly, if low frequencies are important and high frequencies are not, the moving average filters are also not very efficient at passing the low frequencies and attenuating the high frequencies. Improvement is possible by designing better and longer filters, but we do not discuss this here.

We will occasionally use results for multivariate series $\mathbf{x}_t = (x_{t1}, \dots, x_{tp})'$ that are comparable to the simple property shown in [\(4.17\)](#). Consider the matrix filter

$$\mathbf{y}_t = \sum_{j=-\infty}^{\infty} A_j \mathbf{x}_{t-j}, \quad (4.109)$$

where $\{A_j\}$ denotes a sequence of $q \times p$ matrices such that $\sum_{j=-\infty}^{\infty} \|A_j\| < \infty$ and $\|\cdot\|$ denotes any matrix norm, $\mathbf{x}_t = (x_{t1}, \dots, x_{tp})'$ is a $p \times 1$ stationary vector process with mean vector $\boldsymbol{\mu}_x$ and $p \times p$, matrix covariance function $\Gamma_{xx}(h)$ and spectral matrix $f_{xx}(\omega)$, and $\mathbf{y}_t$ is the $q \times 1$ vector output process. Then, we can obtain the following property.

Property 4.8 Output Spectral Matrix of a Linearly Filtered Stationary Vector Series

The spectral matrix of the filtered output $\mathbf{y}_t$ in (4.109) is related to the spectrum of the input $\mathbf{x}_t$ by

$$f_{yy}(\omega) = \mathcal{A}(\omega) f_{xx}(\omega) \mathcal{A}^*(\omega), \quad (4.110)$$

where the matrix frequency response function $\mathcal{A}(\omega)$ is defined by

$$\mathcal{A}(\omega) = \sum_{j=-\infty}^{\infty} A_j \exp(-2\pi i \omega j). \quad (4.111)$$

4.9 Lagged Regression Models

One of the intriguing possibilities offered by the coherence analysis of the relation between the SOI and Recruitment series discussed in [Example 4.19](#) would be extending classical regression to the analysis of lagged regression models of the form

$$y_t = \sum_{r=-\infty}^{\infty} \beta_r x_{t-r} + v_t, \quad (4.112)$$

where v_t is a stationary noise process, x_t is the observed input series, and y_t is the observed output series. We are interested in estimating the filter coefficients β_r relating the adjacent lagged values of x_t to the output series y_t .

In the case of SOI and Recruitment series, we might identify the El Niño driving series, SOI, as the input, x_t , and y_t , the Recruitment series, as the output. In general, there will be more than a single possible input series and we may envision a $q \times 1$ vector of driving series. This multivariate input situation is covered in Chapter 7. The model given by (4.112) is useful under several different scenarios, corresponding to different assumptions that can be made about the components.

We assume that the inputs and outputs have zero means and are jointly stationary with the 2×1 vector process $(x_t, y_t)'$ having a spectral matrix of the form

$$f(\omega) = \begin{pmatrix} f_{xx}(\omega) & f_{xy}(\omega) \\ f_{yx}(\omega) & f_{yy}(\omega) \end{pmatrix}. \quad (4.113)$$

Here, $f_{xy}(\omega)$ is the cross-spectrum relating the input x_t to the output y_t , and $f_{xx}(\omega)$ and $f_{yy}(\omega)$ are the spectra of the input and output series, respectively. Generally, we

observe two series, regarded as input and output and search for regression functions $\{\beta_t\}$ relating the inputs to the outputs. We assume all autocovariance functions satisfy the absolute summability conditions of the form (4.33).

Then, minimizing the mean squared error

$$MSE = E \left(y_t - \sum_{r=-\infty}^{\infty} \beta_r x_{t-r} \right)^2 \quad (4.114)$$

leads to the usual orthogonality conditions

$$E \left[\left(y_t - \sum_{r=-\infty}^{\infty} \beta_r x_{t-r} \right) x_{t-s} \right] = 0 \quad (4.115)$$

for all $s = 0, \pm 1, \pm 2, \dots$. Taking the expectations inside leads to the normal equations

$$\sum_{r=-\infty}^{\infty} \beta_r \gamma_{xx}(s-r) = \gamma_{yx}(s) \quad (4.116)$$

for $s = 0, \pm 1, \pm 2, \dots$. These equations might be solved, with some effort, if the covariance functions were known exactly. If data (x_t, y_t) for $t = 1, \dots, n$ are available, we might use a finite approximation to the above equations with $\hat{\gamma}_{xx}(h)$ and $\hat{\gamma}_{yx}(h)$ substituted into (4.116). If the regression vectors are essentially zero for $|s| \geq M/2$, and $M < n$, the system (4.116) would be of full rank and the solution would involve inverting an $(M-1) \times (M-1)$ matrix.

A frequency domain approximate solution is easier in this case for two reasons. First, the computations depend on spectra and cross-spectra that can be estimated from sample data using the techniques of §4.6. In addition, no matrices will have to be inverted, although the frequency domain ratio will have to be computed for each frequency. In order to develop the frequency domain solution, substitute the representation (4.92) into the normal equations, using the convention defined in (4.113). The left side of (4.116) can then be written in the form

$$\int_{-1/2}^{1/2} \sum_{r=-\infty}^{\infty} \beta_r e^{2\pi i \omega(s-r)} f_{xx}(\omega) d\omega = \int_{-1/2}^{1/2} e^{2\pi i \omega s} B(\omega) f_{xx}(\omega) d\omega,$$

where

$$B(\omega) = \sum_{r=-\infty}^{\infty} \beta_r e^{-2\pi i \omega r} \quad (4.117)$$

is the Fourier transform of the regression coefficients β_t . Now, because $\gamma_{yx}(s)$ is the inverse transform of the cross-spectrum $f_{yx}(\omega)$, we might write the system of equations in the frequency domain, using the uniqueness of the Fourier transform, as

$$B(\omega) f_{xx}(\omega) = f_{yx}(\omega), \quad (4.118)$$

which then become the analogs of the usual normal equations. Then, we may take




Fig. 4.18. Estimated impulse response functions relating SOI to Recruitment (top) and Recruitment to SOI (bottom) $L = 15$, $M = 32$.

$$\widehat{B}(\omega_k) = \frac{\widehat{f}_{yx}(\omega_k)}{\widehat{f}_{xx}(\omega_k)} \quad (4.119)$$

as the estimator for the Fourier transform of the regression coefficients, evaluated at some subset of fundamental frequencies $\omega_k = k/M$ with $M \ll n$. Generally, we assume smoothness of $B(\cdot)$ over intervals of the form $\{\omega_k + \ell/n; \ell = -(L-1)/2, \dots, (L-1)/2\}$. The inverse transform of the function $\widehat{B}(\omega)$ would give $\widehat{\beta}_t$, and we note that the discrete time approximation can be taken as

$$\widehat{\beta}_t = M^{-1} \sum_{k=0}^{M-1} \widehat{B}(\omega_k) e^{2\pi i \omega_k t} \quad (4.120)$$

for $t = 0, \pm 1, \pm 2, \dots, \pm(M/2-1)$. If we were to use (4.120) to define $\widehat{\beta}_t$ for $|t| \geq M/2$, we would end up with a sequence of coefficients that is periodic with a period of M . In practice we define $\widehat{\beta}_t = 0$ for $|t| \geq M/2$ instead. **Problem 4.29** explores the error resulting from this approximation.

Example 4.22 Lagged Regression for SOI and Recruitment

The high coherence between the SOI and Recruitment series noted in **Example 4.19** suggests a lagged regression relation between the two series. A natural direction for the implication in this situation is implied because we feel that the sea surface temperature or SOI should be the input and the Recruitment series should be the output. With this in mind, let x_t be the SOI series and y_t the Recruitment series.

Although we think naturally of the SOI as the input and the Recruitment as the output, two input-output configurations are of interest. With SOI as the input, the model is

$$y_t = \sum_{r=-\infty}^{\infty} a_r x_{t-r} + w_t$$

whereas a model that reverses the two roles would be

$$x_t = \sum_{r=-\infty}^{\infty} b_r y_{t-r} + v_t,$$

where w_t and v_t are white noise processes. Even though there is no plausible environmental explanation for the second of these two models, displaying both possibilities helps to settle on a parsimonious transfer function model.

Based on the script [LagReg](#) in [astsa](#), the estimated regression or impulse response function for SOI, with $M = 32$ and $L = 15$ is

```
LagReg(soi, rec, L=15, M=32, threshold=6)
```

```
lag s      beta(s)
[1,]      5 -18.479306
[2,]      6 -12.263296
[3,]      7 -8.539368
[4,]      8 -6.984553
```

The prediction equation is

```
rec(t) = alpha + sum_s[ beta(s)*soi(t-s) ], where alpha = 65.97
```

```
MSE = 414.08
```

Note the negative peak at a lag of five points in the top of [Figure 4.18](#); in this case, SOI is the input series. The fall-off after lag five seems to be approximately exponential and a possible model is

$$y_t = 66 - 18.5x_{t-5} - 12.3x_{t-6} - 8.5x_{t-7} - 7x_{t-8} + w_t.$$

If we examine the inverse relation, namely, a regression model with the Recruitment series y_t as the input, the bottom of [Figure 4.18](#) implies a much simpler model,

```
LagReg(rec, soi, L=15, M=32, inverse=TRUE, threshold=.01)
```

```
lag s      beta(s)
[1,]      4 0.01593167
[2,]      5 -0.02120013
```

The prediction equation is

```
soi(t) = alpha + sum_s[ beta(s)*rec(t+s) ], where alpha = 0.41
```

```
MSE = 0.07
```

depending on only two coefficients, namely,

$$x_t = .41 + .016y_{t+4} - .02y_{t+5} + v_t.$$

Multiplying both sides by $50B^5$ and rearranging, we have

$$(1 - .8B)y_t = 20.5 - 50B^5x_t + \epsilon_t,$$

where ϵ_t is white noise, as our final, parsimonious model.

The example shows we can get a clean estimator for the transfer functions relating the two series if the coherence $\hat{\rho}_{xy}^2(\omega)$ is large. The reason is that we can write the minimized mean squared error ([4.114](#)) as

$$MSE = E \left[(y_t - \sum_{r=-\infty}^{\infty} \beta_r x_{t-r}) y_t \right] = \gamma_{yy}(0) - \sum_{r=-\infty}^{\infty} \beta_r \gamma_{xy}(-r),$$

using the result about the orthogonality of the data and error term in the Projection theorem. Then, substituting the spectral representations of the autocovariance and cross-covariance functions and identifying the Fourier transform (4.117) in the result leads to

$$\begin{aligned} MSE &= \int_{-1/2}^{1/2} [f_{yy}(\omega) - B(\omega)f_{xy}(\omega)] d\omega \\ &= \int_{-1/2}^{1/2} f_{yy}(\omega)[1 - \rho_{yx}^2(\omega)] d\omega, \end{aligned} \quad (4.121)$$

where $\rho_{yx}^2(\omega)$ is just the squared coherence given by (4.90). The similarity of (4.121) to the usual mean square error that results from predicting y from x is obvious. In that case, we would have

$$E(y - \beta x)^2 = \sigma_y^2(1 - \rho_{xy}^2)$$

for jointly distributed random variables x and y with zero means, variances σ_x^2 and σ_y^2 , and covariance $\sigma_{xy} = \rho_{xy}\sigma_x\sigma_y$. Because the mean squared error in (4.121) satisfies $MSE \geq 0$ with $f_{yy}(\omega)$ a non-negative function, it follows that the coherence satisfies

$$0 \leq \rho_{xy}^2(\omega) \leq 1$$

for all ω . Furthermore, Problem 4.30 shows the squared coherence is one when the output are linearly related by the filter relation (4.112), and there is no noise, i.e., $v_t = 0$. Hence, the multiple coherence gives a measure of the association or correlation between the input and output series as a function of frequency.

The matter of verifying that the F -distribution claimed for (4.100) will hold when the sample coherence values are substituted for theoretical values still remains. Again, the form of the F -statistic is exactly analogous to the usual t -test for no correlation in a regression context. We give an argument leading to this conclusion later using the results in Appendix C, §C.3. Another question that has not been resolved in this section is the extension to the case of multiple inputs $x_{t1}, x_{t2}, \dots, x_{tq}$. Often, more than just a single input series is present that can possibly form a lagged predictor of the output series y_t . An example is the cardiovascular mortality series that depended on possibly a number of pollution series and temperature. We discuss this particular extension as a part of the multivariate time series techniques considered in Chapter 7.

4.10 Signal Extraction and Optimum Filtering

A model closely related to regression can be developed by assuming again that

$$y_t = \sum_{r=-\infty}^{\infty} \beta_r x_{t-r} + v_t, \quad (4.122)$$

but where the β_s are known and x_t is some unknown random signal that is uncorrelated with the noise process v_t . In this case, we observe only y_t and are interested in an estimator for the signal x_t of the form

$$\hat{x}_t = \sum_{r=-\infty}^{\infty} a_r y_{t-r}. \quad (4.123)$$

In the frequency domain, it is convenient to make the additional assumptions that the series x_t and v_t are both mean-zero stationary series with spectra $f_{xx}(\omega)$ and $f_{vv}(\omega)$, often referred to as the signal spectrum and noise spectrum, respectively. Often, the special case $\beta_t = \delta_t$, in which δ_t is the Kronecker delta, is of interest because (4.122) reduces to the simple signal plus noise model

$$y_t = x_t + v_t \quad (4.124)$$

in that case. In general, we seek the set of filter coefficients a_t that minimize the mean squared error of estimation, say,

$$MSE = E \left[\left(x_t - \sum_{r=-\infty}^{\infty} a_r y_{t-r} \right)^2 \right]. \quad (4.125)$$

This problem was originally solved by Kolmogorov (1941) and by Wiener (1949), who derived the result in 1941 and published it in classified reports during World War II.

We can apply the orthogonality principle to write

$$E \left[\left(x_t - \sum_{r=-\infty}^{\infty} a_r y_{t-r} \right) y_{t-s} \right] = 0$$

for $s = 0, \pm 1, \pm 2, \dots$, which leads to

$$\sum_{r=-\infty}^{\infty} a_r \gamma_{yy}(s-r) = \gamma_{xy}(s),$$

to be solved for the filter coefficients. Substituting the spectral representations for the autocovariance functions into the above and identifying the spectral densities through the uniqueness of the Fourier transform produces

$$A(\omega) f_{yy}(\omega) = f_{xy}(\omega), \quad (4.126)$$

where $A(\omega)$ and the optimal filter a_t are Fourier transform pairs for $B(\omega)$ and β_t . Now, a special consequence of the model is that (see [Problem 4.27](#))

$$f_{xy}(\omega) = \overline{B(\omega)} f_{xx}(\omega) \quad (4.127)$$

and

$$f_{yy}(\omega) = |B(\omega)|^2 f_{xx}(\omega) + f_{vv}(\omega), \quad (4.128)$$

implying the optimal filter would be Fourier transform of

$$A(\omega) = \frac{\overline{B(\omega)}}{\left(|B(\omega)|^2 + \frac{f_{vv}(\omega)}{f_{xx}(\omega)}\right)}, \quad (4.129)$$

where the second term in the denominator is just the inverse of the signal to noise ratio, say,

$$\text{SNR}(\omega) = \frac{f_{xx}(\omega)}{f_{vv}(\omega)}. \quad (4.130)$$

The result shows the optimum filters can be computed for this model if the signal and noise spectra are both known or if we can assume knowledge of the signal-to-noise ratio $\text{SNR}(\omega)$ as function of frequency. In Chapter 7, we show some methods for estimating these two parameters in conjunction with random effects analysis of variance models, but we assume here that it is possible to specify the signal-to-noise ratio a priori. If the signal-to-noise ratio is known, the optimal filter can be computed by the inverse transform of the function $A(\omega)$. It is more likely that the inverse transform will be intractable and a finite filter approximation like that used in the previous section can be applied to the data. In this case, we will have

$$a_t^M = M^{-1} \sum_{k=0}^{M-1} A(\omega_k) e^{2\pi i \omega_k t} \quad (4.131)$$

as the estimated filter function. It will often be the case that the form of the specified frequency response will have some rather sharp transitions between regions where the signal-to-noise ratio is high and regions where there is little signal. In these cases, the shape of the frequency response function will have ripples that can introduce frequencies at different amplitudes. An aesthetic solution to this problem is to introduce tapering as was done with spectral estimation in (4.64)-(4.71). We use below the tapered filter $\tilde{a}_t = h_t a_t$ where h_t is the cosine taper given in (4.71). The squared frequency response of the resulting filter will be $|\tilde{A}(\omega)|^2$, where

$$\tilde{A}(\omega) = \sum_{t=-\infty}^{\infty} a_t h_t e^{-2\pi i \omega t}. \quad (4.132)$$

The results are illustrated in the following example that extracts the El Niño component of the sea surface temperature series.

Example 4.23 Estimating the El Niño Signal via Optimal Filters

Figure 4.6 shows the spectrum of the SOI series, and we note that essentially two components have power, the El Niño frequency of about .02 cycles per month (the four-year cycle) and a yearly frequency of about .08 cycles per month (the annual cycle). We assume, for this example, that we wish to preserve the lower frequency as signal and to eliminate the higher order frequencies, and in particular, the annual cycle. In this case, we assume the simple signal plus noise model




Fig. 4.19. Filter coefficients (top) and frequency response functions (bottom) for designed SOI filters.




Fig. 4.20. Original SOI series (top) compared to filtered version showing the estimated El Niño temperature signal (bottom).

$$y_t = x_t + v_t,$$

so that there is no convolving function β_t . Furthermore, the signal-to-noise ratio is assumed to be high to about .06 cycles per month and zero thereafter. The optimal frequency response was assumed to be unity to .05 cycles per point and then to decay linearly to zero in several steps. [Figure 4.19](#) shows the coefficients as specified by (4.131) with $M = 64$, as well as the frequency response function given by (4.132), of the cosine tapered coefficients; recall [Figure 4.10](#), where we demonstrated the need for tapering to avoid severe ripples in the window. The constructed response function is compared to the ideal window in [Figure 4.19](#).

[Figure 4.20](#) shows the original and filtered SOI index, and we see a smooth extracted signal that conveys the essence of the underlying El Niño signal. The frequency response of the designed filter can be compared with that of the symmetric 12-month moving average applied to the same series in [Example 4.20](#). The filtered series, shown in [Figure 4.15](#), shows a good deal of higher frequency chatter riding on the smoothed version, which has been introduced by the higher frequencies that leak through in the squared frequency response, as in [Figure 4.17](#).

The analysis can be replicated using the script [SigExtract](#).

```
SigExtract(soi, L=9, M=64, max.freq=.05)
```

The design of finite filters with a specified frequency response requires some experimentation with various target frequency response functions and we have only touched on the methodology here. The filter designed here, sometimes called a low-pass filter reduces the high frequencies and keeps or passes the low frequencies. Alternately, we could design a high-pass filter to keep high frequencies if that is where the signal is located. An example of a simple high-pass filter is the first difference with a frequency response that is shown in [Figure 4.17](#). We can also design band-pass filters that keep frequencies in specified bands. For example, seasonal adjustment filters are often used in economics to reject seasonal frequencies while keeping both high frequencies, lower frequencies, and trend (see, for example, Grether and Nerlove, 1970).

The filters we have discussed here are all symmetric two-sided filters, because the designed frequency response functions were purely real. Alternatively, we may design recursive filters to produce a desired response. An example of a recursive filter is one that replaces the input x_t by the filtered output

$$y_t = \sum_{k=1}^p \phi_k y_{t-k} + x_t - \sum_{k=1}^q \theta_k x_{t-k}. \quad (4.133)$$

Note the similarity between (4.133) and the ARMA(p, q) model, in which the white noise component is replaced by the input. Transposing the terms involving y_t and using the basic linear filter result in [Property 4.3](#) leads to

$$f_y(\omega) = \frac{|\theta(e^{-2\pi i\omega})|^2}{|\phi(e^{-2\pi i\omega})|^2} f_x(\omega), \quad (4.134)$$

where

$$\phi(e^{-2\pi i\omega}) = 1 - \sum_{k=1}^p \phi_k e^{-2\pi i k \omega}$$

and

$$\theta(e^{-2\pi i\omega}) = 1 - \sum_{k=1}^q \theta_k e^{-2\pi i k \omega}.$$

Recursive filters such as those given by (4.134) distort the phases of arriving frequencies, and we do not consider the problem of designing such filters in any detail.

4.11 Spectral Analysis of Multidimensional Series

Multidimensional series of the form $x_{\mathbf{s}}$, where $\mathbf{s} = (s_1, s_2, \dots, s_r)'$ is an r -dimensional vector of spatial coordinates or a combination of space and time coordinates, were introduced in §1.7. The example given there, shown in Figure 1.17, was a collection of temperature measurements taking on a rectangular field. These data would form a two-dimensional process, indexed by row and column in space. In that section, the multidimensional autocovariance function of an r -dimensional stationary series was given as $\gamma_x(\mathbf{h}) = E[x_{\mathbf{s}+\mathbf{h}}x_{\mathbf{s}}]$, where the multidimensional lag vector is $\mathbf{h} = (h_1, h_2, \dots, h_r)'$.

The multidimensional wavenumber spectrum is given as the Fourier transform of the autocovariance, namely,

$$f_x(\boldsymbol{\omega}) = \sum_{\mathbf{h}} \gamma_x(\mathbf{h}) e^{-2\pi i \boldsymbol{\omega}' \mathbf{h}}. \quad (4.135)$$

Again, the inverse result

$$\gamma_x(\mathbf{h}) = \int_{-1/2}^{1/2} f_x(\boldsymbol{\omega}) e^{2\pi i \boldsymbol{\omega}' \mathbf{h}} d\boldsymbol{\omega} \quad (4.136)$$

holds, where the integral is over the multidimensional range of the vector $\boldsymbol{\omega}$. The wavenumber argument is exactly analogous to the frequency argument, and we have the corresponding intuitive interpretation as the cycling rate ω_i per distance traveled s_i in the i -th direction.

Two-dimensional processes occur often in practical applications, and the representations above reduce to

$$f_x(\omega_1, \omega_2) = \sum_{h_1=-\infty}^{\infty} \sum_{h_2=-\infty}^{\infty} \gamma_x(h_1, h_2) e^{-2\pi i(\omega_1 h_1 + \omega_2 h_2)} \quad (4.137)$$

and

$$\gamma_x(h_1, h_2) = \int_{-1/2}^{1/2} \int_{-1/2}^{1/2} f_x(\omega_1, \omega_2) e^{2\pi i(\omega_1 h_1 + \omega_2 h_2)} d\omega_1 d\omega_2 \quad (4.138)$$

in the case $r = 2$. The notion of linear filtering generalizes easily to the two-dimensional case by defining the impulse response function a_{s_1, s_2} and the spatial filter output as

$$y_{s_1, s_2} = \sum_{u_1} \sum_{u_2} a_{u_1, u_2} x_{s_1 - u_1, s_2 - u_2}. \quad (4.139)$$

The spectrum of the output of this filter can be derived as

$$f_y(\omega_1, \omega_2) = |A(\omega_1, \omega_2)|^2 f_x(\omega_1, \omega_2), \quad (4.140)$$

where

$$A(\omega_1, \omega_2) = \sum_{u_1} \sum_{u_2} a_{u_1, u_2} e^{-2\pi i(\omega_1 u_1 + \omega_2 u_2)}. \quad (4.141)$$

These results are analogous to those in the one-dimensional case, described by **Property 4.3**.

The multidimensional DFT is also a straightforward generalization of the univariate expression. In the two-dimensional case with data on a rectangular grid, $\{x_{s_1, s_2}; s_1 = 1, \dots, n_1, s_2 = 1, \dots, n_2\}$, we will write, for $-1/2 \leq \omega_1, \omega_2 \leq 1/2$,

$$d(\omega_1, \omega_2) = (n_1 n_2)^{-1/2} \sum_{s_1=1}^{n_1} \sum_{s_2=1}^{n_2} x_{s_1, s_2} e^{-2\pi i(\omega_1 s_1 + \omega_2 s_2)} \quad (4.142)$$

as the two-dimensional DFT, where the frequencies ω_1, ω_2 are evaluated at multiples of $(1/n_1, 1/n_2)$ on the spatial frequency scale. The two-dimensional wavenumber spectrum can be estimated by the smoothed sample wavenumber spectrum

$$\bar{f}_x(\omega_1, \omega_2) = (L_1 L_2)^{-1} \sum_{\ell_1, \ell_2} |d(\omega_1 + \ell_1/n_1, \omega_2 + \ell_2/n_2)|^2, \quad (4.143)$$

where the sum is taken over the grid $\{-m_j \leq \ell_j \leq m_j; j = 1, 2\}$, where $L_1 = 2m_1 + 1$ and $L_2 = 2m_2 + 1$. The statistic

$$\frac{2L_1 L_2 \bar{f}_x(\omega_1, \omega_2)}{f_x(\omega_1, \omega_2)} \sim \chi_{2L_1 L_2}^2 \quad (4.144)$$

can be used to set confidence intervals or make approximate tests against a fixed assumed spectrum $f_0(\omega_1, \omega_2)$. We may also extend this analysis to weighted estimation and window estimation as discussed in §4.5.

Example 4.24 Soil Surface Temperatures

As an example, consider the periodogram of the two-dimensional temperature series shown in **Figure 1.17** and analyzed by Bazza et al. (1988). We recall the spatial coordinates in this case will be (s_1, s_2) , which define the spatial coordinates rows and columns so that the frequencies in the two directions will be expressed as cycles per row and cycles per column. **Figure 4.21** shows the periodogram of the two-dimensional temperature series, and we note the ridge of strong spectral peaks running over rows at a column frequency of zero. An obvious periodic component




Fig. 4.21. Two-dimensional periodogram of soil temperature profile showing peak at .0625 cycles/row. The period is 16 rows, and this corresponds to $16 \times 17 \text{ ft} = 272 \text{ ft}$.

appears at frequencies of .0625 and $-.0625$ cycles per row, which corresponds to 16 rows or about 272 ft. On further investigation of previous irrigation patterns over this field, treatment levels of salt varied periodically over columns. This analysis is extended in [Problem 4.21](#), where we recover the salt treatment profile over rows and compare it to a signal, computed by averaging over columns.

[Figure 4.21](#) may be reproduced in R as follows. In the code for this example, the periodogram is computed in one step as `per`; the rest of the code is simply manipulation to obtain a nice graphic.

```
per = Mod(fft(soiltemp-mean(soiltemp))/sqrt(64*36))^2
per2 = cbind(per[1:32,18:2], per[1:32,1:18])
per3 = rbind(per2[32:2,],per2)
par(mar=c(1,2.5,0,0)+.1)
persp(-31:31/64, -17:17/36, per3, phi=30, theta=30, expand=.6,
      ticktype="detailed", xlab="cycles/row", ylab="cycles/column",
      zlab="Periodogram Ordinate")
```

Another application of two-dimensional spectral analysis of agricultural field trials is given in McBratney and Webster (1981), who used it to detect ridge and furrow patterns in yields. The requirement for regular, equally spaced samples on fairly large grids has tended to limit enthusiasm for strict two-dimensional spectral analysis. An exception is when a propagating signal from a given velocity and azimuth

is present so predicting the wavenumber spectrum as a function of velocity and azimuth becomes feasible (see Shumway et al., 1999).

Problems

Section 4.2

4.1 Repeat the simulations and analyses in [Example 4.1](#) and [Example 4.2](#) with the following changes:

- (a) Change the sample size to $n = 128$ and generate and plot the same series as in [Example 4.1](#):

$$x_{t1} = 2 \cos(2\pi \cdot 06 t) + 3 \sin(2\pi \cdot 06 t),$$

$$x_{t2} = 4 \cos(2\pi \cdot 10 t) + 5 \sin(2\pi \cdot 10 t),$$

$$x_{t3} = 6 \cos(2\pi \cdot 40 t) + 7 \sin(2\pi \cdot 40 t),$$

$$x_t = x_{t1} + x_{t2} + x_{t3}.$$

What is the major difference between these series and the series generated in [Example 4.1](#)? (Hint: The answer is *fundamental*. But if your answer is the series are longer, you may be punished severely.)

- (b) As in [Example 4.2](#), compute and plot the periodogram of the series, x_t , generated in (a) and comment.
- (c) Repeat the analyses of (a) and (b) but with $n = 100$ (as in [Example 4.1](#)), and adding noise to x_t ; that is

$$x_t = x_{t1} + x_{t2} + x_{t3} + w_t$$

where $w_t \sim \text{iid } N(0, 25)$. That is, you should simulate and plot the data, and then plot the periodogram of x_t and comment.

4.2 With reference to equations (4.1) and (4.2), let $Z_1 = U_1$ and $Z_2 = -U_2$ be independent, standard normal variables. Consider the polar coordinates of the point (Z_1, Z_2) , that is,

$$A^2 = Z_1^2 + Z_2^2 \quad \text{and} \quad \phi = \tan^{-1}(Z_2/Z_1).$$

- (a) Find the joint density of A^2 and ϕ , and from the result, conclude that A^2 and ϕ are independent random variables, where A^2 is a chi-squared random variable with 2 df, and ϕ is uniformly distributed on $(-\pi, \pi)$.
- (b) Going in reverse from polar coordinates to rectangular coordinates, suppose we assume that A^2 and ϕ are independent random variables, where A^2 is chi-squared with 2 df, and ϕ is uniformly distributed on $(-\pi, \pi)$. With $Z_1 = A \cos(\phi)$ and $Z_2 = A \sin(\phi)$, where A is the positive square root of A^2 , show that Z_1 and Z_2 are independent, standard normal random variables.

4.3 Verify (4.4).

Section 4.3

4.4 A time series was generated by first drawing the white noise series w_t from a normal distribution with mean zero and variance one. The observed series x_t was generated from

$$x_t = w_t - \theta w_{t-1}, \quad t = 0, \pm 1, \pm 2, \dots,$$

where θ is a parameter.

- (a) Derive the theoretical mean value and autocovariance functions for the series x_t and w_t . Are the series x_t and w_t stationary? Give your reasons.
 (b) Give a formula for the power spectrum of x_t , expressed in terms of θ and ω .

4.5 A first-order autoregressive model is generated from the white noise series w_t using the generating equations

$$x_t = \phi x_{t-1} + w_t,$$

where ϕ , for $|\phi| < 1$, is a parameter and the w_t are independent random variables with mean zero and variance σ_w^2 .

- (a) Show that the power spectrum of x_t is given by

$$f_x(\omega) = \frac{\sigma_w^2}{1 + \phi^2 - 2\phi \cos(2\pi\omega)}.$$

- (b) Verify the autocovariance function of this process is

$$\gamma_x(h) = \frac{\sigma_w^2 \phi^{|h|}}{1 - \phi^2},$$

$h = 0, \pm 1, \pm 2, \dots$, by showing that the inverse transform of $\gamma_x(h)$ is the spectrum derived in part (a).

4.6 In applications, we will often observe series containing a signal that has been delayed by some unknown time D , i.e.,

$$x_t = s_t + A s_{t-D} + n_t,$$

where s_t and n_t are stationary and independent with zero means and spectral densities $f_s(\omega)$ and $f_n(\omega)$, respectively. The delayed signal is multiplied by some unknown constant A . Show that

$$f_x(\omega) = [1 + A^2 + 2A \cos(2\pi\omega D)] f_s(\omega) + f_n(\omega).$$

4.7 Suppose x_t and y_t are stationary zero-mean time series with x_t independent of y_s for all s and t . Consider the product series

$$z_t = x_t y_t.$$

Prove the spectral density for z_t can be written as

$$f_z(\omega) = \int_{-1/2}^{1/2} f_x(\omega - \nu) f_y(\nu) d\nu.$$




Fig. 4.22. Smoothed 12-month sunspot numbers ([sunspotz](#)) sampled twice per year.

Section 4.4

4.8 Figure 4.22 shows the biyearly smoothed (12-month moving average) number of sunspots from June 1749 to December 1978 with $n = 459$ points that were taken twice per year; the data are contained in [sunspotz](#). With Example 4.11 as a guide, perform a periodogram analysis identifying the predominant periods and obtaining confidence intervals for the identified periods. Interpret your findings.

4.9 The levels of salt concentration known to have occurred over rows, corresponding to the average temperature levels for the soil science data considered in Figure 1.17 and Figure 1.18, are in [salt](#) and [saltemp](#). Plot the series and then identify the dominant frequencies by performing separate spectral analyses on the two series. Include confidence intervals for the dominant frequencies and interpret your findings.

4.10 Let the observed series x_t be composed of a periodic signal and noise so it can be written as

$$x_t = \beta_1 \cos(2\pi\omega_k t) + \beta_2 \sin(2\pi\omega_k t) + w_t,$$

where w_t is a white noise process with variance σ_w^2 . The frequency ω_k is assumed to be known and of the form k/n in this problem. Suppose we consider estimating β_1 , β_2 and σ_w^2 by least squares, or equivalently, by maximum likelihood if the w_t are assumed to be Gaussian.

(a) Prove, for a fixed ω_k , the minimum squared error is attained by

$$\begin{pmatrix} \hat{\beta}_1 \\ \hat{\beta}_2 \end{pmatrix} = 2n^{-1/2} \begin{pmatrix} d_c(\omega_k) \\ d_s(\omega_k) \end{pmatrix},$$

where the cosine and sine transforms (4.26) and (4.27) appear on the right-hand side.

(b) Prove that the error sum of squares can be written as

$$SSE = \sum_{t=1}^n x_t^2 - 2I_x(\omega_k)$$

so that the value of ω_k that minimizes squared error is the same as the value that maximizes the periodogram $I_x(\omega_k)$ estimator (4.23).

(c) Under the Gaussian assumption and fixed ω_k , show that the F -test of no regression leads to an F -statistic that is a monotone function of $I_x(\omega_k)$.

4.11 Prove the convolution property of the DFT, namely,

$$\sum_{s=1}^n a_s x_{t-s} = \sum_{k=0}^{n-1} d_A(\omega_k) d_x(\omega_k) \exp\{2\pi\omega_k t\},$$

for $t = 1, 2, \dots, n$, where $d_A(\omega_k)$ and $d_x(\omega_k)$ are the discrete Fourier transforms of a_t and x_t , respectively, and we assume that $x_t = x_{t+n}$ is periodic.

Section 4.5

4.12 Repeat **Problem 4.8** using a nonparametric spectral estimation procedure. In addition to discussing your findings in detail, comment on your choice of a spectral estimate with regard to smoothing and tapering.

4.13 Repeat **Problem 4.9** using a nonparametric spectral estimation procedure. In addition to discussing your findings in detail, comment on your choice of a spectral estimate with regard to smoothing and tapering.

4.14 Cepstral Analysis. The periodic behavior of a time series induced by echoes can also be observed in the spectrum of the series; this fact can be seen from the results stated in **Problem 4.6**. Using the notation of that problem, suppose we observe $x_t = s_t + As_{t-D} + n_t$, which implies the spectra satisfy $f_x(\omega) = [1 + A^2 + 2A \cos(2\pi\omega D)]f_s(\omega) + f_n(\omega)$. If the noise is negligible ($f_n(\omega) \approx 0$) then $\log f_x(\omega)$ is approximately the sum of a periodic component, $\log[1 + A^2 + 2A \cos(2\pi\omega D)]$, and $\log f_s(\omega)$. Bogart et al. (1962) proposed treating the detrended log spectrum as a pseudo time series and calculating its spectrum, or *cepstrum*, which should show a peak at a *quefrency* corresponding to $1/D$. The cepstrum can be plotted as a function of quefrency, from which the delay D can be estimated.

For the speech series presented in **Example 1.3**, estimate the pitch period using cepstral analysis as follows. The data are in [speech](#).

- Calculate and display the log-periodogram of the data. Is the periodogram periodic, as predicted?
- Perform a cepstral (spectral) analysis on the detrended logged periodogram, and use the results to estimate the delay D . How does your answer compare with the analysis of **Example 1.25**, which was based on the ACF?

4.15 Use [Property 4.2](#) to verify (4.66). Then verify (4.69) and (4.70).

4.16 Consider two time series

$$x_t = w_t - w_{t-1},$$

$$y_t = \frac{1}{2}(w_t + w_{t-1}),$$

formed from the white noise series w_t with variance $\sigma_w^2 = 1$.

- (a) Are x_t and y_t jointly stationary? Recall the cross-covariance function must also be a function only of the lag h and cannot depend on time.
- (b) Compute the spectra $f_y(\omega)$ and $f_x(\omega)$, and comment on the difference between the two results.
- (c) Suppose sample spectral estimators $\tilde{f}_y(.10)$ are computed for the series using $L = 3$. Find a and b such that

$$P\left\{a \leq \tilde{f}_y(.10) \leq b\right\} = .90.$$

This expression gives two points that will contain 90% of the sample spectral values. Put 5% of the area in each tail.

Section 4.6

4.17 Often, the periodicities in the sunspot series are investigated by fitting an autoregressive spectrum of sufficiently high order. The main periodicity is often stated to be in the neighborhood of 11 years. Fit an autoregressive spectral estimator to the sunspot data using a model selection method of your choice. Compare the result with a conventional nonparametric spectral estimator found in [Problem 4.8](#).

4.18 Fit an autoregressive spectral estimator to the Recruitment series and compare it to the results of [Example 4.14](#).

4.19 Suppose a sample time series with $n = 256$ points is available from the first-order autoregressive model. Furthermore, suppose a sample spectrum computed with $L = 3$ yields the estimated value $\tilde{f}_x(1/8) = 2.25$. Is this sample value consistent with $\sigma_w^2 = 1, \phi = .5$? Repeat using $L = 11$ if we just happen to obtain the same sample value.

4.20 Suppose we wish to test the noise alone hypothesis $H_0 : x_t = n_t$ against the signal-plus-noise hypothesis $H_1 : x_t = s_t + n_t$, where s_t and n_t are uncorrelated zero-mean stationary processes with spectra $f_s(\omega)$ and $f_n(\omega)$. Suppose that we want the test over a band of $L = 2m + 1$ frequencies of the form $\omega_{j:n} + k/n$, for $k = 0, \pm 1, \pm 2, \dots, \pm m$ near some fixed frequency ω . Assume that both the signal and noise spectra are approximately constant over the interval.

- (a) Prove the approximate likelihood-based test statistic for testing H_0 against H_1 is proportional to

$$T = \sum_k |d_x(\omega_{j:n} + k/n)|^2 \left(\frac{1}{f_n(\omega)} - \frac{1}{f_s(\omega) + f_n(\omega)} \right).$$

- (b) Find the approximate distributions of T under H_0 and H_1 .
 (c) Define the false alarm and signal detection probabilities as $P_F = P\{T > K|H_0\}$ and $P_d = P\{T > k|H_1\}$, respectively. Express these probabilities in terms of the signal-to-noise ratio $f_s(\omega)/f_n(\omega)$ and appropriate chi-squared integrals.

Section 4.7

4.21 Analyze the coherency between the temperature and salt data discussed in **Problem 4.9**. Discuss your findings.

4.22 Consider two processes

$$x_t = w_t \quad \text{and} \quad y_t = \phi x_{t-D} + v_t$$

where w_t and v_t are independent white noise processes with common variance σ^2 , ϕ is a constant, and D is a fixed integer delay.

- (a) Compute the coherency between x_t and y_t .
 (b) Simulate $n = 1024$ normal observations from x_t and y_t for $\phi = .9$, $\sigma^2 = 1$, and $D = 0$. Then estimate and plot the coherency between the simulated series for the following values of L and comment:
 (i) $L = 1$, (ii) $L = 3$, (iii) $L = 41$, and (iv) $L = 101$.

Section 4.8

4.23 For the processes in **Problem 4.22**:

- (a) Compute the phase between x_t and y_t .
 (b) Simulate $n = 1024$ observations from x_t and y_t for $\phi = .9$, $\sigma^2 = 1$, and $D = 1$. Then estimate and plot the phase between the simulated series for the following values of L and comment:
 (i) $L = 1$, (ii) $L = 3$, (iii) $L = 41$, and (iv) $L = 101$.

4.24 Consider the bivariate time series records containing monthly U.S. production (**prod**) as measured by the Federal Reserve Board Production Index and the monthly unemployment series (**unemp**).

- (a) Compute the spectrum and the log spectrum for each series, and identify statistically significant peaks. Explain what might be generating the peaks. Compute the coherence, and explain what is meant when a high coherence is observed at a particular frequency.

(b) What would be the effect of applying the filter

$$u_t = x_t - x_{t-1} \quad \text{followed by} \quad v_t = u_t - u_{t-12}$$

to the series given above? Plot the predicted frequency responses of the simple difference filter and of the seasonal difference of the first difference.

(c) Apply the filters successively to one of the two series and plot the output. Examine the output after taking a first difference and comment on whether stationarity is a reasonable assumption. Why or why not? Plot after taking the seasonal difference of the first difference. What can be noticed about the output that is consistent with what you have predicted from the frequency response? Verify by computing the spectrum of the output after filtering.

4.25 Determine the theoretical power spectrum of the series formed by combining the white noise series w_t to form

$$y_t = w_{t-2} + 4w_{t-1} + 6w_t + 4w_{t+1} + w_{t+2}.$$

Determine which frequencies are present by plotting the power spectrum.

4.26 Let $x_t = \cos(2\pi\omega t)$, and consider the output

$$y_t = \sum_{k=-\infty}^{\infty} a_k x_{t-k},$$

where $\sum_k |a_k| < \infty$. Show

$$y_t = |A(\omega)| \cos(2\pi\omega t + \phi(\omega)),$$

where $|A(\omega)|$ and $\phi(\omega)$ are the amplitude and phase of the filter, respectively. Interpret the result in terms of the relationship between the input series, x_t , and the output series, y_t .

4.27 Suppose x_t is a stationary series, and we apply two filtering operations in succession, say,

$$y_t = \sum_r a_r x_{t-r} \quad \text{then} \quad z_t = \sum_s b_s y_{t-s}.$$

(a) Show the spectrum of the output is

$$f_z(\omega) = |A(\omega)|^2 |B(\omega)|^2 f_x(\omega),$$

where $A(\omega)$ and $B(\omega)$ are the Fourier transforms of the filter sequences a_t and b_t , respectively.

(b) What would be the effect of applying the filter

$$u_t = x_t - x_{t-1} \quad \text{followed by} \quad v_t = u_t - u_{t-12}$$

to a time series?

- (c) Plot the predicted frequency responses of the simple difference filter and of the seasonal difference of the first difference. Filters like these are called seasonal adjustment filters in economics because they tend to attenuate frequencies at multiples of the monthly periods. The difference filter tends to attenuate low-frequency trends.

4.28 Suppose we are given a stationary zero-mean series x_t with spectrum $f_x(\omega)$ and then construct the derived series

$$y_t = ay_{t-1} + x_t, \quad t = \pm 1, \pm 2, \dots$$

- (a) Show how the theoretical $f_y(\omega)$ is related to $f_x(\omega)$.
 (b) Plot the function that multiplies $f_x(\omega)$ in part (a) for $a = .1$ and for $a = .8$. This filter is called a recursive filter.

Section 4.9

4.29 Consider the problem of approximating the filter output

$$y_t = \sum_{k=-\infty}^{\infty} a_k x_{t-k}, \quad \sum_{-\infty}^{\infty} |a_k| < \infty,$$

by

$$y_t^M = \sum_{|k| < M/2} a_k^M x_{t-k}$$

for $t = M/2 - 1, M/2, \dots, n - M/2$, where x_t is available for $t = 1, \dots, n$ and

$$a_t^M = M^{-1} \sum_{k=0}^{M-1} A(\omega_k) \exp\{2\pi i \omega_k t\}$$

with $\omega_k = k/M$. Prove

$$E\{(y_t - y_t^M)^2\} \leq 4\gamma_x(0) \left(\sum_{|k| \geq M/2} |a_k| \right)^2.$$

4.30 Prove the squared coherence $\rho_{y,x}^2(\omega) = 1$ for all ω when

$$y_t = \sum_{r=-\infty}^{\infty} a_r x_{t-r},$$

that is, when x_t and y_t can be related exactly by a linear filter.

4.31 The data set [climhyd](#), contains 454 months of measured values for six climatic variables: (i) air temperature [[Temp](#)], (ii) dew point [[DewPt](#)], (iii) cloud cover [[CldCvr](#)], (iv) wind speed [[WndSpd](#)], (v) precipitation [[Precip](#)], and (vi) inflow [[Inflow](#)], at Lake Shasta in California; the data are displayed in [Figure 7.3](#). We would like to look at possible relations among the weather factors and between the weather factors and the inflow to Lake Shasta.

- (a) First transform the inflow and precipitation series as follows: $I_t = \log i_t$, where i_t is inflow, and $P_t = \sqrt{p_t}$, where p_t is precipitation. Then, compute the square coherencies between all the weather variables and transformed inflow and argue that the strongest determinant of the inflow series is (transformed) precipitation. [Tip: If $\mathbf{x}$ contains multiple time series, then the easiest way to display all the squared coherencies is to first make an object of class `spec`; e.g., `u = spectrum(x, span=c(7,7), plot=FALSE)` and then plot the coherencies suppressing the confidence intervals, `plot(u, ci=-1, plot.type="coh").`]
- (b) Fit a lagged regression model of the form

$$I_t = \beta_0 + \sum_{j=0}^{\infty} \beta_j P_{t-j} + w_t,$$

using thresholding, and then comment of the predictive ability of precipitation for inflow.

Section 4.10

4.32 Consider the *signal plus noise* model

$$y_t = \sum_{r=-\infty}^{\infty} \beta_r x_{t-r} + v_t,$$

where the signal and noise series, x_t and v_t are both stationary with spectra $f_x(\omega)$ and $f_v(\omega)$, respectively. Assuming that x_t and v_t are independent of each other for all t , verify (4.127) and (4.128).

4.33 Consider the model

$$y_t = x_t + v_t,$$

where

$$x_t = \phi x_{t-1} + w_t,$$

such that v_t is Gaussian white noise and independent of x_t with $\text{var}(v_t) = \sigma_v^2$, and w_t is Gaussian white noise and independent of v_t , with $\text{var}(w_t) = \sigma_w^2$, and $|\phi| < 1$ and $Ex_0 = 0$. Prove that the spectrum of the observed series y_t is

$$f_y(\omega) = \sigma^2 \frac{|1 - \theta e^{-2\pi i \omega}|^2}{|1 - \phi e^{-2\pi i \omega}|^2},$$

where

$$\theta = \frac{c \pm \sqrt{c^2 - 4}}{2}, \quad \sigma^2 = \frac{\sigma_v^2 \phi}{\theta},$$

and

$$c = \frac{\sigma_w^2 + \sigma_v^2(1 + \phi^2)}{\sigma_v^2 \phi}.$$

4.34 Consider the same model as in the preceding problem.

(a) Prove the optimal smoothed estimator of the form

$$\hat{x}_t = \sum_{s=-\infty}^{\infty} a_s y_{t-s}$$

has

$$a_s = \frac{\sigma_w^2}{\sigma^2} \frac{\theta^{|s|}}{1 - \theta^2}.$$

(b) Show the mean square error is given by

$$E\{(x_t - \hat{x}_t)^2\} = \frac{\sigma_v^2 \sigma_w^2}{\sigma^2(1 - \theta^2)}.$$

(c) Compare mean square error of the estimator in part (b) with that of the optimal finite estimator of the form

$$\hat{x}_t = a_1 y_{t-1} + a_2 y_{t-2}$$

when $\sigma_v^2 = .053$, $\sigma_w^2 = .172$, and $\phi_1 = .9$.

Section 4.11

4.35 Consider the two-dimensional linear filter given as the output (4.139).

- (a) Express the two-dimensional autocovariance function of the output, say, $\gamma_y(h_1, h_2)$, in terms of an infinite sum involving the autocovariance function of x_s and the filter coefficients a_{s_1, s_2} .
- (b) Use the expression derived in (a), combined with (4.138) and (4.141) to derive the spectrum of the filtered output (4.140).

The following problems require supplemental material from Appendix C

4.36 Let w_t be a Gaussian white noise series with variance σ_w^2 . Prove that the results of Theorem C.4 hold without error for the DFT of w_t .

4.37 Show that condition (4.43) implies (C.19) by showing

$$n^{-1/2} \sum_{h \geq 0} h |\gamma(h)| \leq \sigma_w^2 \sum_{k \geq 0} |\psi_k| \sum_{j \geq 0} \sqrt{j} |\psi_j|.$$

4.38 Prove Lemma C.4.

4.39 Finish the proof of Theorem C.5.

4.40 For the zero-mean complex random vector $\mathbf{z} = \mathbf{x}_c - i\mathbf{x}_s$, with $\text{cov}(\mathbf{z}) = \Sigma = \mathbf{C} - i\mathbf{Q}$, with $\Sigma = \Sigma^*$, define

$$w = 2\text{Re}(\mathbf{a}^* \mathbf{z}),$$

where $\mathbf{a} = \mathbf{a}_c - i\mathbf{a}_s$ is an arbitrary non-zero complex vector. Prove

$$\text{cov}(w) = 2\mathbf{a}^* \Sigma \mathbf{a}.$$

Recall $*$ denotes the complex conjugate transpose.

Chapter 5

Additional Time Domain Topics

5.1 Introduction

In this chapter, we present material that may be considered special or advanced topics in the time domain. Chapter 6 is devoted to one of the most useful and interesting time domain topics, state-space models. Consequently, we do not cover state-space models or related topics|of which there are many|in this chapter. This chapter consists of sections of independent topics that may be read in any order. Most of the sections depend on a basic knowledge of ARMA models, forecasting and estimation, which is the material that is covered in [Chapter 3](#). A few sections, for example the section on long memory models, require some knowledge of spectral analysis and related topics covered in [Chapter 4](#). In addition to long memory, we discuss unit root testing, GARCH models, threshold models, regression with autocorrelated errors, lagged regression or transfer functions, and selected topics in multivariate ARMAX models.

5.2 Long Memory ARMA and Fractional Differencing

The conventional ARMA(p, q) process is often referred to as a short-memory process because the coefficients in the representation

$$x_t = \sum_{j=0}^{\infty} \psi_j w_{t-j},$$

obtained by solving

$$\phi(z)\psi(z) = \theta(z),$$

are dominated by exponential decay. As pointed out in [§3.3](#), this result implies the ACF of the short memory process $\rho(h) \rightarrow 0$ exponentially fast as $h \rightarrow \infty$. When the sample ACF of a time series decays slowly, the advice given in Chapter 3 has been to difference the series until it seems stationary. Following this advice with the glacial varve series first presented in [Example 3.32](#) leads to the first difference

of the logarithms of the data being represented as a first-order moving average. In **Example 3.40**, further analysis of the residuals leads to fitting an ARIMA(1, 1, 1) model,

$$\nabla x_t = \phi \nabla x_{t-1} + w_t + \theta w_{t-1},$$

where we understand x_t is the log-transformed varve series. In particular, the estimates of the parameters (and the standard errors) were $\hat{\phi} = .23(.05)$, $\hat{\theta} = -.89(.03)$, and $\hat{\sigma}_w^2 = .23$. The use of the first difference $\nabla x_t = (1 - B)x_t$ can be too severe a modification in the sense that the nonstationary model might represent an overdifferencing of the original process.

Long memory (or persistent) time series were considered in Hosking (1981) and Granger and Joyeux (1980) as intermediate compromises between the short memory ARMA type models and the fully integrated nonstationary processes in the Box–Jenkins class. The easiest way to generate a long memory series is to think of using the difference operator $(1 - B)^d$ for fractional values of d , say, $0 < d < .5$, so a basic long memory series gets generated as

$$(1 - B)^d x_t = w_t, \quad (5.1)$$

where w_t still denotes white noise with variance σ_w^2 . The fractionally differenced series (5.1), for $|d| < .5$, is often called *fractional noise* (except when d is zero). Now, d becomes a parameter to be estimated along with σ_w^2 . Differencing the original process, as in the Box–Jenkins approach, may be thought of as simply assigning a value of $d = 1$. This idea has been extended to the class of fractionally integrated ARMA, or ARFIMA models, where $-.5 < d < .5$; when d is negative, the term antipersistent is used. Long memory processes occur in hydrology (see Hurst, 1951, and McLeod and Hipel, 1978) and in environmental series, such as the varve data we have previously analyzed, to mention a few examples. Long memory time series data tend to exhibit sample autocorrelations that are not necessarily large (as in the case of $d = 1$), but persist for a long time. **Figure 5.1** shows the sample ACF, to lag 100, of the log-transformed varve series, which exhibits classic long memory behavior:

```
u = acf(log(varve), 100, plot=FALSE)
plot(u[1:100], ylim=c(-.1,1), main="log(varve)") # get rid of lag 0
```

To investigate its properties, we can use the binomial expansion ($d > -1$) to write

$$w_t = (1 - B)^d x_t = \sum_{j=0}^{\infty} \pi_j B^j x_t = \sum_{j=0}^{\infty} \pi_j x_{t-j} \quad (5.2)$$

where

$$\pi_j = \frac{\Gamma(j - d)}{\Gamma(j + 1)\Gamma(-d)} \quad (5.3)$$

with $\Gamma(x + 1) = x\Gamma(x)$ being the gamma function. Similarly ($d < 1$), we can write

$$x_t = (1 - B)^{-d} w_t = \sum_{j=0}^{\infty} \psi_j B^j w_t = \sum_{j=0}^{\infty} \psi_j w_{t-j} \quad (5.4)$$




Fig. 5.1. Sample ACF of the log transformed varve series.

where

$$\psi_j = \frac{\Gamma(j+d)}{\Gamma(j+1)\Gamma(d)}. \quad (5.5)$$

When $|d| < .5$, the processes (5.2) and (5.4) are well-defined stationary processes (see Brockwell and Davis, 1991, for details). In the case of fractional differencing, however, the coefficients satisfy $\sum \pi_j^2 < \infty$ and $\sum \psi_j^2 < \infty$ as opposed to the absolute summability of the coefficients in ARMA processes.

Using the representation (5.4)–(5.5), and after some nontrivial manipulations, it can be shown that the ACF of x_t is

$$\rho(h) = \frac{\Gamma(h+d)\Gamma(1-d)}{\Gamma(h-d+1)\Gamma(d)} \sim h^{2d-1} \quad (5.6)$$

for large h . From this we see that for $0 < d < .5$

$$\sum_{h=-\infty}^{\infty} |\rho(h)| = \infty$$

and hence the term *long memory*.

In order to examine a series such as the varve series for a possible long memory pattern, it is convenient to look at ways of estimating d . Using (5.3) it is easy to derive the recursions

$$\pi_{j+1}(d) = \frac{(j-d)\pi_j(d)}{(j+1)}, \quad (5.7)$$

for $j = 0, 1, \dots$, with $\pi_0(d) = 1$. Maximizing the joint likelihood of the errors under normality, say, $w_t(d)$, will involve minimizing the sum of squared errors

$$Q(d) = \sum w_t^2(d).$$




Fig. 5.2. Coefficients $\pi_j(.384)$, $j = 1, 2, \dots, 30$ in the representation (5.7).

The usual Gauss–Newton method, described in §3.6, leads to the expansion

$$w_t(d) = w_t(d_0) + w'_t(d_0)(d - d_0),$$

where

$$w'_t(d_0) = \left. \frac{\partial w_t}{\partial d} \right|_{d=d_0}$$

and d_0 is an initial estimate (guess) at to the value of d . Setting up the usual regression leads to

$$d = d_0 - \frac{\sum_t w'_t(d_0)w_t(d_0)}{\sum_t w'_t(d_0)^2}. \quad (5.8)$$

The derivatives are computed recursively by differentiating (5.7) successively with respect to d :

$$\pi'_{j+1}(d) = \frac{(j-d)\pi'_j(d) - \pi_j(d)}{j+1},$$

where $\pi'_0(d) = 0$. The errors are computed from an approximation to (5.2), namely,

$$w_t(d) = \sum_{j=0}^t \pi_j(d)x_{t-j}. \quad (5.9)$$

It is advisable to omit a number of initial terms from the computation and start the sum, (5.8), at some fairly large value of t to have a reasonable approximation.

Example 5.1 Long Memory Fitting of the Glacial Varve Series

We consider analyzing the glacial varve series discussed in [Example 2.6](#) and [Example 3.32](#). [Figure 2.6](#) shows the original and log-transformed series (which we denote by x_t). In [Example 3.40](#), we noted that x_t could be modeled as an ARIMA(1, 1, 1) process. We fit the fractionally differenced model, (5.1), to the mean-adjusted series, $x_t - \bar{x}$. Applying the Gauss–Newton iterative procedure previously described, starting with $d = .1$ and omitting the first 30 points from the computation, leads to

a final value of $d = .384$, which implies the set of coefficients $\pi_j(.384)$, as given in [Figure 5.2](#) with $\pi_0(.384) = 1$. We can compare roughly the performance of the fractional difference operator with the ARIMA model by examining the autocorrelation functions of the two residual series as shown in [Figure 5.3](#). The ACFs of the two residual series are roughly comparable with the white noise model.

To perform this analysis in R, first download and install the `fracdiff` package. Then use

```
library(fracdiff)
lvarve = log(varve)-mean(log(varve))
varve.fd = fracdiff(lvarve, nar=0, nma=0, M=30)
varve.fd$d # = 0.3841688
varve.fd$stderror.dpq # = 4.589514e-06 (questionable result!!)
p = rep(1,31)
for (k in 1:30){ p[k+1] = (k-varve.fd$d)*p[k]/(k+1) }
plot(1:30, p[-1], ylab=expression(pi(d)), xlab="Index", type="h")
res.fd = diffseries(log(varve), varve.fd$d) # frac diff resids
res.arima = resid(arima(log(varve), order=c(1,1,1))) # arima resids
par(mfrow=c(2,1))
acf(res.arima, 100, xlim=c(4,97), ylim=c(-.2,.2), main="")
acf(res.fd, 100, xlim=c(4,97), ylim=c(-.2,.2), main="")
```

The R package uses a truncated maximum likelihood procedure that was discussed in Haslett and Raftery (1989), which is a little more elaborate than simply zeroing out initial values. The default truncation value in R is $M = 100$. In the default case, the estimate is $\hat{d} = .37$ with approximately the same (questionable) standard error. The standard error is (supposedly) obtained from the Hessian as described in [Example 3.29](#).

Forecasting long memory processes is similar to forecasting ARIMA models. That is, (5.2) and (5.7) can be used to obtain the truncated forecasts

$$\tilde{x}_{n+m}^n = - \sum_{j=1}^n \pi_j(\hat{d}) \tilde{x}_{n+m-j}^n, \quad (5.10)$$

for $m = 1, 2, \dots$. Error bounds can be approximated by using

$$P_{n+m}^n = \hat{\sigma}_w^2 \left(\sum_{j=0}^{m-1} \psi_j^2(\hat{d}) \right) \quad (5.11)$$

where, as in (5.7),

$$\psi_j(\hat{d}) = \frac{(j + \hat{d})\psi_j(\hat{d})}{(j + 1)}, \quad (5.12)$$

with $\psi_0(\hat{d}) = 1$.

No obvious short memory ARMA-type component can be seen in the ACF of the residuals from the fractionally differenced varve series shown in [Figure 5.3](#). It is natural, however, that cases will exist in which substantial short memory-type components will also be present in data that exhibits long memory. Hence, it is natural to define the general ARFIMA(p, d, q), $-.5 < d < .5$ process as




Fig. 5.3. ACF of residuals from the ARIMA(1, 1, 1) fit to the logged varve series (top) and of the residuals from the long memory model fit, $(1 - B)^d x_t = w_t$, with $d = .384$ (bottom).

$$\phi(B)\nabla^d(x_t - \mu) = \theta(B)w_t, \quad (5.13)$$

where $\phi(B)$ and $\theta(B)$ are as given in Chapter 3. Writing the model in the form

$$\phi(B)\pi_d(B)(x_t - \mu) = \theta(B)w_t \quad (5.14)$$

makes it clear how we go about estimating the parameters for the more general model. Forecasting for the ARFIMA(p, d, q) series can be easily done, noting that we may equate coefficients in

$$\phi(z)\psi(z) = (1 - z)^{-d}\theta(z) \quad (5.15)$$

and

$$\theta(z)\pi(z) = (1 - z)^d\phi(z) \quad (5.16)$$

to obtain the representations

$$x_t = \mu + \sum_{j=0}^{\infty} \psi_j w_{t-j}$$

and

$$w_t = \sum_{j=0}^{\infty} \pi_j (x_{t-j} - \mu).$$

We then can proceed as discussed in (5.10) and (5.11).

Comprehensive treatments of long memory time series models are given in the texts by Beran (1994), Palma (2007), and Robinson (2003), and it should be noted that

several other techniques for estimating the parameters, especially, the long memory parameter, can be developed in the frequency domain. In this case, we may think of the equations as generated by an infinite order autoregressive series with coefficients π_j given by (5.7). Using the same approach as before, we obtain

$$\begin{aligned} f_x(\omega) &= \frac{\sigma_w^2}{|\sum_{k=0}^{\infty} \pi_k e^{-2\pi i k \omega}|^2} \\ &= \sigma_w^2 |1 - e^{-2\pi i \omega}|^{-2d} = [4 \sin^2(\pi \omega)]^{-d} \sigma_w^2 \end{aligned} \quad (5.17)$$

as equivalent representations of the spectrum of a long memory process. The long memory spectrum approaches infinity as the frequency $\omega \rightarrow 0$.

The main reason for defining the Whittle approximation to the log likelihood is to propose its use for estimating the parameter d in the long memory case as an alternative to the time domain method previously mentioned. The time domain approach is useful because of its simplicity and easily computed standard errors. One may also use an exact likelihood approach by developing an innovations form of the likelihood as in Brockwell and Davis (1991).

For the approximate approach using the Whittle likelihood (4.81), we consider using the approach of Fox and Taquq (1986) who showed that maximizing the Whittle log likelihood leads to a consistent estimator with the usual asymptotic normal distribution that would be obtained by treating (4.81) as a conventional log likelihood (see also Dahlhaus, 1989; Robinson, 1995; Hurvich et al., 1998). Unfortunately, the periodogram ordinates are not asymptotically independent (Hurvich and Beltrao, 1993), although a quasi-likelihood in the form of the Whittle approximation works well and has good asymptotic properties.

To see how this would work for the purely long memory case, write the long memory spectrum as

$$f_x(\omega_k; d, \sigma_w^2) = \sigma_w^2 g_k^{-d}, \quad (5.18)$$

where

$$g_k = 4 \sin^2(\pi \omega_k). \quad (5.19)$$

Then, differentiating the log likelihood, say,

$$\ln L(\mathbf{x}; d, \sigma_w^2) \approx -m \ln \sigma_w^2 + d \sum_{k=1}^m \ln g_k - \frac{1}{\sigma_w^2} \sum_{k=1}^m g_k^d I(\omega_k) \quad (5.20)$$

at $m = n/2 - 1$ frequencies and solving for σ_w^2 yields

$$\sigma_w^2(d) = \frac{1}{m} \sum_{k=1}^m g_k^d I(\omega_k) \quad (5.21)$$

as the approximate maximum likelihood estimator for the variance parameter. To estimate d , we can use a grid search of the concentrated log likelihood

$$\ln L(\mathbf{x}; d) \approx -m \ln \sigma_w^2(d) + d \sum_{k=1}^m \ln g_k - m \quad (5.22)$$




Fig. 5.4. Long Memory ($d = .380$) [solid line] and autoregressive AR(8) [dashed line] spectral estimators for the paleoclimatic glacial varve series.

over the interval $(-.5, .5)$, followed by a Newton–Raphson procedure to convergence.

Example 5.2 Long Memory Spectra for the Varve Series

In [Example 5.1](#), we fit a long memory model to the glacial varve data via time domain methods. Fitting the same model using frequency domain methods and the Whittle approximation above gives $\hat{d} = .380$, with an estimated standard error of .028. The earlier time domain method gave $\hat{d} = .384$ with $M = 30$ and $\hat{d} = .370$ with $M = 100$. Both estimates obtained via time domain methods had a standard error of about 4.6×10^{-6} , which seems implausible. The error variance estimate in this case is $\hat{\sigma}_w^2 = .2293$; in [Example 5.1](#), we could have used `var(res.fd)` as an estimate, in which case we obtain .2298. The R code to perform this analysis is

```
series = log(varve) # specify series to be analyzed
d0 = .1 # initial value of d
n.per = nextn(length(series))
m = (n.per)/2 - 1
per = Mod(fft(series-mean(series))[-1])^2 # remove 0 freq
per = per/n.per # and scale the periodogram
g = 4*(sin(pi*((1:m)/n.per))^2)
# Function to calculate -log.likelihood
whit.like = function(d){
  g.d=g*d
  sig2 = (sum(g.d*per[1:m])/m)
  log.like = m*log(sig2) - d*sum(log(g)) + m
  return(log.like) }
# Estimation (?optim for details - output not shown)
(est = optim(d0, whit.like, gr=NULL, method="L-BFGS-B", hessian=TRUE,
  lower=-.5, upper=.5, control=list(trace=1,REPORT=1)))
# Results: d.hat = .380, se(dhat) = .028, and sig2hat = .229
```

```
cat("d.hat =", est$par, "se(dhat) = ", 1/sqrt(est$hessian), "\n")
g.dhat = g^est$par; sig2 = sum(g.dhat*per[1:m])/m
cat("sig2hat =", sig2, "\n")
```

One might also consider fitting an autoregressive model to these data using a procedure similar to that used in [Example 4.16](#). Following this approach gave an autoregressive model with $p = 8$ and $\hat{\phi} = (.34, .11, .04, .09, .08, .08, .02, .09)'$, with $\hat{\sigma}_w^2 = .2267$ as the error variance. The two log spectra are plotted in [Figure 5.4](#) for $\omega > 0$, and we note that long memory spectrum will eventually become infinite, whereas the AR(8) spectrum is finite at $\omega = 0$. The R code used for this part of the example (assuming the previous values have been retained) is

```
u = spec.ar(log(varve), plot=FALSE) # produces AR(8)
g = 4*(sin(pi*((1:500)/2000))^2)
fhat = sig2*g^{-est$par} # long memory spectral estimate
plot(1:500/2000, log(fhat), type="l", ylab="log(spectrum)", xlab="frequency")
lines(u$freq[1:250], log(u$spec[1:250]), lty="dashed")
ar.mle(log(varve)) # to get AR(8) estimates
```

Often, time series are not purely long memory. A common situation has the long memory component multiplied by a short memory component, leading to an alternate version of (5.18) of the form

$$f_x(\omega_k; d, \theta) = g_k^{-d} f_0(\omega_k; \theta), \quad (5.23)$$

where $f_0(\omega_k; \theta)$ might be the spectrum of an autoregressive moving average process with vector parameter θ , or it might be unspecified. If the spectrum has a parametric form, the Whittle likelihood can be used. However, there is a substantial amount of semiparametric literature that develops the estimators when the underlying spectrum $f_0(\omega; \theta)$ is unknown. A class of Gaussian semi-parametric estimators simply uses the same Whittle likelihood (5.22), evaluated over a sub-band of low frequencies, say $m' = \sqrt{n}$. There is some latitude in selecting a band that is relatively free from low frequency interference due to the short memory component in (5.23).

Geweke and Porter–Hudak (1983) developed an approximate method for estimating d based on a regression model, derived from (5.22). Note that we may write a simple equation for the logarithm of the spectrum as

$$\ln f_x(\omega_k; d) = \ln f_0(\omega_k; \theta) - d \ln[4 \sin^2(\pi\omega_k)], \quad (5.24)$$

with the frequencies $\omega_k = k/n$ restricted to a range $k = 1, 2, \dots, m'$ near the zero frequency with $m' = \sqrt{n}$ as the recommended value. Relationship (5.24) suggests using a simple linear regression model of the form,

$$\ln I(\omega_k) = \beta_0 - d \ln[4 \sin^2(\pi\omega_k)] + e_k \quad (5.25)$$

for the periodogram to estimate the parameters σ_w^2 and d . In this case, one performs least squares using $\ln I(\omega_k)$ as the dependent variable, and $\ln[4 \sin^2(\pi\omega_k)]$ as the independent variable for $k = 1, 2, \dots, m$. The resulting slope estimate is then used as an estimate of $-d$. For a good discussion of various alternative methods for selecting

m , see Hurvich and Deo (1999). The R package `fracdiff` also provides this method via the command `fdGPH()`; see the help file for further information.

One of the above two procedures works well for estimating the long memory component but there will be cases (such as ARFIMA) where there will be a parameterized short memory component $f_0(\omega_k; \theta)$ that needs to be estimated. If the spectrum is highly parameterized, one might estimate using the Whittle log likelihood (5.19) and

$$f_x(\omega_k; \theta) = g_k^{-d} f_0(\omega_k; \theta)$$

and jointly estimating the parameters d and θ using the Newton–Raphson method. If we are interested in a nonparametric estimator, using the conventional smoothed spectral estimator for the periodogram, adjusted for the long memory component, say $g_k^d I(\omega_k)$ might be a possible approach.

5.3 Unit Root Testing

As discussed in the previous section, the use of the first difference $\nabla x_t = (1 - B)x_t$ can be too severe a modification in the sense that the nonstationary model might represent an overdifferencing of the original process. For example, consider a causal AR(1) process (we assume throughout this section that the noise is Gaussian),

$$x_t = \phi x_{t-1} + w_t. \quad (5.26)$$

Applying $(1 - B)$ to both sides shows that differencing, $\nabla x_t = \phi \nabla x_{t-1} + \nabla w_t$, or

$$y_t = \phi y_{t-1} + w_t - w_{t-1},$$

where $y_t = \nabla x_t$, introduces extraneous correlation and invertibility problems. That is, while x_t is a causal AR(1) process, working with the differenced process y_t will be problematic because it is a non-invertible ARMA(1, 1).

A unit root test provides a way to test whether (5.26) is a random walk (the null case) as opposed to a causal process (the alternative). That is, it provides a procedure for testing

$$H_0: \phi = 1 \quad \text{versus} \quad H_1: |\phi| < 1.$$

An obvious test statistic would be to consider $(\hat{\phi} - 1)$, appropriately normalized, in the hope to develop an asymptotically normal test statistic, where $\hat{\phi}$ is one of the optimal estimators discussed in Chapter 3, §3.6. Unfortunately, the theory of §3.6 will not work in the null case because the process is nonstationary. Moreover, as seen in Example 3.35, estimation near the boundary of stationarity produces highly skewed sample distributions (see Figure 3.12) and this is a good indication that the problem will be atypical.

To examine the behavior of $(\hat{\phi} - 1)$ under the null hypothesis that $\phi = 1$, or more precisely that the model is a random walk, $x_t = \sum_{j=1}^t w_j$, or $x_t = x_{t-1} + w_t$ with $x_0 = 0$, consider the least squares estimator of ϕ . Noting that $\mu_x = 0$, the least squares estimator can be written as

$$\hat{\phi} = \frac{\frac{1}{n} \sum_{t=1}^n x_t x_{t-1}}{\frac{1}{n} \sum_{t=1}^n x_{t-1}^2} = 1 + \frac{\frac{1}{n} \sum_{t=1}^n w_t x_{t-1}}{\frac{1}{n} \sum_{t=1}^n x_{t-1}^2}, \quad (5.27)$$

where we have written $x_t = x_{t-1} + w_t$ in the numerator; recall that $x_0 = 0$ and in the least squares setting, we are regressing x_t on x_{t-1} for $t = 1, \dots, n$. Hence, under H_0 , we have that

$$\hat{\phi} - 1 = \frac{\frac{1}{n\sigma_w^2} \sum_{t=1}^n w_t x_{t-1}}{\frac{1}{n\sigma_w^2} \sum_{t=1}^n x_{t-1}^2}. \quad (5.28)$$

Consider the numerator of (5.28). Note first that by squaring both sides of $x_t = x_{t-1} + w_t$, we obtain $x_t^2 = x_{t-1}^2 + 2x_{t-1}w_t + w_t^2$ so that

$$x_{t-1}w_t = \frac{1}{2}(x_t^2 - x_{t-1}^2 - w_t^2),$$

and summing,

$$\frac{1}{n\sigma_w^2} \sum_{t=1}^n x_{t-1}w_t = \frac{1}{2} \left(\frac{x_n^2}{n\sigma_w^2} - \frac{\sum_{t=1}^n w_t^2}{n\sigma_w^2} \right).$$

Because $x_n = \sum_{t=1}^n w_t$, we have that $x_n \sim N(0, n\sigma_w^2)$, so that $\frac{1}{n\sigma_w^2} x_n^2 \sim \chi_1^2$, the chi-squared distribution with one degree of freedom. Moreover, because w_t is white Gaussian noise, $\frac{1}{n} \sum_{t=1}^n w_t^2 \rightarrow_p \sigma_w^2$, or $\frac{1}{n\sigma_w^2} \sum_{t=1}^n w_t^2 \rightarrow_p 1$. Consequently, ($n \rightarrow \infty$)

$$\frac{1}{n\sigma_w^2} \sum_{t=1}^n x_{t-1}w_t \xrightarrow{d} \frac{1}{2}(\chi_1^2 - 1). \quad (5.29)$$

Next we focus on the denominator of (5.28). First, we introduce standard Brownian motion.

Definition 5.1 A continuous time process $\{W(t); t \geq 0\}$ is called **standard Brownian motion** if it satisfies the following conditions:

- (i) $W(0) = 0$;
- (ii) $\{W(t_2) - W(t_1), W(t_3) - W(t_2), \dots, W(t_n) - W(t_{n-1})\}$ are independent for any collection of points, $0 \leq t_1 < t_2 < \dots < t_n$, and integer $n > 2$;
- (iii) $W(t + \Delta t) - W(t) \sim N(0, \Delta t)$ for $\Delta t > 0$.

The result for the denominator uses the functional central limit theorem, which can be found in Billingsley (1999, §2.8). In particular, if $\xi_1, \dots, \xi_n$ is a sequence of iid random variables with mean 0 and variance 1, then, for $0 \leq t \leq 1$, the continuous time process

$$S_n(t) = \frac{1}{\sqrt{n}} \sum_{j=1}^{\llbracket nt \rrbracket} \xi_j \xrightarrow{d} W(t), \quad (5.30)$$

as $n \rightarrow \infty$, where $\llbracket \cdot \rrbracket$ is the greatest integer function and $W(t)$ is standard Brownian motion on $[0, 1]$. Note the under the null hypothesis, $x_s = w_1 + \dots + w_s \sim N(0, s\sigma_w^2)$,

and based on (5.30), we have $\frac{x_s}{\sigma_w \sqrt{n}} \rightarrow_d W(s)$. From this fact, we can show that ($n \rightarrow \infty$)

$$\sum_{t=1}^n \left(\frac{x_{t-1}}{\sigma_w \sqrt{n}} \right)^2 \frac{1}{n} \xrightarrow{d} \int_0^1 W^2(t) dt. \quad (5.31)$$

The denominator in (5.28) is off from the left side of (5.31) by a factor of n^{-1} , and we adjust accordingly to finally obtain ($n \rightarrow \infty$),

$$n(\hat{\phi} - 1) = \frac{\frac{1}{n\sigma_w^2} \sum_{t=1}^n w_t x_{t-1}}{\frac{1}{n^2\sigma_w^2} \sum_{t=1}^n x_{t-1}^2} \xrightarrow{d} \frac{\frac{1}{2}(\chi_1^2 - 1)}{\int_0^1 W^2(t) dt}. \quad (5.32)$$

The test statistic $n(\hat{\phi} - 1)$ is known as the unit root or Dickey-Fuller (DF) statistic (see Fuller, 1996), although the actual DF test statistic is normalized a little differently. Because the distribution of the test statistic does not have a closed form, quantiles of the distribution must be computed by numerical approximation or by simulation. The R package `tseries` provides this test along with more general tests that we mention briefly.

Toward a more general model, we note that the DF test was established by noting that if $x_t = \phi x_{t-1} + w_t$, then $\nabla x_t = (\phi - 1)x_{t-1} + w_t = \gamma x_{t-1} + w_t$, and one could test $H_0: \gamma = 0$ by regressing ∇x_t on x_{t-1} . They formed a Wald statistic and derived its limiting distribution [the previous derivation based on Brownian motion is due to Phillips (1987)]. The test was extended to accommodate AR(p) models, $x_t = \sum_{j=1}^p \phi_j x_{t-j} + w_t$, as follows. Subtract x_{t-1} from the model to obtain

$$\nabla x_t = \gamma x_{t-1} + \sum_{j=1}^{p-1} \psi_j \nabla x_{t-j} + w_t, \quad (5.33)$$

where $\gamma = \sum_{j=1}^p \phi_j - 1$ and $\psi_j = -\sum_{i=j}^p \phi_i$ for $j = 2, \dots, p$. For a quick check of (5.33) when $p = 2$, note that $x_t = (\phi_1 + \phi_2)x_{t-1} - \phi_2(x_{t-1} - x_{t-2}) + w_t$; now subtract x_{t-1} from both sides. To test the hypothesis that the process has a unit root at 1 (i.e., the AR polynomial $\phi(z) = 0$ when $z = 1$), we can test $H_0: \gamma = 0$ by estimating γ in the regression of ∇x_t on $x_{t-1}, \nabla x_{t-1}, \dots, \nabla x_{t-p+1}$, and forming a Wald test based on $t_\gamma = \hat{\gamma}/\text{se}(\hat{\gamma})$. This test leads to the so-called augmented Dickey-Fuller test (ADF). While the calculations for obtaining the asymptotic null distribution change, the basic ideas and machinery remain the same as in the simple case. The choice of p is crucial, and we will discuss some suggestions in the example. For ARMA(p, q) models, the ADF test can be used by assuming p is large enough to capture the essential correlation structure; another alternative is the Phillips-Perron (PP) test, which differs from the ADF tests mainly in how they deal with serial correlation and heteroskedasticity in the errors.

One can extend the model to include a constant, or even non-stochastic trend. For example, consider the model

$$x_t = \beta_0 + \beta_1 t + \phi x_{t-1} + w_t.$$

If we assume $\beta_1 = 0$, then under the null hypothesis, $\phi = 1$, the process is a random walk with drift β_0 . Under the alternate hypothesis, the process is a causal AR(1) with mean $\mu_x = \beta_0(1 - \phi)$. If we cannot assume $\beta_1 = 0$, then the interest here is testing the null that $(\beta_1, \phi) = (0, 1)$, simultaneously, versus the alternative that $\beta_1 \neq 0$ and $|\phi| < 1$. In this case, the null hypothesis is that the process is a random walk with drift, versus the alternative hypothesis that the process is stationary around a global trend (consider the global temperature series examined in [Example 2.1](#)).

Example 5.3 Testing Unit Roots in the Glacial Varve Series

In this example we use the R package `tseries` to test the null hypothesis that the log of the glacial varve series has a unit root, versus the alternate hypothesis that the process is stationary. We test the null hypothesis using the available DF, ADF and PP tests; note that in each case, the general regression equation incorporates a constant and a linear trend. In the ADF test, the default number of AR components included in the model, say k , is $\llbracket (n - 1)^{\frac{1}{3}} \rrbracket$, which corresponds to the suggested upper bound on the rate at which the number of lags, k , should be made to grow with the sample size for the general ARMA(p, q) setup. For the PP test, the default value of k is $\llbracket .04n^{\frac{1}{4}} \rrbracket$.

```
library(tseries)
adf.test(log(varve), k=0)           # DF test
  Dickey-Fuller = -12.8572, Lag order = 0, p-value < 0.01
  alternative hypothesis: stationary
adf.test(log(varve))               # ADF test
  Dickey-Fuller = -3.5166, Lag order = 8, p-value = 0.04071
  alternative hypothesis: stationary
pp.test(log(varve))                # PP test
  Dickey-Fuller Z(alpha) = -304.5376,
  Truncation lag parameter = 6, p-value < 0.01
  alternative hypothesis: stationary
```

In each test, we reject the null hypothesis that the logged varve series has a unit root. The conclusion of these tests supports the conclusion of the previous section that the logged varve series is long memory rather than integrated.

5.4 GARCH Models

Recent problems in finance have motivated the study of the volatility, or variability, of a time series. Although ARMA models assume a constant variance, models such as the autoregressive conditionally heteroscedastic or ARCH model, first introduced by Engle (1982), were developed to model changes in volatility. These models were later extended to generalized ARCH, or GARCH models by Bollerslev (1986).

In §3.8, we discussed the return or growth rate of a series. For example, if x_t is the value of a stock at time t , then the return or relative gain, y_t , of the stock at time t is

$$y_t = \frac{x_t - x_{t-1}}{x_{t-1}}. \quad (5.34)$$

Definition (5.34) implies that $x_t = (1 + y_t)x_{t-1}$. Thus, based on the discussion in §3.8, if the return represents a small (in magnitude) percentage change then

$$\nabla[\log(x_t)] \approx y_t. \quad (5.35)$$

Either value, $\nabla[\log(x_t)]$ or $(x_t - x_{t-1})/x_{t-1}$, will be called the return, and will be denoted by y_t . It is the study of y_t that is the focus of ARCH, GARCH, and other volatility models. Recently there has been interest in stochastic volatility models and we will discuss these models in Chapter 6 because they are state-space models.

Typically, for financial series, the return y_t , does not have a constant conditional variance, and highly volatile periods tend to be clustered together. In other words, there is a strong dependence of sudden bursts of variability in a return on the series own past. For example, [Figure 1.4](#) shows the daily returns of the New York Stock Exchange (NYSE) from February 2, 1984 to December 31, 1991. In this case, as is typical, the return y_t is fairly stable, except for short-term bursts of high volatility.

The simplest ARCH model, the ARCH(1), models the return as

$$y_t = \sigma_t \epsilon_t \quad (5.36)$$

$$\sigma_t^2 = \alpha_0 + \alpha_1 y_{t-1}^2, \quad (5.37)$$

where ϵ_t is standard Gaussian white noise; that is, $\epsilon_t \sim \text{iid } N(0, 1)$. As with ARMA models, we must impose some constraints on the model parameters to obtain desirable properties. One obvious constraint is that α_1 must not be negative, or else σ_t^2 may be negative.

As we shall see, the ARCH(1) models return as a white noise process with non-constant conditional variance, and that conditional variance depends on the previous return. First, notice that the conditional distribution of y_t given y_{t-1} is Gaussian:

$$y_t \mid y_{t-1} \sim N(0, \alpha_0 + \alpha_1 y_{t-1}^2). \quad (5.38)$$

In addition, it is possible to write the ARCH(1) model as a non-Gaussian AR(1) model in the square of the returns y_t^2 . First, rewrite (5.36)-(5.37) as

$$\begin{aligned} y_t^2 &= \sigma_t^2 \epsilon_t^2 \\ \alpha_0 + \alpha_1 y_{t-1}^2 &= \sigma_t^2, \end{aligned}$$

and subtract the two equations to obtain

$$y_t^2 - (\alpha_0 + \alpha_1 y_{t-1}^2) = \sigma_t^2 \epsilon_t^2 - \sigma_t^2.$$

Now, write this equation as

$$y_t^2 = \alpha_0 + \alpha_1 y_{t-1}^2 + v_t, \quad (5.39)$$

where $v_t = \sigma_t^2(\epsilon_t^2 - 1)$. Because ϵ_t^2 is the square of a $N(0, 1)$ random variable, $\epsilon_t^2 - 1$ is a shifted (to have mean-zero), χ_1^2 random variable.

To explore the properties of ARCH, we define $\mathcal{Y}_s = \{y_s, y_{s-1}, \dots\}$. Then, using (5.38), we immediately see that y_t has a zero mean:

$$E(y_t) = EE(y_t \mid \mathcal{Y}_{t-1}) = EE(y_t \mid y_{t-1}) = 0. \quad (5.40)$$

Because $E(y_t | Y_{t-1}) = 0$, the process y_t is said to be a *martingale difference*.

Because y_t is a martingale difference, it is also an uncorrelated sequence. For example, with $h > 0$,

$$\begin{aligned} \text{cov}(y_{t+h}, y_t) &= E(y_t y_{t+h}) = EE(y_t y_{t+h} | \mathcal{Y}_{t+h-1}) \\ &= E\{y_t E(y_{t+h} | \mathcal{Y}_{t+h-1})\} = 0. \end{aligned} \quad (5.41)$$

The last line of (5.41) follows because y_t belongs to the information set $\mathcal{Y}_{t+h-1}$ for $h > 0$, and, $E(y_{t+h} | \mathcal{Y}_{t+h-1}) = 0$, as determined in (5.40).

An argument similar to (5.40) and (5.41) will establish the fact that the error process v_t in (5.39) is also a martingale difference and, consequently, an uncorrelated sequence. If the variance of v_t is finite and constant with respect to time, and $0 \leq \alpha_1 < 1$, then based on **Property 3.1**, (5.39) specifies a causal AR(1) process for y_t^2 . Therefore, $E(y_t^2)$ and $\text{var}(y_t^2)$ must be constant with respect to time t . This, implies that

$$E(y_t^2) = \text{var}(y_t) = \frac{\alpha_0}{1 - \alpha_1} \quad (5.42)$$

and, after some manipulations,

$$E(y_t^4) = \frac{3\alpha_0^2}{(1 - \alpha_1)^2} \frac{1 - \alpha_1^2}{1 - 3\alpha_1^2}, \quad (5.43)$$

provided $3\alpha_1^2 < 1$. These results imply that the kurtosis, κ , of y_t is

$$\kappa = \frac{E(y_t^4)}{[E(y_t^2)]^2} = 3 \frac{1 - \alpha_1^2}{1 - 3\alpha_1^2}, \quad (5.44)$$

which is always larger than 3 (unless $\alpha_1 = 0$), the kurtosis of the normal distribution. Thus, the marginal distribution of the returns, y_t , is leptokurtic, or has “fat tails.” Thus, if $0 \leq \alpha_1 < 1$, the process y_t itself is white noise and its unconditional distribution is symmetrically distributed around zero; this distribution is leptokurtic. If, in addition, $3\alpha_1^2 < 1$, the square of the process, y_t^2 , follows a causal AR(1) model with ACF given by $\rho_{y^2}(h) = \alpha_1^h \geq 0$, for all $h > 0$. If $3\alpha_1 \geq 1$, but $\alpha_1 < 1$, it can be shown that y_t^2 is strictly stationary with infinite variance (see Douc, et al., 2014).

Estimation of the parameters α_0 and α_1 of the ARCH(1) model is typically accomplished by conditional MLE. The conditional likelihood of the data $y_2, \dots, y_n$ given y_1 , is given by

$$L(\alpha_0, \alpha_1 | y_1) = \prod_{t=2}^n f_{\alpha_0, \alpha_1}(y_t | y_{t-1}), \quad (5.45)$$

where the density $f_{\alpha_0, \alpha_1}(y_t | y_{t-1})$ is the normal density specified in (5.38). Hence, the criterion function to be minimized, $l(\alpha_0, \alpha_1) \propto -\ln L(\alpha_0, \alpha_1 | y_1)$ is given by

$$l(\alpha_0, \alpha_1) = \frac{1}{2} \sum_{t=2}^n \ln(\alpha_0 + \alpha_1 y_{t-1}^2) + \frac{1}{2} \sum_{t=2}^n \left(\frac{y_t^2}{\alpha_0 + \alpha_1 y_{t-1}^2} \right). \quad (5.46)$$




Fig. 5.5. ACF and PACF of the squares of the residuals from the AR(1) fit on U.S. GNP.

Estimation is accomplished by numerical methods, as described in §3.6. In this case, analytic expressions for the gradient vector, $l^{(1)}(\alpha_0, \alpha_1)$, and Hessian matrix, $l^{(2)}(\alpha_0, \alpha_1)$, as described in [Example 3.29](#), can be obtained by straight-forward calculations. For example, the 2×1 gradient vector, $l^{(1)}(\alpha_0, \alpha_1)$, is given by

$$\begin{pmatrix} \partial l / \partial \alpha_0 \\ \partial l / \partial \alpha_1 \end{pmatrix} = \sum_{t=2}^n \begin{pmatrix} 1 \\ y_{t-1}^2 \end{pmatrix} \times \frac{\alpha_0 + \alpha_1 y_{t-1}^2 - y_t^2}{2 (\alpha_0 + \alpha_1 y_{t-1}^2)^2}.$$

The calculation of the Hessian matrix is left as an exercise ([Problem 5.8](#)). The likelihood of the ARCH model tends to be flat unless n is very large. A discussion of this problem can be found in [Shephard \(1996\)](#).

It is also possible to combine a regression or an ARMA model for the mean with an ARCH model for the errors. For example, a regression with ARCH(1) errors model would have the observations x_t as linear function of p regressors, $\mathbf{z}_t = (z_{t1}, \dots, z_{tp})'$, and ARCH(1) noise y_t , say,

$$x_t = \boldsymbol{\beta}' \mathbf{z}_t + y_t,$$

where y_t satisfies (5.36)-(5.37), but, in this case, is unobserved. Similarly, for example, an AR(1) model for data x_t exhibiting ARCH(1) errors would be

$$x_t = \phi_0 + \phi_1 x_{t-1} + y_t.$$

These types of models were explored by [Weiss \(1984\)](#).

Example 5.4 Analysis of U.S. GNP

In [Example 3.38](#), we fit an MA(2) model and an AR(1) model to the U.S. GNP series and we concluded that the residuals from both fits appeared to behave like a white noise process. In [Example 3.42](#) we concluded that the AR(1) is probably the better model in this case. It has been suggested that the U.S. GNP series has ARCH errors, and in this example, we will investigate this claim. If the GNP noise term is ARCH, the squares of the residuals from the fit should behave like a non-Gaussian AR(1) process, as pointed out in (5.39). [Figure 5.5](#) shows the ACF and PACF of the squared residuals it appears that there may be some dependence, albeit small, left in the residuals. The figure was generated in R as follows.

```
gnpgr = diff(log(gnp)) # get the growth rate
sarima(gnpgr, 1, 0, 0) # fit an AR(1)
acf2(innov^2, 24) # get (p)acf of the squared residuals
```

We used the R package `fGarch` to fit an AR(1)-ARCH(1) model to the U.S. GNP returns with the following results. A partial output is shown; we note that `garch(1,0)` specifies an ARCH(1) in the code below (details later).

```
library(fGarch)
summary(garchFit(~arma(1,0)+garch(1,0), gnpgr))
```

	Estimate	Std. Error	t value	Pr(>t)
mu	5.278e-03	8.996e-04	5.867	4.44e-09
ar1	3.666e-01	7.514e-02	4.878	1.07e-06
omega	7.331e-05	9.011e-06	8.135	4.44e-16
alpha1	1.945e-01	9.554e-02	2.035	0.0418

Standardised Residuals Tests:	Statistic	p-Value
Jarque-Bera Test	R Chi^2	9.118036 0.01047234
Shapiro-Wilk Test	R W	0.9842407 0.01433690
Ljung-Box Test	R Q(10)	9.874326 0.4515875
Ljung-Box Test	R Q(15)	17.55855 0.2865844
Ljung-Box Test	R Q(20)	23.41363 0.2689437
Ljung-Box Test	R^2 Q(10)	19.2821 0.03682246
Ljung-Box Test	R^2 Q(15)	33.23648 0.004352736
Ljung-Box Test	R^2 Q(20)	37.74259 0.009518992
LM Arch Test	R TR^2	25.41625 0.01296901

In this example, we obtain $\hat{\phi}_0 = .005$ (called `mu` in the output) and $\hat{\phi}_1 = .367$ (called `ar1`) for the AR(1) parameter estimates; in [Example 3.38](#) the values were .005 and .347, respectively. The ARCH(1) parameter estimates are $\hat{\alpha}_0 = 0$ (called `omega`) for the constant and $\hat{\alpha}_1 = .195$, which is significant with a p-value of about .04. There are a number of tests that are performed on the residuals [R] or the squared residuals [R^2]. For example, the Jarque-Bera statistic tests the residuals of the fit for normality based on the observed skewness and kurtosis, and it appears that the residuals have some non-normal skewness and kurtosis. The Shapiro-Wilk statistic tests the residuals of the fit for normality based on the empirical order statistics. The other tests, primarily based on the Q-statistic, are used on the residuals and their squares.

The ARCH(1) model can be extended to the general ARCH(m) model in an obvious way. That is, (5.36), $y_t = \sigma_t \epsilon_t$, is retained, but (5.37) is extended to

$$\sigma_t^2 = \alpha_0 + \alpha_1 y_{t-1}^2 + \cdots + \alpha_m y_{t-m}^2. \quad (5.47)$$

Estimation for ARCH(m) also follows in an obvious way from the discussion of estimation for ARCH(1) models. That is, the conditional likelihood of the data $y_{m+1}, \dots, y_n$ given $y_1, \dots, y_m$, is given by

$$L(\boldsymbol{\alpha} \mid y_1, \dots, y_m) = \prod_{t=m+1}^n f_{\boldsymbol{\alpha}}(y_t \mid y_{t-1}, \dots, y_{t-m}), \quad (5.48)$$

where $\boldsymbol{\alpha} = (\alpha_0, \alpha_1, \dots, \alpha_m)$ and the conditional densities $f_{\boldsymbol{\alpha}}(\cdot \mid \cdot)$ in (5.48) are normal densities; that is, for $t > m$,

$$y_t \mid y_{t-1}, \dots, y_{t-m} \sim N(0, \alpha_0 + \alpha_1 y_{t-1}^2 + \cdots + \alpha_m y_{t-m}^2).$$

Another extension of ARCH is the generalized ARCH or GARCH model developed by Bollerslev (1986). For example, a GARCH(1, 1) model retains (5.36), $y_t = \sigma_t \epsilon_t$, but extends (5.37) as follows:

$$\sigma_t^2 = \alpha_0 + \alpha_1 y_{t-1}^2 + \beta_1 \sigma_{t-1}^2. \quad (5.49)$$

Under the condition that $\alpha_1 + \beta_1 < 1$, using similar manipulations as in (5.39), the GARCH(1, 1) model, (5.36) and (5.49), admits a non-Gaussian ARMA(1, 1) model for the squared process

$$y_t^2 = \alpha_0 + (\alpha_1 + \beta_1) y_{t-1}^2 + v_t - \beta_1 v_{t-1}, \quad (5.50)$$

where v_t is as defined in (5.39). Representation (5.50) follows by writing (5.36) as

$$\begin{aligned} y_t^2 - \sigma_t^2 &= \sigma_t^2 (\epsilon_t^2 - 1) \\ \beta_1 (y_{t-1}^2 - \sigma_{t-1}^2) &= \beta_1 \sigma_{t-1}^2 (\epsilon_{t-1}^2 - 1), \end{aligned}$$

subtracting the second equation from the first, and using the fact that, from (5.49), $\sigma_t^2 - \beta_1 \sigma_{t-1}^2 = \alpha_0 + \alpha_1 y_{t-1}^2$, on the left-hand side of the result. The GARCH(m, r) model retains (5.36) and extends (5.49) to

$$\sigma_t^2 = \alpha_0 + \sum_{j=1}^m \alpha_j y_{t-j}^2 + \sum_{j=1}^r \beta_j \sigma_{t-j}^2. \quad (5.51)$$

Conditional maximum likelihood estimation of the GARCH(m, r) model parameters is similar to the ARCH(m) case, wherein the conditional likelihood, (5.48), is the product of $N(0, \sigma_t^2)$ densities with σ_t^2 given by (5.51) and where the conditioning is on the first $\max(m, r)$ observations, with $\sigma_1^2 = \cdots = \sigma_r^2 = 0$. Once the parameter estimates are obtained, the model can be used to obtain one-step-ahead forecasts of the volatility, say $\hat{\sigma}_{t+1}^2$, given by

$$\hat{\sigma}_{t+1}^2 = \hat{\alpha}_0 + \sum_{j=1}^m \hat{\alpha}_j y_{t+1-j}^2 + \sum_{j=1}^r \hat{\beta}_j \hat{\sigma}_{t+1-j}^2. \quad (5.52)$$

We explore these concepts in the following example.

Example 5.5 GARCH Analysis of the NYSE Returns

As previously mentioned, the daily returns of the NYSE shown in Figure 1.4 exhibit classic GARCH features. We used the R `fGarch` package to fit a GARCH(1, 1) model to the series with the following results:

```
library(fGarch)
nyse = asts::nyse # fGarch has one too
summary(nyse.g <- garchFit(~garch(1,1), nyse))
```

	Estimate	Std. Error	t value	Pr(>t)
mu	7.369e-04	1.786e-04	4.126	3.69e-05
omega	6.542e-06	1.455e-06	4.495	6.94e-06
alpha1	1.141e-01	1.604e-02	7.114	1.13e-12
beta1	8.061e-01	2.973e-02	27.112	< 2e-16

Standardised Residuals Tests:

		Statistic	p-Value
Jarque-Bera Test	R	Chi^2	3628.415 0
Shapiro-Wilk Test	R	W	0.9515562 0
Ljung-Box Test	R	Q(10)	29.69242 0.0009616813
Ljung-Box Test	R	Q(15)	30.50938 0.01021164
Ljung-Box Test	R	Q(20)	32.81143 0.03538324
Ljung-Box Test	R^2	Q(10)	3.510505 0.9667405
Ljung-Box Test	R^2	Q(15)	4.408852 0.9960585
Ljung-Box Test	R^2	Q(20)	6.68935 0.9975864
LM Arch Test	R	TR^2	3.967784 0.9840107

To explore the GARCH predictions of volatility, we calculated and plotted the 100 observations from the middle of the data (which includes the October 19, 1987 crash) along with the one-step-ahead predictions of the corresponding volatility, σ_t^2 . The results are displayed as the data $\pm 2\hat{\sigma}_t$ as a dashed line surrounding the data in Figure 5.6.

```
u = nyse.g@sigma.t
plot(window(nyse, start=900, end=1000), ylim=c(-.22,.2), ylab="NYSE Returns")
lines(window(nyse-2*u, start=900, end=1000), lty=2, col=4)
lines(window(nyse+2*u, start=900, end=1000), lty=2, col=4)
```

Some key points can be gleaned from the examples of this section. First, it is apparent that the conditional distribution of the returns is rarely normal. `fGarch` allows for various distributions to be fit to the data; see the help file for information. Some drawbacks of the GARCH model are: (i) the model assumes positive and negative returns have the same effect because volatility depends on squared returns; (ii) the model is restrictive because of the tight constraints on the model parameters (e.g., for an ARCH(1), $0 \leq \alpha_1^2 < \frac{1}{3}$); (iii) the likelihood is flat unless n is very large; (iv) the model tends to overpredict volatility because it responds slowly to large isolated returns.

Various extensions to the original model have been proposed to overcome some of the shortcomings we have just mentioned. For example, we have already discussed the fact that the S-PLUS Garch module will fit some non-normal, albeit symmetric, distributions. For asymmetric return dynamics, one can use the EGARCH (exponential GARCH) model, which is a complex model that has different components for positive returns and for negative returns. In the case of persistence in volatility, the integrated GARCH (IGARCH) model may be used. Recall (5.50) where we showed




Fig. 5.6. GARCH predictions of the NYSE volatility, $\pm 2\hat{\sigma}_t$, displayed as dashed lines.

the GARCH(1, 1) model can be written as

$$y_t^2 = \alpha_0 + (\alpha_1 + \beta_1)y_{t-1}^2 + v_t - \beta_1 v_{t-1}$$

and y_t^2 is stationary if $\alpha_1 + \beta_1 < 1$. The IGARCH model sets $\alpha_1 + \beta_1 = 1$, in which case the IGARCH(1, 1) model is

$$y_t = \sigma_t \epsilon_t \quad \text{and} \quad \sigma_t^2 = \alpha_0 + (1 - \beta_1)y_{t-1}^2 + \beta_1 \sigma_{t-1}^2.$$

There are many different extensions to the basic ARCH model that were developed to handle the various situations noticed in practice. Interested readers might find the general discussions in Engle et al. (1994) and Shephard (1996) worthwhile reading. Also, Gouriéroux (1997) gives a detailed presentation of ARCH and related models with financial applications and contains an extensive bibliography. Two excellent texts on financial time series analysis are Chan (2002) and Tsay (2002).

Finally, we briefly discuss stochastic volatility models; a detailed treatment of these models is given in Chapter 6. The volatility component, σ_t^2 , in the GARCH model is conditionally nonstochastic. In the ARCH(1) model for example, any time the previous return is zero, i.e., $y_{t-1} = 0$, it must be the case that $\sigma_t^2 = \alpha_0$, and so on. This assumption seems a bit unrealistic. The stochastic volatility model adds a stochastic component to the volatility in the following way. In the GARCH model, a return, say y_t , is

$$y_t = \sigma_t \epsilon_t \quad \Rightarrow \quad \log y_t^2 = \log \sigma_t^2 + \log \epsilon_t^2. \quad (5.53)$$

Thus, the observations $\log y_t^2$ are generated by two components, the unobserved volatility $\log \sigma_t^2$ and the unobserved noise $\log \epsilon_t^2$. While, for example, the GARCH(1, 1) models volatility without error, $\sigma_{t+1}^2 = \alpha_0 + \alpha_1 y_t^2 + \beta_1 \sigma_t^2$, the basic stochastic volatility model assumes the latent variable is an autoregressive process,

$$\log \sigma_{t+1}^2 = \phi_0 + \phi_1 \log \sigma_t^2 + w_t \quad (5.54)$$

where $w_t \sim \text{iid } N(0, \sigma_w^2)$. The introduction of the noise term w_t makes the latent volatility process stochastic. Together (5.53) and (5.54) comprise the stochastic volatility model. Given n observations, the goals are to estimate the parameters ϕ_0 , ϕ_1 and σ_w^2 , and then predict future observations $\log y_{n+m}^2$. Details are provided in §6.10.

5.5 Threshold Models

In §3.5 we discussed the fact that, for a stationary time series, best linear prediction forward in time is the same as best linear prediction backward in time. This result followed from the fact that the variance–covariance matrix of $\mathbf{x}_{1:n} = (x_1, x_2, \dots, x_n)'$, say, $\Gamma = \{\gamma(i-j)\}_{i,j=1}^n$, is the same as the variance–covariance matrix of $\mathbf{x}_{n:1} = (x_n, x_{n-1}, \dots, x_1)'$. In addition, if the process is Gaussian, the distributions of $\mathbf{x}_{1:n}$ and $\mathbf{x}_{n:1}$ are identical. In this case, a time plot of $\mathbf{x}_{1:n}$ (that is, the data plotted forward in time) should look similar to a time plot of $\mathbf{x}_{n:1}$ (that is, the data plotted backward in time).

There are, however, many series that do not fit into this category. For example, Figure 5.7 shows a plot of monthly pneumonia and influenza deaths per 10,000 in the U.S. for 11 years, 1968 to 1978. Typically, the number of deaths tends to increase slower than it decreases. Thus, if the data were plotted backward in time, the backward series would tend to increase faster than it decreases. Also, if monthly pneumonia and influenza deaths followed a linear Gaussian process, we would not expect to see such large bursts of positive and negative changes that occur periodically in this series. Moreover, although the number of deaths is typically largest during the winter months, the data are not perfectly seasonal. That is, although the peak of the series often occurs in January, in other years, the peak occurs in February or in March.

Many approaches to modeling nonlinear series exist that could be used (see Priestley, 1988); here, we focus on the class of threshold autoregressive models presented in Tong (1983, 1990). The basic idea of these models is that of fitting local linear $\text{AR}(p)$ models, and their appeal is that we can use the intuition from fitting global linear $\text{AR}(p)$ models. Suppose we know p , and given the vectors $\mathbf{x}_{t-1} = (x_{t-1}, \dots, x_{t-p})'$, we can identify r mutually exclusive and exhaustive regions for $\mathbf{x}_{t-1}$, say, $R_1, \dots, R_r$, where the dynamics of the system changes. The threshold model is then written as r $\text{AR}(p)$ models,

$$x_t = \alpha^{(j)} + \phi_1^{(j)} x_{t-1} + \dots + \phi_p^{(j)} x_{t-p} + w_t^{(j)}, \quad \mathbf{x}_{t-1} \in R_j, \quad (5.55)$$

for $j = 1, \dots, r$. In (5.55), the $w_t^{(j)}$ are independent white noise series, each with variance σ_j^2 , for $j = 1, \dots, r$. Model estimation, identification, and diagnostics proceed as in the case in which $r = 1$.




Fig. 5.7. U.S. monthly pneumonia and influenza deaths per 10,000.

Example 5.6 Threshold Modeling of the Influenza Series

As previously discussed, examination of Figure 5.7 leads us to believe that the monthly pneumonia and influenza deaths time series, say flu_t , is not linear. It is also evident from Figure 5.7 that there is a slight negative trend in the data. We have found that the most convenient way to fit a threshold model to this data set, while removing the trend, is to work with the first difference of the data. The differenced data,

$$x_t = \text{flu}_t - \text{flu}_{t-1}$$

is exhibited in Figure 5.8 as the dark solid line with circles representing observations. The dashed line with squares in Figure 5.8 are the one-month-ahead predictions, and we will discuss this series later.

The nonlinearity of the data is more pronounced in the plot of the first differences, x_t . Clearly, the change in the numbers of deaths, x_t , slowly rises for some months and, then, sometime in the winter, has a possibility of jumping to a large number once x_t exceeds about .05. If the process does make a large jump, then a subsequent significant decrease occurs in flu deaths. As an initial analysis, we fit the following threshold model

$$\begin{aligned} x_t &= \alpha^{(1)} + \sum_{j=1}^p \phi_j^{(1)} x_{t-j} + w_t^{(1)}, & x_{t-1} < .05 \\ x_t &= \alpha^{(2)} + \sum_{j=1}^p \phi_j^{(2)} x_{t-j} + w_t^{(2)}, & x_{t-1} \geq .05, \end{aligned} \quad (5.56)$$




Fig. 5.8. First differenced U.S. monthly pneumonia and influenza deaths (triangles); one-month-ahead predictions (dashed red line) with ± 2 prediction error bounds (solid blue lines).

with $p = 6$, assuming this would be larger than necessary.

Model (5.56) is easy to fit using two linear regression runs. That is, let $\delta_t^{(1)} = 12$ if $x_{t-1} < .05$, and zero otherwise, and let $\delta_t^{(2)} = 1$ if $x_{t-1} \geq .05$, and zero otherwise. Then, using the notation of §2.2, for $t = p + 1, \dots, n$, either equation in (5.56) can be written as

$$y_t = \boldsymbol{\beta}' \mathbf{z}_t + w_t$$

where, for $i = 1, 2$,

$$y_t = \delta_t^{(i)} x_t, \quad \mathbf{z}_t' = \delta_t^{(i)} (1, x_{t-1}, \dots, x_{t-p}), \quad w_t = \delta_t^{(i)} w_t^{(i)},$$

and

$$\boldsymbol{\beta}' = (\alpha^{(i)}, \phi_1^{(i)}, \phi_2^{(i)}, \dots, \phi_p^{(i)}).$$

Parameter estimates can then be obtained using the regression techniques of §2.2 twice, once for $i = 1$ and again for $i = 2$.

An order $p = 4$ was finally selected for each part of the model. The final model was

$$\hat{x}_t = .51_{(.08)}x_{t-1} - .20_{(.06)}x_{t-2} + .12_{(.05)}x_{t-3} \\ - .11_{(.05)}x_{t-4} + \hat{w}_t^{(1)}, \quad \text{when } x_{t-1} < .05$$

$$\hat{x}_t = .40 - .75_{(.17)}x_{t-1} - 1.03_{(.21)}x_{t-2} - 2.05_{(1.05)}x_{t-3} \\ - 6.71_{(1.25)}x_{t-4} + \hat{w}_t^{(2)}, \quad \text{when } x_{t-1} \geq .05,$$

where $\hat{\sigma}_1 = .05$ and $\hat{\sigma}_2 = .07$. The threshold of .05 was exceeded 17 times. Using the final model, one-month-ahead predictions can be made, and these are shown in [Figure 5.8](#) as a dashed line with squares. The model does extremely well at predicting a flu epidemic; the peak at 1976, however, was missed by this model. When we fit a model with a smaller threshold of .04, flu epidemics were somewhat underestimated, but the flu epidemic in the eighth year was predicted one month early. We chose the model with a threshold of .05 because the residual diagnostics showed no obvious departure from the model assumption (except for one outlier at 1976); the model with a threshold of .04 still had some correlation left in the residuals and there was more than one outlier. Finally, prediction beyond one-month-ahead for this model is complicated, but some approximate techniques exist (see Tong, 1983). The following commands can be used to perform this analysis in R. The first three lines generate [Figure 5.7](#).

```
# Plot data with month initials as points
plot(flu, type="c")
Months = c("J", "F", "M", "A", "M", "J", "J", "A", "S", "O", "N", "D")
points(flu, pch=Months, cex=.8, font=2)
# Start analysis
dflu = diff(flu)
thrsh = .05 # threshold
Z = ts.intersect(dflu, lag(dflu,-1), lag(dflu,-2), lag(dflu,-3),
  lag(dflu,-4))
ind1 = ifelse(Z[,2] < thrsh, 1, NA) # indicator < thrsh
ind2 = ifelse(Z[,2] < thrsh, NA, 1) # indicator >= thrsh
X1 = Z[,1]*ind1
X2 = Z[,1]*ind2
summary(fit1<-lm(X1~Z[,2:5])) # case 1
summary(fit2<-lm(X2~Z[,2:5])) # case 2
D = cbind(rep(1, nrow(Z)), Z[,2:5]) # get predictions
b1 = fit1$coef
b2 = fit2$coef
p1 = D%*%b1
p2 = D%*%b2
prd = ifelse(Z[,2] < thrsh, p1, p2)
plot(dflu, type="p", pch=2, ylim=c(-.5,.5))
lines(prd, lty=4, col=2)
prde1 = sqrt(sum(resid(fit1)^2)/df.residual(fit1))
prde2 = sqrt(sum(resid(fit2)^2)/df.residual(fit2))
prde = ifelse(Z[,2] < thrsh, prde1, prde2)
lines(prd + 2*prde, col=4)
lines(prd - 2*prde, col=4)
```

The R package `tsDyn` can be used to fit these models; in this case, a similar analysis may be obtained as follows:

```
require(tsDyn) # load package - install it if you don't have it
```

```
# vignette("tsDyn") # for package details
dflu = diff(flu)
lag1.plot(dflu, 1, corr=FALSE) # see the nonlinearity here
(u = setar(dflu, m=4, thDelay=0)) # fit model and view results
BIC(u); AIC(u) # if you want to try other models; m=3 works well too
plot(u) # graphics - ?plot.setar for information
```

5.6 Lagged Regression: Transfer Function Modeling

In §3.10, we considered lagged regression in a frequency domain approach based on coherency. In this section we focus on a time domain approach to the same problem. In the previous section, we looked at autocorrelated errors but, still regarded the input series z_t as being fixed unknown functions of time. This consideration made sense for the time argument t , but was less satisfactory for the other inputs, which are probably stochastic processes. For example, consider the SOI and Recruitment series that were presented in Example 1.5. The series are displayed in Figure 1.5. In this case, the interest is in predicting the output Recruitment series, say, y_t , from the input SOI, say x_t . We might consider the lagged regression model

$$y_t = \sum_{j=0}^{\infty} \alpha_j x_{t-j} + \eta_t = \alpha(B)x_t + \eta_t, \quad (5.57)$$

where $\sum_j |\alpha_j| < \infty$. We assume the input process x_t and noise process η_t in (5.57) are both stationary and mutually independent. The coefficients $\alpha_0, \alpha_1, \dots$ describe the weights assigned to past values of x_t used in predicting y_t and we have used the notation

$$\alpha(B) = \sum_{j=0}^{\infty} \alpha_j B^j. \quad (5.58)$$

In the Box and Jenkins (1970) formulation, we assign ARIMA models, say, $\text{ARIMA}(p, d, q)$ and $\text{ARIMA}(p_\eta, d_\eta, q_\eta)$, to the series x_t and η_t , respectively. The components of (5.57) in backshift notation, for the case of simple $\text{ARMA}(p, q)$ modeling of the input and noise, would have the representation

$$\phi(B)x_t = \theta(B)w_t \quad (5.59)$$

and

$$\phi_\eta(B)\eta_t = \theta_\eta(B)z_t, \quad (5.60)$$

where w_t and z_t are independent white noise processes with variances σ_w^2 and σ_z^2 , respectively. Box and Jenkins (1970) proposed that systematic patterns often observed in the coefficients α_j , for $j = 1, 2, \dots$, could often be expressed as a ratio of polynomials involving a small number of coefficients, along with a specified delay, d , so

$$\alpha(B) = \frac{\delta(B)B^d}{\omega(B)}, \quad (5.61)$$

where

$$\omega(B) = 1 - \omega_1 B - \omega_2 B^2 - \cdots - \omega_r B^r \quad (5.62)$$

and

$$\delta(B) = \delta_0 + \delta_1 B + \cdots + \delta_s B^s \quad (5.63)$$

are the indicated operators; in this section, we find it convenient to represent the inverse of an operator, say, $[\omega(B)]^{-1}$, as $1/\omega(B)$.

Determining a parsimonious model involving a simple form for $\alpha(B)$ and estimating all of the parameters in the above model are the main tasks in the transfer function methodology. Because of the large number of parameters, it is necessary to develop a sequential methodology. Suppose we focus first on finding the ARIMA model for the input x_t and apply this operator to both sides of (5.57), obtaining the new model

$$\tilde{y}_t = \frac{\phi(B)}{\theta(B)} y_t = \alpha(B) w_t + \frac{\phi(B)}{\theta(B)} \eta_t = \alpha(B) w_t + \tilde{\eta}_t,$$

where w_t and the transformed noise $\tilde{\eta}_t$ are independent.

The series w_t is a prewhitened version of the input series, and its cross-correlation with the transformed output series $\tilde{y}_t$ will be just

$$\gamma_{\tilde{y}w}(h) = E[\tilde{y}_{t+h} w_t] = E \left[\sum_{j=0}^{\infty} \alpha_j w_{t+h-j} w_t \right] = \sigma_w^2 \alpha_h, \quad (5.64)$$

because the autocovariance function of white noise will be zero except when $j = h$ in (5.64). Hence, by computing the cross-correlation between the prewhitened input series and the transformed output series should yield a rough estimate of the behavior of $\alpha(B)$.

Example 5.7 Relating the Prewhitened SOI to the Transformed Recruitment Series

We give a simple example of the suggested procedure for the SOI and the Recruitment series. Figure 5.9 shows the sample ACF and PACF of the detrended SOI, and it is clear, from the PACF, that an autoregressive series with $p = 1$ will do a reasonable job. Fitting the series gave $\hat{\phi} = .588$ with $\hat{\sigma}_w^2 = .092$, and we applied the operator $(1 - .588B)$ to both x_t and y_t and computed the cross-correlation function, which is shown in Figure 5.10. Noting the apparent shift of $d = 5$ months and the decrease thereafter, it seems plausible to hypothesize a model of the form

$$\alpha(B) = \delta_0 B^5 (1 + \omega_1 B + \omega_1^2 B^2 + \cdots) = \frac{\delta_0 B^5}{1 - \omega_1 B}$$

for the transfer function. In this case, we would expect ω_1 to be negative. The following R code was used for this example.

```
acf2(soi)
(fit = arima(soi, xreg=time(soi), order=c(1, 0, 0)))
Coefficients:   ar1 intercept time(soi)
              0.5875   13.7507   -0.0069
```





Fig. 5.9. Sample ACF and PACF of SOI.

```

s.e.  0.0379      6.1775      0.0031
sigma^2 estimated as 0.0918
ar1 = as.numeric(fit$coef[1])    # = 0.5875387
soi.pw = resid(fit)
rec.d = resid(lm(rec~time(rec), na.action=NULL))
rec.fil = filter(rec.d, filter=c(1, -ar1), method="conv", sides=1)
ccf(soi.pw, rec.fil, main="", ylab="CCF", na.action=na.omit)

```

In the code above, `soi.pw` is the prewhitened detrended SOI series, `rec.d` is the detrended Recruitment series, and `rec.fil` is the filtered, detrended Recruitment series. In the `ccf` calculation, `na.action=na.omit` is used because `rec.fil[1]` is `NA`.

In some cases, we may postulate the form of the separate components $\delta(B)$ and $\omega(B)$, so we might write the equation

$$y_t = \frac{\delta(B)B^d}{\omega(B)}x_t + \eta_t$$

as

$$\omega(B)y_t = \delta(B)B^d x_t + \omega(B)\eta_t,$$

or in regression form

$$y_t = \sum_{k=1}^r \omega_k y_{t-k} + \sum_{k=0}^s \delta_k x_{t-d-k} + u_t, \quad (5.65)$$

where




Fig. 5.10. Sample CCF of the prewhitened, detrended SOI and the similarly transformed Recruitment series; negative lags indicate that SOI leads Recruitment.

$$u_t = \omega(B)\eta_t. \quad (5.66)$$

The form of (5.65) suggests doing a regression on the lagged versions of both the input and output series to obtain $\hat{\beta}$, the estimate of the $(r + s + 1) \times 1$ regression vector

$$\beta = (\omega_1, \dots, \omega_r, \delta_0, \delta_1, \dots, \delta_s)'$$

The residuals from the regression above, say,

$$\hat{u}_t = y_t - \hat{\beta}'z_t,$$

where

$$z_t = (y_{t-1}, \dots, y_{t-r}, x_{t-d}, \dots, x_{t-d-s})'$$

denotes the usual vector of independent variables, could be used to approximate the best ARMA model for the noise process η_t , because we can compute an estimator for that process from (5.66), using $\hat{u}_t$ and $\hat{\omega}(B)$ and applying the moving average operator to get $\hat{\eta}_t$. Fitting an ARMA(p_η, q_η) model to this estimated noise then completes the specification. The preceding suggests the following sequential procedure for fitting the transfer function model to data.

- (i) Fit an ARMA model to the input series x_t to estimate the parameters $\phi_1, \dots, \phi_p, \theta_1, \dots, \theta_q, \sigma_w^2$ in the specification (5.59). Retain ARMA coefficients for use in step (ii) and the fitted residuals $\hat{w}_t$ for use in Step (iii).
- (ii) Apply the operator determined in step (i), that is,

$$\hat{\phi}(B)y_t = \hat{\theta}(B)\tilde{y}_t,$$

to determine the transformed output series $\tilde{y}_t$.

- (iii) Use the cross-correlation function between $\tilde{y}_t$ and $\widehat{w}_t$ in steps (i) and (ii) to suggest a form for the components of the polynomial

$$\alpha(B) = \frac{\delta(B)B^d}{\omega(B)}$$

and the estimated time delay d .

- (iv) Obtain $\widehat{\beta} = (\widehat{\omega}_1, \dots, \widehat{\omega}_r, \widehat{\delta}_0, \widehat{\delta}_1, \dots, \widehat{\delta}_s)$ by fitting a linear regression of the form (5.65). Retain the residuals $\widehat{u}_t$ for use in step (v).
- (v) Apply the moving average transformation (5.66) to the residuals $\widehat{u}_t$ to find the noise series $\widehat{\eta}_t$, and fit an ARMA model to the noise, obtaining the estimated coefficients in $\widehat{\phi}_\eta(B)$ and $\widehat{\theta}_\eta(B)$.

The above procedure is fairly reasonable, but does not have any recognizable overall optimality. Simultaneous least squares estimation, based on the observed x_t and y_t , can be accomplished by noting that the transfer function model can be written as

$$y_t = \frac{\delta(B)B^d}{\omega(B)}x_t + \frac{\theta_\eta(B)}{\phi_\eta(B)}z_t,$$

which can be put in the form

$$\omega(B)\phi_\eta(B)y_t = \phi_\eta(B)\delta(B)B^d x_t + \omega(B)\theta_\eta(B)z_t, \quad (5.67)$$

and it is clear that we may use least squares to minimize $\sum_t z_t^2$, as in earlier sections. We may also express the transfer function in state-space form (see Brockwell and Davis, 1991, Chapter 12). It is often easier to fit a transfer function model in the spectral domain as presented in §4.10.

Example 5.8 Transfer Function Model for SOI and Recruitment

We illustrate the procedure for fitting a transfer function model of the form suggested in Example 5.7 to the detrended SOI series (x_t) and the detrended Recruitment series (y_t). The results reported here can be compared with the results obtained from the frequency domain approach used in Example 4.22. Note first that steps (i)–(iii). have already been applied to determine the ARMA model

$$(1 - .588B)x_t = w_t,$$

where $\widehat{\sigma}_w^2 = .092$. Using the model determined in Example 5.7, we run the regression

$$y_t = \omega_1 y_{t-1} + \delta_0 x_{t-5} + u_t,$$

yielding $\widehat{\omega}_1 = .853$, $\widehat{\delta}_0 = -18.86$, where the residuals satisfy

$$\widehat{u}_t = (1 - .853B)\eta_t.$$

This completes step (iv). To complete the specification, we apply the moving average operator above to estimate the original noise series η_t and fit a second-order autoregressive model, based on the ACF and PACF shown in Figure 5.11. We obtain




Fig. 5.11. ACF and PACF of the estimated noise $\hat{\eta}_t$ departures from the transfer function model.

$$(1 + .4271B - .3286B^2)\eta_t = z_t,$$

with $\hat{\sigma}_z^2 = 51.3$ as the estimated error variance.

The following R code was used for this example.

```
rec.d = resid(lm(rec~time(rec), na.action=NULL))
soi.d = resid(lm(soi~time(soi), na.action=NULL))
fish = ts.intersect(rec.d, rec.d1=lag(rec.d,-1), soi.d5=lag(soi,-5))
summary(fish.fit <- lm(rec.d~0+rec.d1+soi.d5, data=fish))
      Estimate Std Error t value Pr(>t)
rec.d1   0.8531    0.0144   59.23  <2e-16
soi.d5 -18.8627    1.0038  -18.79  <2e-16
Residual standard error: 8.02 on 446 degrees of freedom
om1 = as.numeric(fish.fit$coef[1])
eta.hat = filter(resid(fish.fit), filter=-om1, method="recur", sides=1)
acf2(eta.hat)
(eta.fit <- arima(eta.hat, order=c(2,0,0)))
      ar1      ar2  intercept
      -0.4271  0.3286      0.7337
s.e.    0.0446  0.0446      0.3079
sigma^2 estimated as 51.3
```

5.7 Multivariate ARMAX Models

To understand multivariate time series models and their capabilities, we first present an introduction to multivariate time series regression techniques. A useful extension of the basic univariate regression model presented in §2.1 is the case in which we have more than one output series, that is, multivariate regression analysis. Suppose, instead of a single output variable y_t , a collection of k output variables $y_{t1}, y_{t2}, \dots, y_{tk}$ exist that are related to the inputs as

$$y_{ti} = \beta_{i1}z_{t1} + \beta_{i2}z_{t2} + \cdots + \beta_{ir}z_{tr} + w_{ti} \quad (5.68)$$

for each of the $i = 1, 2, \dots, k$ output variables. We assume the w_{ti} variables are correlated over the variable identifier i , but are still independent over time. Formally, we assume $\text{cov}\{w_{si}, w_{tj}\} = \sigma_{ij}$ for $s = t$ and is zero otherwise. Then, writing (5.68) in matrix notation, with $\mathbf{y}_t = (y_{t1}, y_{t2}, \dots, y_{tk})'$ being the vector of outputs, and $\mathcal{B} = \{\beta_{ij}\}$, $i = 1, \dots, k$, $j = 1, \dots, r$ being a $k \times r$ matrix containing the regression coefficients, leads to the simple looking form

$$\mathbf{y}_t = \mathcal{B}\mathbf{z}_t + \mathbf{w}_t. \quad (5.69)$$

Here, the $k \times 1$ vector process $\mathbf{w}_t$ is assumed to be a collection of independent vectors with common covariance matrix $E\{\mathbf{w}_t\mathbf{w}_t'\} = \Sigma_w$, the $k \times k$ matrix containing the covariances σ_{ij} . The maximum likelihood estimator, under the assumption of normality, for the regression matrix in this case is

$$\widehat{\mathcal{B}} = Y'Z(Z'Z)^{-1}, \quad (5.70)$$

where $Z' = [\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_n]$ is as before and $Y' = [\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_n]$. The error covariance matrix Σ_w is estimated by

$$\widehat{\Sigma}_w = \frac{1}{n-r} \sum_{t=1}^n (\mathbf{y}_t - \widehat{\mathcal{B}}\mathbf{z}_t)(\mathbf{y}_t - \widehat{\mathcal{B}}\mathbf{z}_t)'. \quad (5.71)$$

The uncertainty in the estimators can be evaluated from

$$\text{se}(\widehat{\beta}_{ij}) = \sqrt{c_{ii}\widehat{\sigma}_{jj}}, \quad (5.72)$$

for $i = 1, \dots, r$, $j = 1, \dots, k$, where se denotes estimated standard error, $\widehat{\sigma}_{jj}$ is the j -th diagonal element of $\widehat{\Sigma}_w$, and c_{ii} is the i -th diagonal element of $(\sum_{t=1}^n \mathbf{z}_t\mathbf{z}_t')^{-1}$.

Also, the information theoretic criterion changes to

$$\text{AIC} = \ln |\widehat{\Sigma}_w| + \frac{2}{n} \left(kr + \frac{k(k+1)}{2} \right). \quad (5.73)$$

and BIC replaces the second term in (5.73) by $K \ln n/n$ where $K = kr + k(k+1)/2$. Bedrick and Tsai (1994) have given a corrected form for AIC in the multivariate case as

$$\text{AICc} = \ln |\widehat{\Sigma}_w| + \frac{k(r+n)}{n-k-r-1}. \quad (5.74)$$

Many data sets involve more than one time series, and we are often interested in the possible dynamics relating all series. In this situation, we are interested in modeling and forecasting $k \times 1$ vector-valued time series $\mathbf{x}_t = (x_{t1}, \dots, x_{tk})'$, $t = 0, \pm 1, \pm 2, \dots$. Unfortunately, extending univariate ARMA models to the multivariate case is not so simple. The multivariate autoregressive model, however, is a straight-forward extension of the univariate AR model.

For the first-order vector autoregressive model, VAR(1), we take

$$\mathbf{x}_t = \boldsymbol{\alpha} + \boldsymbol{\Phi}\mathbf{x}_{t-1} + \mathbf{w}_t, \quad (5.75)$$

where $\boldsymbol{\Phi}$ is a $k \times k$ transition matrix that expresses the dependence of $\mathbf{x}_t$ on $\mathbf{x}_{t-1}$. The vector white noise process $\mathbf{w}_t$ is assumed to be multivariate normal with mean-zero and covariance matrix

$$E(\mathbf{w}_t\mathbf{w}_t') = \boldsymbol{\Sigma}_w. \quad (5.76)$$

The vector $\boldsymbol{\alpha} = (\alpha_1, \alpha_2, \dots, \alpha_k)'$ appears as the constant in the regression setting. If $E(\mathbf{x}_t) = \boldsymbol{\mu}$, then $\boldsymbol{\alpha} = (I - \boldsymbol{\Phi})\boldsymbol{\mu}$.

Note the similarity between the VAR model and the multivariate linear regression model (5.69). The regression formulas carry over, and we can, on observing $\mathbf{x}_1, \dots, \mathbf{x}_n$, set up the model (5.75) with $\mathbf{y}_t = \mathbf{x}_t$, $\mathcal{B} = (\boldsymbol{\alpha}, \boldsymbol{\Phi})$ and $\mathbf{z}_t = (1, \mathbf{x}_{t-1}')'$. Then, write the solution as (5.70) with the conditional maximum likelihood estimator for the covariance matrix given by

$$\widehat{\boldsymbol{\Sigma}}_w = (n-1)^{-1} \sum_{t=2}^n (\mathbf{x}_t - \widehat{\boldsymbol{\alpha}} - \widehat{\boldsymbol{\Phi}}\mathbf{x}_{t-1})(\mathbf{x}_t - \widehat{\boldsymbol{\alpha}} - \widehat{\boldsymbol{\Phi}}\mathbf{x}_{t-1})'. \quad (5.77)$$

The special form assumed for the constant component, $\boldsymbol{\alpha}$, of the vector AR model in (5.75) can be generalized to include a fixed $r \times 1$ vector of inputs, $\mathbf{u}_t$. That is, we could have proposed the vector ARX model,

$$\mathbf{x}_t = \boldsymbol{\Gamma}\mathbf{u}_t + \sum_{j=1}^p \boldsymbol{\Phi}_j\mathbf{x}_{t-j} + \mathbf{w}_t, \quad (5.78)$$

where $\boldsymbol{\Gamma}$ is a $p \times r$ parameter matrix. The X in ARX refers to the exogenous vector process we have denoted here by $\mathbf{u}_t$. The introduction of exogenous variables through replacing $\boldsymbol{\alpha}$ by $\boldsymbol{\Gamma}\mathbf{u}_t$ does not present any special problems in making inferences and we will often drop the X for being superfluous.

Example 5.9 Pollution, Weather, and Mortality

For example, for the three-dimensional series composed of cardiovascular mortality x_{t1} , temperature x_{t2} , and particulate levels x_{t3} , introduced in Example 2.2, take $\mathbf{x}_t = (x_{t1}, x_{t2}, x_{t3})'$ as a vector of dimension $k = 3$. We might envision dynamic relations among the three series defined as the first order relation,

$$x_{t1} = \alpha_1 + \beta_1 t + \phi_{11}x_{t-1,1} + \phi_{12}x_{t-1,2} + \phi_{13}x_{t-1,3} + w_{t1},$$

which expresses the current value of mortality as a linear combination of trend and its immediate past value and the past values of temperature and particulate levels. Similarly,

$$x_{t2} = \alpha_2 + \beta_2 t + \phi_{21}x_{t-1,1} + \phi_{22}x_{t-1,2} + \phi_{23}x_{t-1,3} + w_{t2}$$

and

$$x_{t3} = \alpha_3 + \beta_3 t + \phi_{31}x_{t-1,1} + \phi_{32}x_{t-1,2} + \phi_{33}x_{t-1,3} + w_{t3}$$

express the dependence of temperature and particulate levels on the other series. Of course, methods for the preliminary identification of these models exist, and we will discuss these methods shortly. The model in the form of (5.78) is

$$\mathbf{x}_t = \Gamma \mathbf{u}_t + \Phi \mathbf{x}_{t-1} + \mathbf{w}_t,$$

where, in obvious notation, $\Gamma = [\boldsymbol{\alpha} \mid \boldsymbol{\beta}]$ is 3×2 and $\mathbf{u}_t = (1, t)'$ is 2×1 .

Throughout much of this section we will use the R package `vars` to fit vector AR models via least squares. For this particular example, we have (partial output shown):

```
library(vars)
x = cbind(cmort, tempr, part)
summary(VAR(x, p=1, type="both")) # "both" fits constant + trend
Estimation results for equation cmort: (other equations not shown)
cmort = cmort.l1 + tempr.l1 + part.l1 + const + trend
```

	Estimate	Std. Error	t value	p.value
cmort.l1	0.464824	0.036729	12.656	< 2e-16
tempr.l1	-0.360888	0.032188	-11.212	< 2e-16
part.l1	0.099415	0.019178	5.184	3.16e-07
const	73.227292	4.834004	15.148	< 2e-16
trend	-0.014459	0.001978	-7.308	1.07e-12

```
Residual standard error: 5.583 on 502 degrees of freedom
Multiple R-Squared: 0.6908, Adjusted R-squared: 0.6883
F-statistic: 280.3 on 4 and 502 DF, p-value: < 2.2e-16
```

```
Covariance matrix of residuals:
```

	cmort	tempr	part
cmort	31.172	5.975	16.65
tempr	5.975	40.965	42.32
part	16.654	42.323	144.26

Note that t here is `time(cmort)`, which is $1970 + \Delta(t - 1)$, where $\Delta = 1/52$, for $t = 1, \dots, 508$, ending at 1979.75. For this particular case, we obtain

$$\hat{\boldsymbol{\alpha}} = (73.23, 67.59, 67.46)', \quad \hat{\boldsymbol{\beta}} = (-0.014, -0.007, -0.005)',$$

$$\hat{\Phi} = \begin{pmatrix} .46(.04) & -.36(.03) & .10(.02) \\ -.24(.04) & .49(.04) & -.13(.02) \\ -.12(.08) & -.48(.07) & .58(.04) \end{pmatrix},$$

where the standard errors, computed as in (5.72), are given in parentheses. The estimate of Σ_w seen to be

$$\hat{\Sigma}_w = \begin{pmatrix} 31.17 & 5.98 & 16.65 \\ 5.98 & 40.965 & 42.32 \\ 16.65 & 42.32 & 144.26 \end{pmatrix}.$$

For the vector $(x_{t1}, x_{t2}, x_{t3}) = (M_t, T_t, P_t)$, with M_t, T_t and P_t denoting mortality, temperature, and particulate level, respectively, we obtain the prediction equation for mortality,

$$\widehat{M}_t = 73.23 - .014t + .46M_{t-1} - .36T_{t-1} + .10P_{t-1}.$$

Comparing observed and predicted mortality with this model leads to an R^2 of about .69.

It is easy to extend the VAR(1) process to higher orders, VAR(p). To do this, we use the notation of (5.69) and write the vector of regressors as

$$\mathbf{z}_t = (1, \mathbf{x}'_{t-1}, \mathbf{x}'_{t-2}, \dots, \mathbf{x}'_{t-p})'$$

and the regression matrix as $\mathcal{B} = (\boldsymbol{\alpha}, \Phi_1, \Phi_2, \dots, \Phi_p)$. Then, this regression model can be written as

$$\mathbf{x}_t = \boldsymbol{\alpha} + \sum_{j=1}^p \Phi_j \mathbf{x}_{t-j} + \mathbf{w}_t \quad (5.79)$$

for $t = p + 1, \dots, n$. The $k \times k$ error sum of products matrix becomes

$$SSE = \sum_{t=p+1}^n (\mathbf{x}_t - \mathcal{B}\mathbf{z}_t)(\mathbf{x}_t - \mathcal{B}\mathbf{z}_t)', \quad (5.80)$$

so that the conditional maximum likelihood estimator for the error covariance matrix Σ_w is

$$\widehat{\Sigma}_w = SSE/(n - p), \quad (5.81)$$

as in the multivariate regression case, except now only $n - p$ residuals exist in (5.80). For the multivariate case, we have found that the Schwarz criterion

$$\text{BIC} = \log |\widehat{\Sigma}_w| + k^2 p \ln n/n, \quad (5.82)$$

gives more reasonable classifications than either AIC or corrected version AICc. The result is consistent with those reported in simulations by Lütkepohl (1985).

Example 5.10 Pollution, Weather, and Mortality (cont)

We used the R package first to select a VAR(p) model and then fit the model. The selection criteria used in the package are AIC, Hannan-Quinn (HQ; Hannan & Quinn, 1979), BIC (SC), and Final Prediction Error (FPE). The Hannan-Quinn procedure is similar to BIC, but with $\ln n$ replaced by $2 \ln(\ln(n))$ in the penalty term. FPE finds the model that minimizes the approximate mean squared one-step-ahead prediction error (see Akaike, 1969 for details); it is rarely used.

```
VARselect(x, lag.max=10, type="both")
```

```
$selection
AIC(n)  HQ(n)  SC(n)  FPE(n)
   9      5     2     9

$criteria
      1      2      3      4
AIC(n) 11.73780 11.30185 11.26788 11.23030
HQ(n)  11.78758 11.38149 11.37738 11.36967
SC(n)   11.86463 11.50477 11.54689 11.58541
FPE(n) 125216.91717 80972.28678 78268.19568 75383.73647
```


	5	6	7	8
AIC(n)	11.17634	11.15266	11.15247	11.12878
HQ(n)	11.34557	11.35176	11.38144	11.38760
SC(n)	11.60755	11.65996	11.73587	11.78827
FPE(n)	71426.10041	69758.25113	69749.89175	68122.40518

	9	10
AIC(n)	11.11915	11.12019
HQ(n)	11.40784	11.43874
SC(n)	11.85473	11.93187
FPE(n)	67476.96374	67556.45243

Note that BIC picks the order $p = 2$ model while AIC and FPE pick an order $p = 9$ model and Hannan-Quinn selects an order $p = 5$ model.

Fitting the model selected by BIC we obtain

$$\hat{\alpha} = (56.1, 49.9, 59.6)', \quad \hat{\beta} = (-0.011, -0.005, -0.008)',$$

$$\hat{\Phi}_1 = \begin{pmatrix} .30(.04) & -.20(.04) & .04(.02) \\ -.11(.05) & .26(.05) & -.05(.03) \\ .08(.09) & -.39(.09) & .39(.05) \end{pmatrix},$$

$$\hat{\Phi}_2 = \begin{pmatrix} .28(.04) & -.08(.03) & .07(.03) \\ -.04(.05) & .36(.05) & -.10(.03) \\ -.33(.09) & .05(.09) & .38(.05) \end{pmatrix},$$

where the standard errors are given in parentheses. The estimate of Σ_w is

$$\hat{\Sigma}_w = \begin{pmatrix} 28.03 & 7.08 & 16.33 \\ 7.08 & 37.63 & 40.88 \\ 16.33 & 40.88 & 123.45 \end{pmatrix}.$$

To fit the model using the `vars` package use the following line (partial results displayed):

```
summary(fit <- VAR(x, p=2, type="both"))
cmort = cmort.l1 + tempr.l1 + part.l1 + cmort.l2 + tempr.l2
+ part.l2 + const + trend
```

	Estimate	Std. Error	t value	p.value
cmort.l1	0.297059	0.043734	6.792	3.15e-11
tempr.l1	-0.199510	0.044274	-4.506	8.23e-06
part.l1	0.042523	0.024034	1.769	0.07745
cmort.l2	0.276194	0.041938	6.586	1.15e-10
tempr.l2	-0.079337	0.044679	-1.776	0.07639
part.l2	0.068082	0.025286	2.692	0.00733
const	56.098652	5.916618	9.482	< 2e-16
trend	-0.011042	0.001992	-5.543	4.84e-08

Covariance matrix of residuals:

	cmort	tempr	part
cmort	28.034	7.076	16.33
tempr	7.076	37.627	40.88
part	16.325	40.880	123.45




Fig. 5.12. ACFs (diagonals) and CCFs (off-diagonals) for the residuals of the three-dimensional VAR(2) fit to the LA mortality – pollution data set. On the off-diagonals, the second-named series is the one that leads.

Using the notation of the previous example, the prediction model for cardiovascular mortality is estimated to be

$$\hat{M}_t = 56 - .01t + .3M_{t-1} - .2T_{t-1} + .04P_{t-1} + .28M_{t-2} - .08T_{t-2} + .07P_{t-2}.$$

To examine the residuals, we can plot the cross-correlations of the residuals and examine the multivariate version of the Q-test as follows:

```
acf(resid(fit), 52)
serial.test(fit, lags.pt=12, type="PT.adjusted")
Portmanteau Test (adjusted)
data: Residuals of VAR object fit
Chi-squared = 162.3502, df = 90, p-value = 4.602e-06
```

The cross-correlation matrix is shown in [Figure 5.12](#). The figure shows the ACFs of the individual residual series along the diagonal. For example, the first diagonal graph is the ACF of $M_t - \hat{M}_t$, and so on. The off diagonals display the CCFs between

pairs of residual series. If the title of the off-diagonal plot is x & y , then y leads in the graphic; that is, on the upper-diagonal, the plot shows `corr[x(t+Lag), y(t)]` whereas in the lower-diagonal, if the title is x & y , you get a plot of `corr[x(t+Lag), y(t)]` (yes, it is the same thing, but the lags are negative in the lower diagonal). The graphic is labeled in a strange way, just remember the second named series is the one that leads. In Figure 5.12 we notice that most of the correlations in the residual series are negligible, however, the zero-order correlations of mortality with temperature residuals is about .22 and mortality with particulate residuals is about .28 (type `acf(resid(fit), 52)$acf`) to see the actual values. This means that the AR model is not capturing the concurrent effect of temperature and pollution on mortality (recall the data evolves over a week). It is possible to fit simultaneous models; see Reinsel (1997) for further details. Thus, not unexpectedly, the Q-test rejects the null hypothesis that the noise is white. The Q-test statistic is given by

$$Q = n^2 \sum_{h=1}^H \frac{1}{n-h} \text{tr} [\hat{\Gamma}_w(h) \hat{\Gamma}_w(0)^{-1} \hat{\Gamma}_w(h) \hat{\Gamma}_w(0)^{-1}], \quad (5.83)$$

where

$$\hat{\Gamma}_w(h) = n^{-1} \sum_{t=1}^{n-h} \hat{\mathbf{w}}_{t+h} \hat{\mathbf{w}}_t',$$

and $\hat{\mathbf{w}}_t$ is the residual process. Under the null that $\mathbf{w}_t$ is white noise, (5.83) has an asymptotic χ^2 distribution with $k^2(H-p)$ degrees of freedom.

Finally, prediction follows in a straight forward manner from the univariate case. Using the R package `vars`, use the `predict` command and the `fanchart` command, which produces a nice graphic:

```
(fit.pr = predict(fit, n.ahead = 24, ci = 0.95)) # 4 weeks ahead
fanchart(fit.pr) # plot prediction + error
```

The results are displayed in Figure 5.13; we note that the package stripped time when plotting the fanchart and the horizontal axis is labeled 1, 2, 3, ...

For pure VAR(p) models, the autocovariance structure leads to the multivariate version of the Yule–Walker equations:

$$\Gamma(h) = \sum_{j=1}^p \Phi_j \Gamma(h-j), \quad h = 1, 2, \dots, \quad (5.84)$$

$$\Gamma(0) = \sum_{j=1}^p \Phi_j \Gamma(-j) + \Sigma_w. \quad (5.85)$$

where $\Gamma(h) = \text{cov}(\mathbf{x}_{t+h}, \mathbf{x}_t)$ is a $k \times k$ matrix and $\Gamma(-h) = \Gamma(h)'$.

Estimation of the autocovariance matrix is similar to the univariate case, that is, with $\bar{\mathbf{x}} = n^{-1} \sum_{t=1}^n \mathbf{x}_t$, as an estimate of $\boldsymbol{\mu} = E\mathbf{x}_t$,

$$\hat{\Gamma}(h) = n^{-1} \sum_{t=1}^{n-h} (\mathbf{x}_{t+h} - \bar{\mathbf{x}})(\mathbf{x}_t - \bar{\mathbf{x}})', \quad h = 0, 1, 2, \dots, n-1, \quad (5.86)$$




Fig. 5.13. Predictions from a VAR(2) fit to the LA mortality – pollution data.

and $\widehat{F}(-h) = \widehat{F}(h)'$. If $\widehat{\gamma}_{i,j}(h)$ denotes the element in the i -th row and j -th column of $\widehat{F}(h)$, the cross-correlation functions (CCF), as discussed in (1.35), are estimated by

$$\widehat{\rho}_{i,j}(h) = \frac{\widehat{\gamma}_{i,j}(h)}{\sqrt{\widehat{\gamma}_{i,i}(0)}\sqrt{\widehat{\gamma}_{j,j}(0)}} \quad h = 0, 1, 2, \dots, n-1. \quad (5.87)$$

When $i = j$ in (5.87), we get the estimated autocorrelation function (ACF) of the individual series.

Although least squares estimation was used in [Example 5.9](#) and [Example 5.10](#), we could have also used Yule-Walker estimation, conditional or unconditional maximum likelihood estimation. As in the univariate case, the Yule-Walker estimators, the maximum likelihood estimators, and the least squares estimators are asymptotically equivalent. To exhibit the asymptotic distribution of the autoregression parameter estimators, we write

$$\boldsymbol{\phi} = \text{vec}(\boldsymbol{\Phi}_1, \dots, \boldsymbol{\Phi}_p),$$

where the vec operator stacks the columns of a matrix into a vector. For example, for a bivariate AR(2) model,

$$\boldsymbol{\phi} = \text{vec}(\boldsymbol{\Phi}_1, \boldsymbol{\Phi}_2) = (\boldsymbol{\Phi}_{11}, \boldsymbol{\Phi}_{12}, \boldsymbol{\Phi}_{21}, \boldsymbol{\Phi}_{22})',$$

where $\Phi_{\ell ij}$ is the ij -th element of Φ_ℓ , $\ell = 1, 2$. Because $(\Phi_1, \dots, \Phi_p)$ is a $k \times kp$ matrix, $\boldsymbol{\phi}$ is a $k^2 p \times 1$ vector. We now state the following property.

Property 5.1 Large-Sample Distribution of the Vector Autoregression Estimators

Let $\hat{\boldsymbol{\phi}}$ denote the vector of parameter estimators (obtained via Yule–Walker, least squares, or maximum likelihood) for a k -dimensional AR(p) model. Then,

$$\sqrt{n}(\hat{\boldsymbol{\phi}} - \boldsymbol{\phi}) \sim AN(\mathbf{0}, \Sigma_w \otimes \Gamma_{pp}^{-1}), \quad (5.88)$$

where $\Gamma_{pp} = \{\Gamma(i - j)\}_{i,j=1}^p$ is a $kp \times kp$ matrix and $\Sigma_w \otimes \Gamma_{pp}^{-1} = \{\sigma_{ij} \Gamma_{pp}^{-1}\}_{i,j=1}^k$ is a $k^2 p \times k^2 p$ matrix with σ_{ij} denoting the ij -th element of Σ_w .

The variance–covariance matrix of the estimator $\hat{\boldsymbol{\phi}}$ is approximated by replacing Σ_w by $\hat{\Sigma}_w$, and replacing $\Gamma(h)$ by $\hat{\Gamma}(h)$ in Γ_{pp} . The square root of the diagonal elements of $\hat{\Sigma}_w \otimes \hat{\Gamma}_{pp}^{-1}$ divided by $\sqrt{n}$ gives the individual standard errors. For the mortality data example, the estimated standard errors for the VAR(2) fit are listed in [Example 5.10](#); although those standard errors were taken from a regression run, they could have also been calculated using [Property 5.1](#).

A $k \times 1$ vector-valued time series $\mathbf{x}_t$, for $t = 0, \pm 1, \pm 2, \dots$, is said to be VARMA(p, q) if $\mathbf{x}_t$ is stationary and

$$\mathbf{x}_t = \boldsymbol{\alpha} + \Phi_1 \mathbf{x}_{t-1} + \dots + \Phi_p \mathbf{x}_{t-p} + \mathbf{w}_t + \Theta_1 \mathbf{w}_{t-1} + \dots + \Theta_q \mathbf{w}_{t-q}, \quad (5.89)$$

with $\Phi_p \neq 0$, $\Theta_q \neq 0$, and $\Sigma_w > 0$ (that is, Σ_w is positive definite). The coefficient matrices Φ_j ; $j = 1, \dots, p$ and Θ_j ; $j = 1, \dots, q$ are, of course, $k \times k$ matrices. If $\mathbf{x}_t$ has mean $\boldsymbol{\mu}$ then $\boldsymbol{\alpha} = (I - \Phi_1 - \dots - \Phi_p)\boldsymbol{\mu}$. As in the univariate case, we will have to place a number of conditions on the multivariate ARMA model to ensure the model is unique and has desirable properties such as causality. These conditions will be discussed shortly.

As in the VAR model, the special form assumed for the constant component can be generalized to include a fixed $r \times 1$ vector of inputs, $\mathbf{u}_t$. That is, we could have proposed the vector ARMAX model,

$$\mathbf{x}_t = \Gamma \mathbf{u}_t + \sum_{j=1}^p \Phi_j \mathbf{x}_{t-j} + \sum_{k=1}^q \Theta_k \mathbf{w}_{t-k} + \mathbf{w}_t, \quad (5.90)$$

where Γ is a $p \times r$ parameter matrix.

While extending univariate AR (or pure MA) models to the vector case is fairly easy, extending univariate ARMA models to the multivariate case is not a simple matter. Our discussion will be brief, but interested readers can get more details in Lütkepohl (1993), Reinsel (1997), and Tiao and Tsay (1989).

In the multivariate case, the autoregressive operator is

$$\Phi(B) = I - \Phi_1 B - \dots - \Phi_p B^p, \quad (5.91)$$

and the moving average operator is

$$\Theta(B) = I + \Theta_1 B + \cdots + \Theta_q B^q, \quad (5.92)$$

The zero-mean VARMA(p, q) model is then written in the concise form as

$$\Phi(B)\mathbf{x}_t = \Theta(B)\mathbf{w}_t. \quad (5.93)$$

The model is said to be causal if the roots of $|\Phi(z)|$ (where $|\cdot|$ denotes determinant) are outside the unit circle, $|z| > 1$; that is, $|\Phi(z)| \neq 0$ for any value z such that $|z| \leq 1$. In this case, we can write

$$\mathbf{x}_t = \Psi(B)\mathbf{w}_t,$$

where $\Psi(B) = \sum_{j=0}^{\infty} \Psi_j B^j$, $\Psi_0 = I$, and $\sum_{j=0}^{\infty} \|\Psi_j\| < \infty$. The model is said to be invertible if the roots of $|\Theta(z)|$ lie outside the unit circle. Then, we can write

$$\mathbf{w}_t = \Pi(B)\mathbf{x}_t,$$

where $\Pi(B) = \sum_{j=0}^{\infty} \Pi_j B^j$, $\Pi_0 = I$, and $\sum_{j=0}^{\infty} \|\Pi_j\| < \infty$. Analogous to the univariate case, we can determine the matrices Ψ_j by solving $\Psi(z) = \Phi(z)^{-1}\Theta(z)$, $|z| \leq 1$, and the matrices Π_j by solving $\Pi(z) = \Theta(z)^{-1}\Phi(z)$, $|z| \leq 1$.

For a causal model, we can write $\mathbf{x}_t = \Psi(B)\mathbf{w}_t$ so the general autocovariance structure of an ARMA(p, q) model is

$$\Gamma(h) = \text{cov}(\mathbf{x}_{t+h}, \mathbf{x}_t) = \sum_{j=0}^{\infty} \Psi_{j+h} \Sigma_w \Psi_j'. \quad (5.94)$$

and $\Gamma(-h) = \Gamma(h)'$. For pure MA(q) processes, (5.94) becomes

$$\Gamma(h) = \sum_{j=0}^{q-h} \Theta_{j+h} \Sigma_w \Theta_j', \quad (5.95)$$

where $\Theta_0 = I$. Of course, (5.95) implies $\Gamma(h) = 0$ for $h > q$.

As in the univariate case, we will need conditions for model uniqueness. These conditions are similar to the condition in the univariate case that the autoregressive and moving average polynomials have no common factors. To explore the uniqueness problems that we encounter with multivariate ARMA models, consider a bivariate AR(1) process, $\mathbf{x}_t = (x_{t,1}, x_{t,2})'$, given by

$$\begin{aligned} x_{t,1} &= \phi x_{t-1,2} + w_{t,1}, \\ x_{t,2} &= w_{t,2}, \end{aligned}$$

where $w_{t,1}$ and $w_{t,2}$ are independent white noise processes and $|\phi| < 1$. Both processes, $x_{t,1}$ and $x_{t,2}$ are causal and invertible. Moreover, the processes are jointly stationary because $\text{cov}(x_{t+h,1}, x_{t,2}) = \phi \text{cov}(x_{t+h-1,2}, x_{t,2}) \equiv \phi \gamma_{2,2}(h-1) = \phi \sigma_{w_2}^2 \delta_1^h$ does not depend on t ; note, $\delta_1^h = 1$ when $h = 1$, otherwise, $\delta_1^h = 0$. In matrix notation, we can write this model as

$$\mathbf{x}_t = \Phi \mathbf{x}_{t-1} + \mathbf{w}_t, \quad \text{where} \quad \Phi = \begin{bmatrix} 0 & \phi \\ 0 & 0 \end{bmatrix}. \quad (5.96)$$

We can write (5.96) in operator notation as

$$\Phi(B)\mathbf{x}_t = \mathbf{w}_t \quad \text{where} \quad \Phi(z) = \begin{bmatrix} 1 & -\phi z \\ 0 & 1 \end{bmatrix}.$$

In addition, model (5.96) can be written as a bivariate ARMA(1,1) model

$$\mathbf{x}_t = \Phi_1 \mathbf{x}_{t-1} + \Theta_1 \mathbf{w}_{t-1} + \mathbf{w}_t, \quad (5.97)$$

where

$$\Phi_1 = \begin{bmatrix} 0 & \phi + \theta \\ 0 & 0 \end{bmatrix} \quad \text{and} \quad \Theta_1 = \begin{bmatrix} 0 & -\theta \\ 0 & 0 \end{bmatrix},$$

and θ is arbitrary. To verify this, we write (5.97), as $\Phi_1(B)\mathbf{x}_t = \Theta_1(B)\mathbf{w}_t$, or

$$\Theta_1(B)^{-1}\Phi_1(B)\mathbf{x}_t = \mathbf{w}_t,$$

where

$$\Phi_1(z) = \begin{bmatrix} 1 & -(\phi + \theta)z \\ 0 & 1 \end{bmatrix} \quad \text{and} \quad \Theta_1(z) = \begin{bmatrix} 1 & -\theta z \\ 0 & 1 \end{bmatrix}.$$

Then,

$$\Theta_1(z)^{-1}\Phi_1(z) = \begin{bmatrix} 1 & \theta z \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & -(\phi + \theta)z \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} 1 & -\phi z \\ 0 & 1 \end{bmatrix} = \Phi(z),$$

where $\Phi(z)$ is the polynomial associated with the bivariate AR(1) model in (5.96). Because θ is arbitrary, the parameters of the ARMA(1,1) model given in (5.97) are not identifiable. No problem exists, however, in fitting the AR(1) model given in (5.96).

The problem in the previous discussion was caused by the fact that both $\Theta(B)$ and $\Theta(B)^{-1}$ are finite; such a matrix operator is called unimodular. If $U(B)$ is unimodular, $|U(z)|$ is constant. It is also possible for two seemingly different multivariate ARMA(p, q) models, say, $\Phi(B)\mathbf{x}_t = \Theta(B)\mathbf{w}_t$ and $\Phi_*(B)\mathbf{x}_t = \Theta_*(B)\mathbf{w}_t$, to be related through a unimodular operator, $U(B)$ as $\Phi_*(B) = U(B)\Phi(B)$ and $\Theta_*(B) = U(B)\Theta(B)$, in such a way that the orders of $\Phi(B)$ and $\Theta(B)$ are the same as the orders of $\Phi_*(B)$ and $\Theta_*(B)$, respectively. For example, consider the bivariate ARMA(1,1) models given by

$$\Phi \mathbf{x}_t \equiv \begin{bmatrix} 1 & -\phi B \\ 0 & 1 \end{bmatrix} \mathbf{x}_t = \begin{bmatrix} 1 & \theta B \\ 0 & 1 \end{bmatrix} \mathbf{w}_t \equiv \Theta \mathbf{w}_t$$

and

$$\Phi_*(B)\mathbf{x}_t \equiv \begin{bmatrix} 1 & (\alpha - \phi)B \\ 0 & 1 \end{bmatrix} \mathbf{x}_t = \begin{bmatrix} 1 & (\alpha + \theta)B \\ 0 & 1 \end{bmatrix} \mathbf{w}_t \equiv \Theta_*(B)\mathbf{w}_t,$$

where α , ϕ , and θ are arbitrary constants. Note,

$$\Phi_*(B) \equiv \begin{bmatrix} 1 & (\alpha - \phi)B \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} 1 & \alpha B \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & -\phi B \\ 0 & 1 \end{bmatrix} \equiv U(B)\Phi(B)$$

and

$$\Theta_*(B) \equiv \begin{bmatrix} 1 & (\alpha + \theta)B \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} 1 & \alpha B \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & \theta B \\ 0 & 1 \end{bmatrix} \equiv U(B)\Theta(B).$$

In this case, both models have the same infinite MA representation $\mathbf{x}_t = \Psi(B)\mathbf{w}_t$, where

$$\Psi(B) = \Phi(B)^{-1}\Theta(B) = \Phi(B)^{-1}U(B)^{-1}U(B)\Theta(B) = \Phi_*(B)^{-1}\Theta_*(B).$$

This result implies the two models have the same autocovariance function $\Gamma(h)$. Two such ARMA(p, q) models are said to be observationally equivalent.

As previously mentioned, in addition to requiring causality and invertibility, we will need some additional assumptions in the multivariate case to make sure that the model is unique. To ensure the identifiability of the parameters of the multivariate ARMA(p, q) model, we need the following additional two conditions: (i) the matrix operators $\Phi(B)$ and $\Theta(B)$ have no common left factors other than unimodular ones [that is, if $\Phi(B) = U(B)\Phi_*(B)$ and $\Theta(B) = U(B)\Theta_*(B)$, the common factor must be unimodular] and (ii) with q as small as possible and p as small as possible for that q , the matrix $[\Phi_p, \Theta_q]$ must be full rank, k . One suggestion for avoiding most of the aforementioned problems is to fit only vector AR(p) models in multivariate situations. Although this suggestion might be reasonable for many situations, this philosophy is not in accordance with the law of parsimony because we might have to fit a large number of parameters to describe the dynamics of a process.

Asymptotic inference for the general case of vector ARMA models is more complicated than pure AR models; details can be found in Reinsel (1997) or Lütkepohl (1993), for example. We also note that estimation for VARMA models can be recast into the problem of estimation for state-space models that will be discussed in Chapter 6.

A simple algorithm for fitting vector ARMA models from Spliid (1983) is worth mentioning because it repeatedly uses the multivariate regression equations. Consider a general ARMA(p, q) model for a time series with a nonzero mean

$$\mathbf{x}_t = \boldsymbol{\alpha} + \Phi_1\mathbf{x}_{t-1} + \cdots + \Phi_p\mathbf{x}_{t-p} + \mathbf{w}_t + \Theta_1\mathbf{w}_{t-1} + \cdots + \Theta_q\mathbf{w}_{t-q}. \quad (5.98)$$

If $\boldsymbol{\mu} = E\mathbf{x}_t$, then $\boldsymbol{\alpha} = (I - \Phi_1 - \cdots - \Phi_p)\boldsymbol{\mu}$. If $\mathbf{w}_{t-1}, \dots, \mathbf{w}_{t-q}$ were observed, we could rearrange (5.98) as a multivariate regression model

$$\mathbf{x}_t = \mathcal{B}\mathbf{z}_t + \mathbf{w}_t, \quad (5.99)$$

with

$$\mathbf{z}_t = (1, \mathbf{x}'_{t-1}, \dots, \mathbf{x}'_{t-p}, \mathbf{w}'_{t-1}, \dots, \mathbf{w}'_{t-q})' \quad (5.100)$$

and

$$\mathcal{B} = [\boldsymbol{\alpha}, \Phi_1, \dots, \Phi_p, \Theta_1, \dots, \Theta_q], \quad (5.101)$$

for $t = p + 1, \dots, n$. Given an initial estimator $\mathcal{B}_0$, of $\mathcal{B}$, we can reconstruct $\{\mathbf{w}_{t-1}, \dots, \mathbf{w}_{t-q}\}$ by setting

$$\mathbf{w}_{t-j} = \mathbf{x}_{t-j} - \mathcal{B}_0 \mathbf{z}_{t-j}, \quad t = p + 1, \dots, n, \quad j = 1, \dots, q, \quad (5.102)$$

where, if $q > p$, we put $\mathbf{w}_{t-j} = \mathbf{0}$ for $t - j \leq 0$. The new values of $\{\mathbf{w}_{t-1}, \dots, \mathbf{w}_{t-q}\}$ are then put into the regressors $\mathbf{z}_t$ and a new estimate, say, $\mathcal{B}_1$, is obtained. The initial value, $\mathcal{B}_0$, can be computed by fitting a pure autoregression of order p or higher, and taking $\theta_1 = \dots = \theta_q = 0$. The procedure is then iterated until the parameter estimates stabilize. The algorithm often converges, but not to the maximum likelihood estimators. Experience suggests the estimators can be reasonably close to the maximum likelihood estimators. The algorithm can be considered as a quick and easy way to fit an initial VARMA model as a starting point to using maximum likelihood estimation, which is best done via state-space models covered in the next chapter.

Problems

Section 5.2

5.1 The data set [arf](#) is 1000 simulated observations from an ARFIMA(1, 1, 0) model with $\phi = .75$ and $d = .4$.

- Plot the data and comment.
- Plot the ACF and PACF of the data and comment.
- Estimate the parameters and test for the significance of the estimates $\hat{\phi}$ and $\hat{d}$.
- Explain why, using the results of parts (a) and (b), it would seem reasonable to difference the data prior to the analysis. That is, if x_t represents the data, explain why we might choose to fit an ARMA model to ∇x_t .
- Plot the ACF and PACF of ∇x_t and comment.
- Fit an ARMA model to ∇x_t and comment.

5.2 Compute the sample ACF of the absolute values of the NYSE returns displayed in [Figure 1.4](#) up to lag 200, and comment on whether the ACF indicates long memory. Fit an ARFIMA model to the absolute values and comment.

Section 5.3

5.3 Plot the global temperature series, [gtemp](#), and then test whether there is a unit root versus the alternative that the process is stationary using the three tests, DF, ADF, and PP, discussed in [Example 5.3](#). Comment.

5.4 Plot the GNP series, [gnp](#), and then test for a unit root against the alternative that the process is explosive. State your conclusion.

5.5 Verify (5.33).

Section 5.4

5.6 Weekly crude oil spot prices in dollars per barrel are in [oil](#); see Problem [Problem 2.11](#) and Appendix R for more details. Investigate whether the growth rate of the weekly oil price exhibits GARCH behavior. If so, fit an appropriate model to the growth rate.

5.7 The [stats](#) package of R contains the daily closing prices of four major European stock indices; type [help\(EuStockMarkets\)](#) for details. Fit a GARCH model to the returns of one of these series and discuss your findings. (Note: The data set contains actual values, and not returns. Hence, the data must be transformed prior to the model fitting.)

5.8 The 2×1 gradient vector, $l^{(1)}(\alpha_0, \alpha_1)$, given for an ARCH(1) model was displayed in [\(5.47\)](#). Verify [\(5.47\)](#) and then use the result to calculate the 2×2 Hessian matrix

$$l^{(2)}(\alpha_0, \alpha_1) = \begin{pmatrix} \partial^2 l / \partial \alpha_0^2 & \partial^2 l / \partial \alpha_0 \partial \alpha_1 \\ \partial^2 l / \partial \alpha_0 \partial \alpha_1 & \partial^2 l / \partial \alpha_1^2 \end{pmatrix}.$$

Section 5.5

5.9 The sunspot data ([sunspotz](#)) are plotted in Chapter 4, [Figure 4.22](#). From a time plot of the data, discuss why it is reasonable to fit a threshold model to the data, and then fit a threshold model.

Section 5.6

5.10 The data in [climhyd](#) have 454 months of measured values for the climatic variables air temperature, dew point, cloud cover, wind speed, precipitation (p_t), and inflow (i_t), at Lake Shasta; the data are displayed in [Figure 7.3](#). We would like to look at possible relations between the weather factors and the inflow to Lake Shasta.

- Fit $\text{ARIMA}(0, 0, 0) \times (0, 1, 1)_{12}$ models to (i) transformed precipitation $P_t = \sqrt{p_t}$ and (ii) transformed inflow $I_t = \log i_t$.
- Apply the ARIMA model fitted in part (a) for transformed precipitation to the flow series to generate the prewhitened flow residuals assuming the precipitation model. Compute the cross-correlation between the flow residuals using the precipitation ARIMA model and the precipitation residuals using the precipitation model and interpret. Use the coefficients from the ARIMA model to construct the transformed flow residuals.

5.11 For the [climhyd](#) data set, consider predicting the transformed flows $I_t = \log i_t$ from transformed precipitation values $P_t = \sqrt{p_t}$ using a transfer function model of the form

$$(1 - B^{12})I_t = \alpha(B)(1 - B^{12})P_t + n_t,$$

where we assume that seasonal differencing is a reasonable thing to do. You may think of it as fitting

$$y_t = \alpha(B)x_t + n_t,$$

where y_t and x_t are the seasonally differenced transformed flows and precipitations.

- (a) Argue that x_t can be fitted by a first-order seasonal moving average, and use the transformation obtained to prewhiten the series x_t .
- (b) Apply the transformation applied in (a) to the series y_t , and compute the cross-correlation function relating the prewhitened series to the transformed series. Argue for a transfer function of the form

$$\alpha(B) = \frac{\delta_0}{1 - \omega_1 B}.$$

- (c) Write the overall model obtained in regression form to estimate δ_0 and ω_1 . Note that you will be minimizing the sums of squared residuals for the transformed noise series $(1 - \bar{\omega}_1 B)n_t$. Retain the residuals for further modeling involving the noise n_t . The observed residual is $u_t = (1 - \bar{\omega}_1 B)n_t$.
- (d) Fit the noise residuals obtained in (c) with an ARMA model, and give the final form suggested by your analysis in the previous parts.
- (e) Discuss the problem of forecasting y_{t+m} using the infinite past of y_t and the present and infinite past of x_t . Determine the predicted value and the forecast variance.

Section 5.7

5.12 Consider the data set [econ5](#) containing quarterly U.S. unemployment, GNP, consumption, and government and private investment from 1948-III to 1988-II. The seasonal component has been removed from the data. Concentrating on unemployment (U_t), GNP (G_t), and consumption (C_t), fit a vector ARMA model to the data after first logging each series, and then removing the linear trend. That is, fit a vector ARMA model to $\mathbf{x}_t = (x_{1t}, x_{2t}, x_{3t})'$, where, for example, $x_{1t} = \log(U_t) - \hat{\beta}_0 - \hat{\beta}_1 t$, where $\hat{\beta}_0$ and $\hat{\beta}_1$ are the least squares estimates for the regression of $\log(U_t)$ on time, t . Run a complete set of diagnostics on the residuals.

Chapter 6

State-Space Models

6.1 Introduction

A very general model that seems to subsume a whole class of special cases of interest in much the same way that linear regression does is the state-space model or the dynamic linear model, which was introduced in Kalman (1960) and Kalman and Bucy (1961). Although the model was originally introduced as a method primarily for use in aerospace-related research, it has been applied to modeling data from economics (Harrison and Stevens, 1976; Harvey and Pierse, 1984; Harvey and Todd, 1983; Kitagawa and Gersch 1984, Shumway and Stoffer, 1982), medicine (Jones, 1984) and the soil sciences (Shumway, 1988, §3.4.5). An excellent modern treatment of time series analysis based on the state space model is the text by Durbin and Koopman (2001).

The state-space model or dynamic linear model (DLM), in its basic form, employs an order one, vector autoregression as the state equation,

$$x_t = \Phi x_{t-1} + w_t, \quad (6.1)$$

where the state equation determines the rule for the generation of the $p \times 1$ state vector x_t from the past $p \times 1$ state x_{t-1} , for time points $t = 1, \dots, n$. We assume the w_t are $p \times 1$ independent and identically distributed, zero-mean normal vectors with covariance matrix Q . In the DLM, we assume the process starts with a normal vector x_0 that has mean μ_0 and $p \times p$ covariance matrix Σ_0 .

The DLM, however, adds an additional component to the model in assuming we do not observe the state vector x_t directly, but only a linear transformed version of it with noise added, say

$$y_t = A_t x_t + v_t, \quad (6.2)$$

where A_t is a $q \times p$ measurement or observation matrix; (6.2) is called the observation equation. The model arose originally in the space tracking setting, where the state equation defines the motion equations for the position or state of a spacecraft with location x_t and y_t reflects information that can be observed from a tracking device such as velocity and azimuth. The observed data vector, y_t , is q -dimensional, which




Fig. 6.1. Longitudinal series of blood parameter levels monitored, log (white blood count) [WBC; top], log (platelet) [PLT; middle], and hematocrit [HCT; bottom], after a bone marrow transplant ($n = 91$ days).

can be larger than or smaller than p , the state dimension. The additive observation noise v_t is assumed to be white and Gaussian with $q \times q$ covariance matrix R . In addition, we initially assume, for simplicity, x_0 , $\{w_t\}$ and $\{v_t\}$ are uncorrelated; this assumption is not necessary, but it helps in the explanation of first concepts. The case of correlated errors is discussed in §6.7.

As in the ARMAX model of §5.7, exogenous variables, or fixed inputs, may enter into the states or into the observations. In this case, we suppose we have an $r \times 1$ vector of inputs u_t , and write the model as

$$x_t = \Phi x_{t-1} + \Upsilon u_t + w_t \quad (6.3)$$

$$y_t = A_t x_t + \Gamma u_t + v_t \quad (6.4)$$

where Υ is $p \times r$ and Γ is $q \times r$.

Example 6.1 A Biomedical Example

Suppose we consider the problem of monitoring the level of several biomedical markers after a cancer patient undergoes a bone marrow transplant. The data in Figure 6.1, used by Jones (1984), are measurements made for 91 days on three variables, log(white blood count) [WBC], log(platelet) [PLT], and hematocrit [HCT], denoted y_{t1} , y_{t2} , and y_{t3} , respectively. Approximately 40% of the values are missing, with missing values occurring primarily after the 35th day. The main objectives are to model the three variables using the state-space approach, and to estimate the

missing values. According to Jones, "Platelet count at about 100 days post transplant has previously been shown to be a good indicator of subsequent long term survival." For this particular situation, we model the three variables in terms of the state equation (6.1); that is,

$$\begin{pmatrix} x_{t1} \\ x_{t2} \\ x_{t3} \end{pmatrix} = \begin{pmatrix} \phi_{11} & \phi_{12} & \phi_{13} \\ \phi_{21} & \phi_{22} & \phi_{23} \\ \phi_{31} & \phi_{32} & \phi_{33} \end{pmatrix} \begin{pmatrix} x_{t-1,1} \\ x_{t-1,2} \\ x_{t-1,3} \end{pmatrix} + \begin{pmatrix} w_{t1} \\ w_{t2} \\ w_{t3} \end{pmatrix}. \quad (6.5)$$

The 3×3 observation matrix, A_t , is either the identity matrix, or the identity matrix with all zeros in a row when that variable is missing. The covariance matrices R and Q are 3×3 matrices with $R = \text{diag}\{r_{11}, r_{22}, r_{33}\}$, a diagonal matrix, required for a simple approach when data are missing. The following R code was used to produce Figure 6.1.

```
plot(blood, type="o", pch=19, xlab="day", main="")
```

The model given in (6.1) involving only a single lag is not unduly restrictive. A multivariate model with m lags, such as the VAR(m) discussed in §5.8, could be developed by replacing the $p \times 1$ state vector, x_t , by the $pm \times 1$ state vector $X_t = (x'_t, x'_{t-1}, \dots, x'_{t-m+1})'$ and the transition matrix by

$$\Phi = \begin{pmatrix} \Phi_1 & \Phi_2 & \dots & \Phi_{m-1} & \Phi_m \\ I & 0 & \dots & 0 & 0 \\ 0 & I & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & I & 0 \end{pmatrix}. \quad (6.6)$$

Letting $W_t = (w'_t, 0', \dots, 0')'$ be the new $pm \times 1$ state error vector, the new state equation will be

$$X_t = \Phi X_{t-1} + W_t, \quad (6.7)$$

where $\text{var}(W_t)$ is a $pm \times pm$ matrix with Q in the upper left-hand corner and zeros elsewhere. The observation equation can then be written as

$$y_t = [A_t \mid 0 \mid \dots \mid 0] X_t + v_t. \quad (6.8)$$

This simple recoding shows one way of handling higher order lags within the context of the single lag structure. It is not necessary and often not desirable to have a singular W_t process in the state equation, (6.7). Further discussion of this topic is given in §6.6.

The real advantages of the state space formulation, however, do not really come through in the simple example given above. The special forms that can be developed for various versions of the matrix A_t and for the transition scheme defined by the matrix Φ allow fitting more parsimonious structures with fewer parameters needed to describe a multivariate time series. We will give some examples of structural models in §6.5, but the simple example shown below is instructive.




Fig. 6.2. Annual global temperature deviation series, measured in degrees centigrade, 1880–2009. The solid line is the land-marine series (`gtemp`) whereas the dashed line shows the land-based series (`gtemp2`).

Example 6.2 Global Warming

Figure 6.2 shows two different estimators for the global temperature series from 1880 to 2009. The solid line is `gtemp`, which was considered in the first chapter, and are the global mean land-ocean temperature index data. The second series, `gtemp2`, are the surface air temperature index data using only meteorological station data. Precise details may be obtained from <http://data.giss.nasa.gov/gistemp/graphs/>. Conceptually, both series should be measuring the same underlying climatic signal, and we may consider the problem of extracting this underlying signal. The R code to generate the figure is

```
ts.plot(gtemp, gtemp2, lty=1:2, ylab="Temperature Deviations")
```

We suppose both series are observing the same signal with different noises; that is,

$$y_{t1} = x_t + v_{t1} \quad \text{and} \quad y_{t2} = x_t + v_{t2},$$

or more compactly as

$$\begin{pmatrix} y_{t1} \\ y_{t2} \end{pmatrix} = \begin{pmatrix} 1 \\ 1 \end{pmatrix} x_t + \begin{pmatrix} v_{t1} \\ v_{t2} \end{pmatrix}, \quad (6.9)$$

where

$$R = \text{var} \begin{pmatrix} v_{t1} \\ v_{t2} \end{pmatrix} = \begin{pmatrix} r_{11} & r_{12} \\ r_{21} & r_{22} \end{pmatrix}.$$

It is reasonable to suppose that the unknown common signal, x_t , can be modeled as a random walk with drift of the form

$$x_t = \delta + x_{t-1} + w_t, \quad (6.10)$$

with $Q = \text{var}(w_t)$. In this example, $p = 1$, $q = 2$, $\Phi = 1$, and $\Upsilon = \delta$ with $u_t \equiv 1$.

The introduction of the state-space approach as a tool for modeling data in the social and biological sciences requires model identification and parameter estimation because there is rarely a well-defined differential equation describing the state transition. The questions of general interest for the dynamic linear model (6.3) and (6.4) relate to estimating the unknown parameters contained in $\Phi, \gamma, Q, \Gamma, A_t$, and R , that define the particular model, and estimating or forecasting values of the underlying unobserved process x_t . The advantages of the state-space formulation are in the ease with which we can treat various missing data configurations and in the incredible array of models that can be generated from (6.1) and (6.2). The analogy between the observation matrix A_t and the design matrix in the usual regression and analysis of variance setting is a useful one. We can generate fixed and random effect structures that are either constant or vary over time simply by making appropriate choices for the matrix A_t and the transition structure Φ . We give a few examples in this chapter; for further examples, see Durbin and Koopman (2001), Harvey (1993), Jones (1993), or Shumway (1988) to mention a few.

Before continuing our investigation of the more complex model, it is instructive to consider a simple univariate state-space model wherein an AR(1) process is observed using a noisy instrument.

Example 6.3 An AR(1) Process with Observational Noise

Consider a univariate state-space model where the observations are noisy,

$$y_t = x_t + v_t, \quad (6.11)$$

and the signal (state) is an AR(1) process,

$$x_t = \phi x_{t-1} + w_t, \quad (6.12)$$

for $t = 1, 2, \dots, n$, where $v_t \sim \text{iid } N(0, \sigma_v^2)$, $w_t \sim \text{iid } N(0, \sigma_w^2)$, and $x_0 \sim N(0, \sigma_w^2/(1 - \phi^2))$; $\{v_t\}, \{w_t\}, x_0$ are independent.

In Chapter 3, we investigated the properties of the state, x_t , because it is a stationary AR(1) process (recall Problem 3.2e). For example, we know the autocovariance function of x_t is

$$\gamma_x(h) = \frac{\sigma_w^2}{1 - \phi^2} \phi^h, \quad h = 0, 1, 2, \dots \quad (6.13)$$

But here, we must investigate how the addition of observation noise affects the dynamics. Although it is not a necessary assumption, we have assumed in this example that x_t is stationary. In this case, the observations are also stationary because y_t is the sum of two independent stationary components x_t and v_t . We have

$$\gamma_y(0) = \text{var}(y_t) = \text{var}(x_t + v_t) = \frac{\sigma_w^2}{1 - \phi^2} + \sigma_v^2, \quad (6.14)$$

and, when $h \geq 1$,

$$\gamma_y(h) = \text{cov}(y_t, y_{t-h}) = \text{cov}(x_t + v_t, x_{t-h} + v_{t-h}) = \gamma_x(h). \quad (6.15)$$

Consequently, for $h \geq 1$, the ACF of the observations is

$$\rho_y(h) = \frac{\gamma_y(h)}{\gamma_y(0)} = \left(1 + \frac{\sigma_v^2}{\sigma_w^2}(1 - \phi^2)\right)^{-1} \phi^h. \quad (6.16)$$

It should be clear from the correlation structure given by (6.16) that the observations, y_t , are not AR(1) unless $\sigma_v^2 = 0$. In addition, the autocorrelation structure of y_t is identical to the autocorrelation structure of an ARMA(1,1) process, as presented in Example 3.13. Thus, the observations can also be written in an ARMA(1,1) form,

$$y_t = \phi y_{t-1} + \theta u_{t-1} + u_t,$$

where u_t is Gaussian white noise with variance σ_u^2 , and with θ and σ_u^2 suitably chosen. We leave the specifics of this problem alone for now and defer the discussion to §6.7; in particular, see Example 6.12.

Although an equivalence exists between stationary ARMA models and stationary state-space models (see §6.7), it is sometimes easier to work with one form than another. As previously mentioned, in the case of missing data, complex multivariate systems, mixed effects, and certain types of nonstationarity, it is easier to work in the framework of state-space models; in this chapter, we explore some of these situations.

6.2 Filtering, Smoothing, and Forecasting

From a practical view, the primary aims of any analysis involving the state-space model as defined by (6.1)-(6.2), or by (6.3)-(6.4), would be to produce estimators for the underlying unobserved signal x_t , given the data $\mathcal{Y}_s = \{y_1, \dots, y_s\}$, to time s . When $s < t$, the problem is called forecasting or prediction. When $s = t$, the problem is called filtering, and when $s > t$, the problem is called smoothing. In addition to these estimates, we would also want to measure their precision. The solution to these problems is accomplished via the Kalman filter and smoother and is the focus of this section.

Throughout this chapter, we will use the following definitions:

$$x_t^s = E(x_t \mid \mathcal{Y}_s) \quad (6.17)$$

and

$$P_{t_1, t_2}^s = E \left\{ (x_{t_1} - x_{t_1}^s)(x_{t_2} - x_{t_2}^s)' \right\}. \quad (6.18)$$

When $t_1 = t_2 (= t \text{ say})$ in (6.18), we will write P_t^s for convenience.

In obtaining the filtering and smoothing equations, we will rely heavily on the Gaussian assumption. Some knowledge of the material covered in Appendix B will be helpful in understanding the details of this section (although these details may be skipped on a casual reading of the material). Even in the non-Gaussian case, the

estimators we obtain are the minimum mean-squared error estimators within the class of linear estimators. That is, we can think of E in (6.17) as the projection operator in the sense of §B.1 rather than expectation and $\mathcal{Y}_s$ as the space of linear combinations of $\{y_1, \dots, y_s\}$; in this case, P_t^s is the corresponding mean-squared error. When we assume, as in this section, the processes are Gaussian, (6.18) is also the conditional error covariance; that is,

$$P_{t_1, t_2}^s = E \left\{ (x_{t_1} - x_{t_1}^s)(x_{t_2} - x_{t_2}^s)' \mid \mathcal{Y}_s \right\}.$$

This fact can be seen, for example, by noting the covariance matrix between $(x_t - x_t^s)$ and $\mathcal{Y}_s$, for any t and s , is zero; we could say they are orthogonal in the sense of §B.1. This result implies that $(x_t - x_t^s)$ and $\mathcal{Y}_s$ are independent (because of the normality), and hence, the conditional distribution of $(x_t - x_t^s)$ given $\mathcal{Y}_s$ is the unconditional distribution of $(x_t - x_t^s)$. Derivations of the filtering and smoothing equations from a Bayesian perspective are given in Meinhold and Singpurwalla (1983); more traditional approaches based on the concept of projection and on multivariate normal distribution theory are given in Jazwinski (1970) and Anderson and Moore (1979).

First, we present the Kalman filter, which gives the filtering and forecasting equations. The name filter comes from the fact that x_t^t is a linear filter of the observations $y_1, \dots, y_t$; that is, $x_t^t = \sum_{s=1}^t B_s y_s$ for suitably chosen $p \times q$ matrices B_s . The advantage of the Kalman filter is that it specifies how to update the filter from x_{t-1}^{t-1} to x_t^t once a new observation y_t is obtained, without having to reprocess the entire data set $y_1, \dots, y_t$.

Property 6.1 The Kalman Filter

For the state-space model specified in (6.3) and (6.4), with initial conditions $x_0^0 = \mu_0$ and $P_0^0 = \Sigma_0$, for $t = 1, \dots, n$,

$$x_t^{t-1} = \Phi x_{t-1}^{t-1} + \Upsilon u_t, \quad (6.19)$$

$$P_t^{t-1} = \Phi P_{t-1}^{t-1} \Phi' + Q, \quad (6.20)$$

with

$$x_t^t = x_t^{t-1} + K_t(y_t - A_t x_t^{t-1} - \Gamma u_t), \quad (6.21)$$

$$P_t^t = [I - K_t A_t] P_t^{t-1}, \quad (6.22)$$

where

$$K_t = P_t^{t-1} A_t' [A_t P_t^{t-1} A_t' + R]^{-1} \quad (6.23)$$

is called the Kalman gain. Prediction for $t > n$ is accomplished via (6.19) and (6.20) with initial conditions x_n^n and P_n^n . Important byproducts of the filter are the innovations (prediction errors)

$$\epsilon_t = y_t - E(y_t \mid \mathcal{Y}_{t-1}) = y_t - A_t x_t^{t-1} - \Gamma u_t, \quad (6.24)$$

and the corresponding variance-covariance matrices

$$\Sigma_t \stackrel{\text{def}}{=} \text{var}(\epsilon_t) = \text{var}[A_t(x_t - x_t^{t-1}) + v_t] = A_t P_t^{t-1} A_t' + R \quad (6.25)$$

for $t = 1, \dots, n$

Proof: The derivations of (6.19) and (6.20) follow from straight forward calculations, because from (6.3) we have

$$x_t^{t-1} = E(x_t \mid \mathcal{Y}_{t-1}) = E(\Phi x_{t-1} + \Upsilon u_t + w_t \mid \mathcal{Y}_{t-1}) = \Phi x_{t-1}^{t-1} + \Upsilon u_t,$$

and thus

$$\begin{aligned} P_t^{t-1} &= E \left\{ (x_t - x_t^{t-1})(x_t - x_t^{t-1})' \right\} \\ &= E \left\{ \left[\Phi(x_{t-1} - x_{t-1}^{t-1}) + w_t \right] \left[\Phi(x_{t-1} - x_{t-1}^{t-1}) + w_t \right]' \right\} \\ &= \Phi P_{t-1}^{t-1} \Phi' + Q. \end{aligned}$$

To derive (6.21), we note that $\text{cov}(\epsilon_t, y_s) = 0$ for $s < t$, which in view of the fact the innovation sequence is a Gaussian process, implies that the innovations are independent of the past observations. Furthermore, the conditional covariance between x_t and ϵ_t given $\mathcal{Y}_{t-1}$ is

$$\begin{aligned} \text{cov}(x_t, \epsilon_t \mid \mathcal{Y}_{t-1}) &= \text{cov}(x_t, y_t - A_t x_t^{t-1} - \Gamma u_t \mid \mathcal{Y}_{t-1}) \\ &= \text{cov}(x_t - x_t^{t-1}, y_t - A_t x_t^{t-1} - \Gamma u_t \mid \mathcal{Y}_{t-1}) \\ &= \text{cov}[x_t - x_t^{t-1}, A_t(x_t - x_t^{t-1}) + v_t] \\ &= P_t^{t-1} A_t'. \end{aligned} \tag{6.26}$$

Using these results we have that the joint conditional distribution of x_t and ϵ_t given $\mathcal{Y}_{t-1}$ is normal

$$\begin{pmatrix} x_t \\ \epsilon_t \end{pmatrix} \mid \mathcal{Y}_{t-1} \sim N \left(\begin{bmatrix} x_t^{t-1} \\ 0 \end{bmatrix}, \begin{bmatrix} P_t^{t-1} & P_t^{t-1} A_t' \\ A_t P_t^{t-1} & \Sigma_t \end{bmatrix} \right). \tag{6.27}$$

Thus, using (B.9) of Appendix B, we can write

$$x_t^t = E(x_t \mid y_1, \dots, y_{t-1}, y_t) = E(x_t \mid \mathcal{Y}_{t-1}, \epsilon_t) = x_t^{t-1} + K_t \epsilon_t, \tag{6.28}$$

where

$$K_t = P_t^{t-1} A_t' \Sigma_t^{-1} = P_t^{t-1} A_t' (A_t P_t^{t-1} A_t' + R)^{-1}.$$

The evaluation of P_t^t is easily computed from (6.27) [see (B.10)] as

$$P_t^t = \text{cov}(x_t \mid \mathcal{Y}_{t-1}, \epsilon_t) = P_t^{t-1} - P_t^{t-1} A_t' \Sigma_t^{-1} A_t P_t^{t-1},$$

which simplifies to (6.22). □

Nothing in the proof of **Property 6.1** precludes the cases where some or all of the parameters vary with time, or where the observation dimension changes with time, which leads to the following corollary.

Corollary 6.1 Kalman Filter: The Time-Varying Case

If, in the DLM (6.3) and (6.4), any or all of the parameters are time dependent, $\Phi = \Phi_t$, $\Upsilon = \Upsilon_t$, $Q = Q_t$ in the state equation or $\Gamma = \Gamma_t$, $R = R_t$ in the observation equation, or the dimension of the observational equation is time dependent, $q = q_t$, **Property 6.1** holds with the appropriate substitutions.

Next, we explore the model, prediction, and filtering from a density point of view. To ease the notation, we will drop the inputs from the model. There are two key ingredients to the state space model. Letting $p_{\Theta}(\cdot)$ denote a generic density function with parameters represented by Θ , we have the state process is Markovian:

$$p_{\Theta}(x_t \mid x_{t-1}, x_{t-2}, \dots, x_0) = p_{\Theta}(x_t \mid x_{t-1}), \quad (6.29)$$

and the observations are conditionally independent given the state:

$$p_{\Theta}(y_t \mid x_t, y_1, \dots, y_{t-1}) = p_{\Theta}(y_t \mid x_t), \quad (6.30)$$

In this chapter, we focus primarily on the linear Gaussian model, (6.1)–(6.2), in which case,

$$x_t \mid x_{t-1} \sim N_p(\Phi x_{t-1}, Q) \quad \text{and} \quad y_t \mid x_t \sim N_q(A_t x_t, R)$$

with initial condition $x_0 \sim N_p(\mu_0, \Sigma_0)$. If we let $f_w(\cdot)$ denote the p -variate normal density with mean zero and variance-covariance matrix Q , then

$$p_{\Theta}(x_t \mid x_{t-1}) = f_w(x_t - \Phi x_{t-1}),$$

and if we let $f_v(\cdot)$ denote the q -variate normal density with mean zero and variance-covariance matrix R , then

$$p_{\Theta}(y_t \mid x_t) = f_v(y_t - A_t x_t).$$

In terms of densities, the Kalman filter can be seen as a simple updating scheme, where, to determine the forecast densities, we have

$$\begin{aligned} p_{\Theta}(x_t \mid \mathcal{Y}_{t-1}) &= \int_{R^p} p_{\Theta}(x_t, x_{t-1} \mid \mathcal{Y}_{t-1}) dx_{t-1} \\ &= \int_{R^p} p_{\Theta}(x_t \mid x_{t-1}) p_{\Theta}(x_{t-1} \mid \mathcal{Y}_{t-1}) dx_{t-1} \\ &= \int_{R^p} f_w(x_t - \Phi x_{t-1}) p_{\Theta}(x_{t-1} \mid \mathcal{Y}_{t-1}) dx_{t-1}. \end{aligned} \quad (6.31)$$

Since we seeking an iterative procedure, we introduced x_{t-1} in (6.31) because we have (presumably) previously evaluated the filter $p_{\Theta}(x_{t-1} \mid \mathcal{Y}_{t-1})$. Once we have the predictor, the filter density is obtained as

$$\begin{aligned} p_{\Theta}(x_t \mid \mathcal{Y}_t) &= p_{\Theta}(x_t \mid y_t, \mathcal{Y}_{t-1}) \\ &\propto p_{\Theta}(y_t \mid x_t) p_{\Theta}(x_t \mid \mathcal{Y}_{t-1}), \\ &= f_v(y_t - A_t x_t) p_{\Theta}(x_t \mid \mathcal{Y}_{t-1}). \end{aligned} \quad (6.32)$$

Since we have a linear Gaussian state space model, it is relatively easy to deduce that $p_{\Theta}(x_t \mid \mathcal{Y}_{t-1})$ and $p_{\Theta}(x_t \mid \mathcal{Y}_t)$ are Gaussian using multivariate normal distribution theory. In fact, we know from **Property 6.1** that the densities are $N_p(x_t^{t-1}, P_t^{t-1})$ and $N_p(x_t^t, P_t^t)$, respectively. The following example illustrates these ideas for a simple univariate case.

Example 6.4 Local Level Model

In this example, we suppose that we observe a univariate series y_t that consists of a trend component, μ_t , and a noise component, v_t , where

$$y_t = \mu_t + v_t \quad (6.33)$$

and $v_t \sim \text{iid } N(0, \sigma_v^2)$. In particular, we assume the trend is a random walk given by

$$\mu_t = \mu_{t-1} + w_t \quad (6.34)$$

where $w_t \sim \text{iid } N(0, \sigma_w^2)$ is independent of $\{v_t\}$. Recall **Example 6.2**, where we suggested this type of trend model for the global temperature series.

The model is, of course, a state-space model with (6.33) being the observation equation, and (6.34) being the state equation. We will use the following notation introduced in Blight (1974). Let

$$\{x; \mu, \sigma^2\} = \exp \left\{ -\frac{1}{2\sigma^2}(x - \mu)^2 \right\}, \quad (6.35)$$

then simple manipulation shows

$$\{x; \mu, \sigma^2\} = \{\mu; x, \sigma^2\} \quad (6.36)$$

and

$$\begin{aligned} \{x; \mu_1, \sigma_1^2\} \{x; \mu_2, \sigma_2^2\} &= \left\{ x; \frac{\mu_1/\sigma_1^2 + \mu_2/\sigma_2^2}{1/\sigma_1^2 + 1/\sigma_2^2}, (1/\sigma_1^2 + 1/\sigma_2^2)^{-1} \right\} \\ &\times \left\{ \mu_1; \mu_2, \sigma_1^2 + \sigma_2^2 \right\}. \end{aligned} \quad (6.37)$$

Thus, using (6.31), (6.36) and (6.37) we have

$$\begin{aligned} p(\mu_t \mid \mathcal{Y}_{t-1}) &\propto \int \left\{ \mu_t; \mu_{t-1}, \sigma_w^2 \right\} \left\{ \mu_{t-1}; \mu_{t-1}^{t-1}, P_{t-1}^{t-1} \right\} d\mu_{t-1} \\ &= \int \left\{ \mu_{t-1}; \mu_t, \sigma_w^2 \right\} \left\{ \mu_{t-1}; \mu_{t-1}^{t-1}, P_{t-1}^{t-1} \right\} d\mu_{t-1} \\ &= \left\{ \mu_t; \mu_{t-1}^{t-1}, P_{t-1}^{t-1} + \sigma_w^2 \right\}. \end{aligned} \quad (6.38)$$

From (6.38) we conclude that

$$\mu_t \mid \mathcal{Y}_{t-1} \sim N(\mu_{t-1}^{t-1}, P_{t-1}^{t-1}) \quad (6.39)$$

where

$$\mu_t^{t-1} = \mu_{t-1}^{t-1} \quad \text{and} \quad P_t^{t-1} = P_{t-1}^{t-1} + \sigma_w^2 \quad (6.40)$$

which agrees with the first part of **Property 6.1**.

To derive the filter density using (6.32) and (6.36) we have

$$\begin{aligned}
p(\mu_t \mid \mathcal{Y}_t) &\propto \{y_t; \mu_t, \sigma_v^2\} \{\mu_t; \mu_t^{t-1}, P_t^{t-1}\} \\
&= \{\mu_t; y_t, \sigma_v^2\} \{\mu_t; \mu_t^{t-1}, P_t^{t-1}\}.
\end{aligned} \tag{6.41}$$

An application of (6.37) gives

$$\mu_t \mid \mathcal{Y}_t \sim N(\mu_t^t, P_t^t) \tag{6.42}$$

with

$$\mu_t^t = \frac{\sigma_v^2 \mu_t^{t-1}}{P_t^{t-1} + \sigma_v^2} + \frac{P_t^{t-1} y_t}{P_t^{t-1} + \sigma_v^2} = \mu_t^{t-1} + K_t(y_t - \mu_t^{t-1}), \tag{6.43}$$

where we have defined

$$K_t = \frac{P_t^{t-1}}{P_t^{t-1} + \sigma_v^2}, \tag{6.44}$$

and

$$P_t^t = \left(\frac{1}{\sigma_v^2} + \frac{1}{P_t^{t-1}} \right)^{-1} = \frac{\sigma_v^2 P_t^{t-1}}{P_t^{t-1} + \sigma_v^2} = (1 - K_t) P_t^{t-1}. \tag{6.45}$$

The filter for this specific case, of course, agrees with **Property 6.1**.

Next, we consider the problem of obtaining estimators for x_t based on the entire data sample $y_1, \dots, y_n$, where $t \leq n$, namely, x_t^n . These estimators are called smoothers because a time plot of the sequence $\{x_t^n; t = 1, \dots, n\}$ is typically smoother than the forecasts $\{x_t^{t-1}; t = 1, \dots, n\}$ or the filters $\{x_t^t; t = 1, \dots, n\}$. As is obvious from the above remarks, smoothing implies that each estimated value is a function of the present, future, and past, whereas the filtered estimator depends on the present and past. The forecast depends only on the past, as usual.

Property 6.2 The Kalman Smoother

For the state-space model specified in (6.3) and (6.4), with initial conditions x_n^n and P_n^n obtained via **Property 6.1**, for $t = n, n-1, \dots, 1$,

$$x_{t-1}^n = x_{t-1}^{t-1} + J_{t-1} (x_t^n - x_t^{t-1}), \tag{6.46}$$

$$P_{t-1}^n = P_{t-1}^{t-1} + J_{t-1} (P_t^n - P_t^{t-1}) J_{t-1}', \tag{6.47}$$

where

$$J_{t-1} = P_{t-1}^{t-1} \Phi' [P_t^{t-1}]^{-1}. \tag{6.48}$$

Proof: The smoother can be derived in many ways. Here we provide a proof that was given in Ansley and Kohn (1982). First, for $1 \leq t \leq n$, define

$$\mathcal{Y}_{t-1} = \{y_1, \dots, y_{t-1}\} \quad \text{and} \quad \eta_t = \{v_t, \dots, v_n, w_{t+1}, \dots, w_n\},$$

with $\mathcal{Y}_0$ being empty, and let

$$m_{t-1} = E\{x_{t-1} \mid \mathcal{Y}_{t-1}, x_t - x_t^{t-1}, \eta_t\}.$$

Then, because $\mathcal{Y}_{t-1}$, $\{x_t - x_t^{t-1}\}$, and η_t are mutually independent, and x_{t-1} and η_t are independent, using (B.9) we have

$$m_{t-1} = x_{t-1}^{t-1} + J_{t-1}(x_t - x_t^{t-1}), \quad (6.49)$$

where

$$J_{t-1} = \text{cov}(x_{t-1}, x_t - x_t^{t-1})[P_t^{t-1}]^{-1} = P_{t-1}^{t-1}\Phi'[P_t^{t-1}]^{-1}.$$

Finally, because $\mathcal{Y}_{t-1}$, $x_t - x_t^{t-1}$, and η_t generate $\mathcal{Y}_n = \{y_1, \dots, y_n\}$,

$$x_{t-1}^n = E\{x_{t-1} \mid \mathcal{Y}_n\} = E\{m_{t-1} \mid \mathcal{Y}_n\} = x_{t-1}^{t-1} + J_{t-1}(x_t^n - x_t^{t-1}),$$

which establishes (6.46).

The recursion for the error covariance, P_{t-1}^n , is obtained by straight-forward calculation. Using (6.46) we obtain

$$x_{t-1} - x_{t-1}^n = x_{t-1} - x_{t-1}^{t-1} - J_{t-1}(x_t^n - \Phi x_{t-1}^{t-1}),$$

or

$$(x_{t-1} - x_{t-1}^n) + J_{t-1}x_t^n = (x_{t-1} - x_{t-1}^{t-1}) + J_{t-1}\Phi x_{t-1}^{t-1}. \quad (6.50)$$

Multiplying each side of (6.50) by the transpose of itself and taking expectation, we have

$$P_{t-1}^n + J_{t-1}E(x_t^n x_t^{n'})J_{t-1}' = P_{t-1}^{t-1} + J_{t-1}\Phi E(x_{t-1}^{t-1} x_{t-1}^{t-1'})\Phi' J_{t-1}', \quad (6.51)$$

using the fact the cross-product terms are zero. But,

$$E(x_t^n x_t^{n'}) = E(x_t x_t') - P_t^n = \Phi E(x_{t-1} x_{t-1}')\Phi' + Q - P_t^n,$$

and

$$E(x_{t-1}^{t-1} x_{t-1}^{t-1'}) = E(x_{t-1} x_{t-1}') - P_{t-1}^{t-1},$$

so (6.51) simplifies to (6.47). $\square$

Next, we discuss smoothing from a density point-of-view. Recall that in (6.31) we introduced x_{t-1} because we wanted to develop a forward recursion. Because smoothing involves a backward recursion, we introduce x_{t+1} instead, as follows. For $t = n - 1, \dots, 0$,

$$\begin{aligned} p_{\Theta}(x_t \mid \mathcal{Y}_n) &= \int_{R^p} p_{\Theta}(x_t, x_{t+1} \mid \mathcal{Y}_n) dx_{t+1} \\ &= \int_{R^p} p_{\Theta}(x_{t+1} \mid \mathcal{Y}_n) p_{\Theta}(x_t \mid x_{t+1}, \mathcal{Y}_t) dx_{t+1}. \end{aligned} \quad (6.52)$$

Because we are iterating backward, we already have $p_{\Theta}(x_{t+1} \mid \mathcal{Y}_n)$; e.g., when $t = n - 1$, we evaluated $p_{\Theta}(x_n \mid \mathcal{Y}_n)$ from the filter, and so on. Next,

$$p_{\Theta}(x_t \mid x_{t+1}, \mathcal{Y}_t) = p_{\Theta}(x_t \mid \mathcal{Y}_t) \frac{p_{\Theta}(x_{t+1} \mid x_t)}{p_{\Theta}(x_{t+1} \mid \mathcal{Y}_t)},$$




Fig. 6.3. Displays for [Example 6.5](#). The simulated values of μ_t , for $t = 1, \dots, 50$, given by (6.54) are shown as points. *Top:* The predictions μ_t^{t-1} obtained via the Kalman filter are shown as a line. Error bounds, $\mu_t^{t-1} \pm 2\sqrt{P_t^{t-1}}$, are shown as dashed lines. *Middle:* The filter μ_t^t obtained via the Kalman filter are shown as a line. Error bounds, $\mu_t^t \pm 2\sqrt{P_t^t}$, are shown as dashed lines. *Bottom:* The smoothers μ_t^n obtained via the Kalman smoother are shown as a line. Error bounds, $\mu_t^n \pm 2\sqrt{P_t^n}$, are shown as dashed lines.

and the values on the right-hand-side are all known quantities obtained from filtering (forward recursion) and the model. Hence (6.52) can be written as

$$p_{\theta}(x_t | \mathcal{Y}_n) = p_{\theta}(x_t | \mathcal{Y}_t) \int_{R^p} p_{\theta}(x_{t+1} | \mathcal{Y}_n) \frac{p_{\theta}(x_{t+1} | x_t)}{p_{\theta}(x_{t+1} | \mathcal{Y}_t)} dx_{t+1}. \quad (6.53)$$

As in the prediction and filtering cases, (6.31)-(6.32), because we are dealing with the linear Gaussian state space model, (6.53) will simplify to a normal density; i.e., $x_t | \mathcal{Y}_n \sim N_p(x_t^n, P_t^n)$ with values given in [Property 6.2](#).

Example 6.5 Prediction, Filtering and Smoothing for the Local Level Model

For this example, we simulated $n = 50$ observations from the local level trend model discussed in [Example 6.4](#). We generated a random walk

$$\mu_t = \mu_{t-1} + w_t \quad (6.54)$$

with $w_t \sim \text{iid } N(0, 1)$ and $\mu_0 \sim N(0, 1)$. We then supposed that we observe a univariate series y_t consisting of the trend component, μ_t , and a noise component,

Table 6.1. First 10 Observations of [Example 6.5](#)

t	y_t	μ_t	μ_t^{t-1}	P_t^{t-1}	μ_t^t	P_t^t	μ_t^n	P_t^n
0	—	-.63	—	—	.00	1.00	-.32	.62
1	-1.05	-.44	.00	2.00	-.70	.67	-.65	.47
2	-.94	-1.28	-.70	1.67	-.85	.63	-.57	.45
3	-.81	.32	-.85	1.63	-.83	.62	-.11	.45
4	2.08	.65	-.83	1.62	.97	.62	1.04	.45
5	1.81	-.17	.97	1.62	1.49	.62	1.16	.45
6	-.05	.31	1.49	1.62	.53	.62	.63	.45
7	.01	1.05	.53	1.62	.21	.62	.78	.45
8	2.20	1.63	.21	1.62	1.44	.62	1.70	.45
9	1.19	1.32	1.44	1.62	1.28	.62	2.12	.45
10	5.24	2.83	1.28	1.62	3.73	.62	3.48	.45

$v_t \sim \text{iid } N(0, 1)$, where

$$y_t = \mu_t + v_t. \quad (6.55)$$

The sequences $\{w_t\}$, $\{v_t\}$ and μ_0 were generated independently. We then ran the Kalman filter and smoother, [Property 6.1](#) and [Property 6.2](#), using the actual parameters. The top panel of [Figure 6.3](#) shows the actual values of μ_t as points, and the predictions μ_t^{t-1} , for $t = 1, 2, \dots, 50$, superimposed on the graph as a line. In addition, we display $\mu_t^{t-1} \pm 2\sqrt{P_t^{t-1}}$ as dashed lines on the plot. The middle panel displays the filter, μ_t^t , for $t = 1, \dots, 50$, as a line with $\mu_t^t \pm 2\sqrt{P_t^t}$ as dashed lines. The bottom panel of [Figure 6.3](#) shows a similar plot for the smoother μ_t^n .

[Table 6.1](#) shows the first 10 observations as well as the corresponding state values, the predictions, filters and smoothers. Note that one-step-ahead prediction is more uncertain than the corresponding filtered value, which, in turn, is more uncertain than the corresponding smoother value (that is $P_t^{t-1} \geq P_t^t \geq P_t^n$). Also, in each case, the error variances stabilize quickly.

The R code for this example is as follows. In the example we use `Ksmooth0`, which calls `Kfilter0` for the filtering part (see Appendix R). In the returned values from `Ksmooth0`, the letters `p`, `f`, `s` denote prediction, filter, and smooth, respectively (e.g., `xp` is x_t^{t-1} , `xf` is x_t^t , `xs` is x_t^n , and so on). These scripts use a Cholesky-type decomposition¹ of Q and R ; they are denoted by `cQ` and `cR`. Practically, the scripts only require that Q or R may be reconstructed as `t(cQ)%*(cQ)` or `t(cR)%*(cR)`, respectively, which allows more flexibility. For example, the model (6.7) - (6.8) does not pose a problem even though the state noise covariance matrix is not positive definite.

```
# generate data
set.seed(1); num = 50
w = rnorm(num+1,0,1); v = rnorm(num,0,1)
```

¹ Given a positive definite matrix A , its Cholesky decomposition is an upper triangular matrix U with strictly positive diagonal entries such that $A = U'U$. In R, use `chol(A)`. For the univariate case, it is simply the positive square root of A .

```

mu = cumsum(w) # state: mu[0], mu[1], ..., mu[50]
y = mu[-1] + v # obs: y[1], ..., y[50]
# filter and smooth
mu0 = 0; sigma0 = 1; phi = 1; cQ = 1; cR = 1
ks = Ksmooth0(num, y, 1, mu0, sigma0, phi, cQ, cR)
# start figure
par(mfrow=c(3,1)); Time = 1:num
plot(Time, mu[-1], main="Prediction", ylim=c(-5,10))
  lines(ks$xp)
  lines(ks$xp+2*sqrt(ks$Pp), lty="dashed", col="blue")
  lines(ks$xp-2*sqrt(ks$Pp), lty="dashed", col="blue")
plot(Time, mu[-1], main="Filter", ylim=c(-5,10))
  lines(ks$xf)
  lines(ks$xf+2*sqrt(ks$Pf), lty="dashed", col="blue")
  lines(ks$xf-2*sqrt(ks$Pf), lty="dashed", col="blue")
plot(Time, mu[-1], main="Smoother", ylim=c(-5,10))
  lines(ks$xs)
  lines(ks$xs+2*sqrt(ks$Ps), lty="dashed", col="blue")
  lines(ks$xs-2*sqrt(ks$Ps), lty="dashed", col="blue")
mu[1]; ks$x0n; sqrt(ks$P0n) # initial value info

```

When we discuss maximum likelihood estimation via the EM algorithm in the next section, we will need a set of recursions for obtaining $P_{t,t-1}^n$, as defined in (6.18). We give the necessary recursions in the following property.

Property 6.3 The Lag-One Covariance Smoother

For the state-space model specified in (6.3) and (6.4), with K_t , J_t ($t = 1, \dots, n$), and P_n^n obtained from *Property 6.1* and *Property 6.2*, and with initial condition

$$P_{n,n-1}^n = (I - K_n A_n) \Phi P_{n-1}^{n-1}, \quad (6.56)$$

for $t = n, n-1, \dots, 2$,

$$P_{t-1,t-2}^n = P_{t-1}^{t-1} J'_{t-2} + J_{t-1} \left(P_{t,t-1}^n - \Phi P_{t-1}^{t-1} \right) J'_{t-2}. \quad (6.57)$$

Proof: Because we are computing covariances, we may assume $u_t \equiv 0$ without loss of generality. To derive the initial term (6.56), we first define

$$\tilde{x}_t^s = x_t - x_t^s.$$

Then, using (6.21) and (6.46), we write

$$\begin{aligned}
 P_{t,t-1}^t &= E \left(\tilde{x}_t^t \tilde{x}_{t-1}^{t'} \right) \\
 &= E \left\{ [\tilde{x}_t^{t-1} - K_t(y_t - A_t x_{t-1}^{t-1})][\tilde{x}_{t-1}^{t-1} - J_{t-1} K_t(y_t - A_t x_{t-1}^{t-1})]' \right\} \\
 &= E \left\{ [\tilde{x}_t^{t-1} - K_t(A_t \tilde{x}_{t-1}^{t-1} + v_t)][\tilde{x}_{t-1}^{t-1} - J_{t-1} K_t(A_t \tilde{x}_{t-1}^{t-1} + v_t)]' \right\}.
 \end{aligned}$$

Expanding terms and taking expectation, we arrive at

$$P_{t,t-1}^t = P_{t,t-1}^{t-1} - P_{t-1}^{t-1} A_t' K_t' J_{t-1}' - K_t A_t P_{t,t-1}^{t-1} + K_t (A_t P_{t-1}^{t-1} A_t' + R) K_t' J_{t-1}',$$

noting $E(\tilde{x}_t^{t-1} v_t') = 0$. The final simplification occurs by realizing that $K_t(A_t P_{t-1}^{t-1} A_t' + R) = P_t^{t-1} A_t'$, and $P_{t,t-1}^{t-1} = \Phi P_{t-1}^{t-1}$. These relationships hold for any $t = 1, \dots, n$, and (6.56) is the case $t = n$.

We give the basic steps in the derivation of (6.57). The first step is to use (6.46) to write

$$\tilde{x}_{t-1}^n + J_{t-1} x_t^n = \tilde{x}_{t-1}^{t-1} + J_{t-1} \Phi x_{t-1}^{t-1} \quad (6.58)$$

and

$$\tilde{x}_{t-2}^n + J_{t-2} x_{t-1}^n = \tilde{x}_{t-2}^{t-2} + J_{t-2} \Phi x_{t-2}^{t-2}. \quad (6.59)$$

Next, multiply the left-hand side of (6.58) by the transpose of the left-hand side of (6.59), and equate that to the corresponding result of the right-hand sides of (6.58) and (6.59). Then, taking expectation of both sides, the left-hand side result reduces to

$$P_{t-1,t-2}^n + J_{t-1} E(x_t^n x_{t-1}^{n'}) J_{t-2}' \quad (6.60)$$

and the right-hand side result reduces to

$$\begin{aligned} P_{t-1,t-2}^{t-2} - K_{t-1} A_{t-1} P_{t-1,t-2}^{t-2} + J_{t-1} \Phi K_{t-1} A_{t-1} P_{t-1,t-2}^{t-2} \\ + J_{t-1} \Phi E(x_{t-1}^{t-1} x_{t-2}^{t-2'}) \Phi' J_{t-2}'. \end{aligned} \quad (6.61)$$

In (6.60), write

$$E(x_t^n x_{t-1}^{n'}) = E(x_t x_{t-1}') - P_{t,t-1}^n = \Phi E(x_{t-1} x_{t-2}') \Phi' + \Phi Q - P_{t,t-1}^n,$$

and in (6.61), write

$$E(x_{t-1}^{t-1} x_{t-2}^{t-2'}) = E(x_{t-1}^{t-2} x_{t-2}^{t-2'}) = E(x_{t-1} x_{t-2}') - P_{t-1,t-2}^{t-2}.$$

Equating (6.60) to (6.61) using these relationships and simplifying the result leads to (6.57). $\square$

6.3 Maximum Likelihood Estimation

The estimation of the parameters that specify the state space model, (6.3) and (6.4), is quite involved. We use θ to represent the vector of unknown parameters in the initial mean and covariance μ_0 and Σ_0 , the transition matrix Φ , and the state and observation covariance matrices Q and R and the input coefficient matrices, Υ and Γ . We use maximum likelihood under the assumption that the initial state is normal, $x_0 \sim N_p(\mu_0, \Sigma_0)$, and the errors are normal, $w_t \sim \text{iid } N_p(0, Q)$ and $v_t \sim \text{iid } N_q(0, R)$. We continue to assume, for simplicity, $\{w_t\}$ and $\{v_t\}$ are uncorrelated.

The likelihood is computed using the innovations $\epsilon_1, \epsilon_2, \dots, \epsilon_n$, defined by (6.24),

$$\epsilon_t = y_t - A_t x_t^{t-1} - \Gamma u_t.$$

The innovations form of the likelihood function, which was first given by Schweppe (1965), is obtained using an argument similar to the one leading to (3.116) and proceeds by

noting the innovations are independent Gaussian random vectors with zero means and, as shown in (6.25), covariance matrices

$$\Sigma_t = A_t P_t^{t-1} A_t' + R. \quad (6.62)$$

Hence, ignoring a constant, we may write the likelihood, $L_Y(\theta)$, as

$$-\ln L_Y(\theta) = \frac{1}{2} \sum_{t=1}^n \log |\Sigma_t(\theta)| + \frac{1}{2} \sum_{t=1}^n \epsilon_t(\theta)' \Sigma_t(\theta)^{-1} \epsilon_t(\theta), \quad (6.63)$$

where we have emphasized the dependence of the innovations on the parameters θ . Of course, (6.63) is a highly nonlinear and complicated function of the unknown parameters. The usual procedure is to fix x_0 and then develop a set of recursions for the log likelihood function and its first two derivatives (for example, Gupta and Mehra, 1974). Then, a Newton–Raphson algorithm (see 3.29) can be used successively to update the parameter values until the negative of the log likelihood is minimized. This approach is advocated, for example, by Jones (1980), who developed ARMA estimation by putting the ARMA model in state-space form. For the univariate case, (6.63) is identical, in form, to the likelihood for the ARMA model given in (3.116).

The steps involved in performing a Newton–Raphson estimation procedure are as follows.

- (i) Select initial values for the parameters, say, $\theta^{(0)}$.
- (ii) Run the Kalman filter, **Property 6.1**, using the initial parameter values, $\theta^{(0)}$, to obtain a set of innovations and error covariances, say, $\{\epsilon_t^{(0)}; t = 1, \dots, n\}$ and $\{\Sigma_t^{(0)}; t = 1, \dots, n\}$.
- (iii) Run one iteration of a Newton–Raphson procedure with $-\ln L_Y(\theta)$ as the criterion function (refer to **Example 3.29** for details), to obtain a new set of estimates, say $\theta^{(1)}$.
- (iv) At iteration j , ($j = 1, 2, \dots$), repeat step 2 using $\theta^{(j)}$ in place of $\theta^{(j-1)}$ to obtain a new set of innovation values $\{\epsilon_t^{(j)}; t = 1, \dots, n\}$ and $\{\Sigma_t^{(j)}; t = 1, \dots, n\}$. Then repeat step 3 to obtain a new estimate $\theta^{(j+1)}$. Stop when the estimates or the likelihood stabilize; for example, stop when the values of $\theta^{(j+1)}$ differ from $\theta^{(j)}$, or when $L_Y(\theta^{(j+1)})$ differs from $L_Y(\theta^{(j)})$, by some predetermined, but small amount.

Example 6.6 Newton–Raphson for Example 6.3

In this example, we generated $n = 100$ observations, $y_1, \dots, y_{100}$, using the model in **Example 6.3**, to perform a Newton–Raphson estimation of the parameters ϕ , σ_w^2 , and σ_v^2 . In the notation of §6.2, we would have $\Phi = \phi$, $Q = \sigma_w^2$ and $R = \sigma_v^2$. The actual values of the parameters are $\phi = .8$, $\sigma_w^2 = \sigma_v^2 = 1$.

In the simple case of an AR(1) with observational noise, initial estimation can be accomplished using the results of **Example 6.3**. For example, using (6.16), we set

$$\phi^{(0)} = \widehat{\rho}_y(2)/\widehat{\rho}_y(1).$$

Similarly, from (6.15), $\gamma_x(1) = \gamma_y(1) = \phi\sigma_w^2/(1 - \phi^2)$, so that, initially, we set

$$\sigma_w^{2(0)} = (1 - \phi^{(0)2})\widehat{\gamma}_y(1)/\phi^{(0)}.$$

Finally, using (6.14) we obtain an initial estimate of σ_v^2 , namely,

$$\sigma_v^{2(0)} = \widehat{\gamma}_y(0) - [\sigma_w^{2(0)}/(1 - \phi^{(0)2})].$$

Newton–Raphson estimation was accomplished using the R program `optim`. The code used for this example is given below. In that program, we must provide an evaluation of the function to be minimized, namely, $-\ln L_Y(\theta)$. In this case, the function call combines steps 2 and 3, using the current values of the parameters, $\theta^{(j-1)}$, to obtain first the filtered values, then the innovation values, and then calculating the criterion function, $-\ln L_Y(\theta^{(j-1)})$, to be minimized. We can also provide analytic forms of the gradient or score vector, $-\partial \ln L_Y(\theta)/\partial \theta$, and the Hessian matrix, $-\partial^2 \ln L_Y(\theta)/\partial \theta \partial \theta'$, in the optimization routine, or allow the program to calculate these values numerically. In this example, we let the program proceed numerically and we note the need to be cautious when calculating gradients numerically. For better stability, we can also provide an iterative solution for obtaining analytic gradients and Hessians of the log likelihood function; for details, see Problem 6.9 and Problem 6.10 and Gupta and Mehra (1974).

```
# Generate Data
set.seed(999); num = 100; N = num+1
x = arima.sim(n=N, list(ar = .8, sd=1))
y = ts(x[-1] + rnorm(num,0,1))
# Initial Estimates
u = ts.intersect(y, lag(y,-1), lag(y,-2))
varu = var(u); coru = cor(u)
phi = coru[1,3]/coru[1,2]
q = (1-phi^2)*varu[1,2]/phi; r = varu[1,1] - q/(1-phi^2)
(init.par = c(phi, sqrt(q), sqrt(r))) # = .91, .51, 1.03
# Function to evaluate the likelihood
Linn=function(para){
  phi = para[1]; sigw = para[2]; sigv = para[3]
  Sigma0 = (sigw^2)/(1-phi^2); Sigma0[Sigma0<0]=0
  kf = Kfilter0(num, y, 1, mu0=0, Sigma0, phi, sigw, sigv)
  return(kf$like) }
# Estimation (partial output shown)
(est = optim(init.par, Linn, gr=NULL, method="BFGS", hessian=TRUE,
  control=list(trace=1, REPORT=1)))
SE = sqrt(diag(solve(est$hessian)))
cbind(estimate=c(phi=est$par[1],sigw=est$par[2],sigv=est$par[3]),SE)
      estimate      SE
phi  0.8137623 0.08060636
sigw 0.8507863 0.17528895
sigv 0.8743968 0.14293192
```

As seen from the output, the final estimates, along with their standard errors (in parentheses), are $\widehat{\phi} = .81 (.08)$, $\widehat{\sigma}_w = .85 (.18)$, $\widehat{\sigma}_v = .87 (.14)$. The report from `optim` yielded the following results of the estimation procedure:

```
initial value 79.601468
iter    2 value 79.060391
```




Fig. 6.4. Plot for **Example 6.7**. The dashed lines are the two average global temperature deviations shown in **Figure 6.2**. The solid line is the estimated smoother $\hat{x}_t^n$, and the corresponding two root mean square error bounds is the gray area.

```

iter   3 value 79.034121
iter   4 value 79.032615
iter   5 value 79.014817
iter   6 value 79.014453
final  value 79.014452

```

Note that the algorithm converged in six (or seven?) steps with the final value of the negative of the log likelihood being 79.014452. The standard errors are a byproduct of the estimation procedure, and we will discuss their evaluation later in this section, after **Property 6.4**.

Example 6.7 Newton–Raphson for **Example 6.2**

In **Example 6.2**, we considered two different global temperature series of $n = 130$ observations each, and they are plotted in **Figure 6.2**. In that example, we argued that both series should be measuring the same underlying climatic signal, x_t , which we model as a random walk with drift,

$$x_t = \delta + x_{t-1} + w_t.$$

Recall that the observation equation was written as

$$\begin{pmatrix} y_{t1} \\ y_{t2} \end{pmatrix} = \begin{pmatrix} 1 \\ 1 \end{pmatrix} x_t + \begin{pmatrix} v_{t1} \\ v_{t2} \end{pmatrix},$$

and the model covariance matrices are given by $Q = q_{11}$ and

$$R = \begin{pmatrix} r_{11} & r_{12} \\ r_{21} & r_{22} \end{pmatrix}.$$

Hence, there are five parameters to estimate, δ , the drift, and the variance components, $q_{11}, r_{11}, r_{12}, r_{22}$, noting that $r_{21} = r_{12}$. We hold the the initial state parameters fixed in this example at $\mu_0 = -.26$ and $\Sigma_0 = .01$.

The final estimates are $\hat{\delta} = .006, \hat{q}_{11} = .033, \hat{r}_{11} = .007, \hat{r}_{12} = .01, \hat{r}_{22} = .02$, with all values being significant. The observations and the smoothed estimate of the signal, $\hat{x}_t^n$, are displayed in Figure 6.4. The R code, which uses `Kfilter1` and `Ksmooth1`, is as follows.

```
# Setup
y = cbind(gtemp,gtemp2); num = nrow(y); input = rep(1,num)
A = array(rep(1,2), dim=c(2,1,num))
mu0 = -.26; Sigma0 = .01; Phi = 1
# Function to Calculate Likelihood
Linn=function(para){
  cQ = para[1]      # sigma_w
  cR1 = para[2]     # 11 element of chol(R)
  cR2 = para[3]     # 22 element of chol(R)
  cR12 = para[4]    # 12 element of chol(R)
  cR = matrix(c(cR1,0,cR12,cR2),2) # put the matrix together
  drift = para[5]
  kf = Kfilter1(num,y,A,mu0,Sigma0,Phi,drift,0,cQ,cR,input)
  return(kf$like) }
# Estimation
init.par = c(.1,.1,.1,0,.05) # initial values of parameters
(est = optim(init.par, Linn, NULL, method="BFGS", hessian=TRUE,
  control=list(trace=1,REPORT=1))) # output not shown
SE = sqrt(diag(solve(est$hessian)))
# display estimates
u = cbind(estimate=est$par, SE)
rownames(u)=c("sigw", "cR11", "cR22", "cR12", "drift"); u
      estimate      SE
sigw  0.032730315 0.007473594
cR11  0.084752492 0.007815219
cR22  0.070864957 0.005732578
cR12  0.122458872 0.014867006
drift  0.005852047 0.002919058
# Smooth (first set parameters to their final estimates)
cQ=est$par[1]
cR1=est$par[2]
cR2=est$par[3]
cR12=est$par[4]
cR = matrix(c(cR1,0,cR12,cR2), 2)
(R = t(cR)%*%cR) # to view the estimated R matrix
drift = est$par[5]
ks = Ksmooth1(num,y,A,mu0,Sigma0,Phi,drift,0,cQ,cR,input)
# Plot
xsm = ts(as.vector(ks$xs), start=1880)
rmse = ts(sqrt(as.vector(ks$Ps)), start=1880)
plot(xsm, ylim=c(-.5,.8), ylab="Temperature Deviations")
xx = c(time(xsm), rev(time(xsm)))
yy = c(xsm-2*rmse, rev(xsm+2*rmse))
polygon(xx, yy, border=NA, col=rgb(.6,.6,.6,alpha=.25))
lines(gtemp, type='o', pch=2, col=rgb(0,0,.8), lty=2)
lines(gtemp2, type='o', pch=3, col=rgb(0,.65,0), lty=2)
```

In addition to Newton–Raphson, Shumway and Stoffer (1982) presented a conceptually simpler estimation procedure based on the EM (expectation-maximization) algorithm (Dempster et al., 1977). For the sake of brevity, we ignore the inputs and consider the model in the form of (6.1) and (6.2). The basic idea is that if we could observe the states, $\mathcal{X}_n = \{x_0, x_1, \dots, x_n\}$, in addition to the observations $\mathcal{Y}_n = \{y_1, \dots, y_n\}$, then we would consider $\{\mathcal{X}_n, \mathcal{Y}_n\}$ as the complete data, with the joint density

$$f_{\boldsymbol{\theta}}(\mathcal{X}_n, \mathcal{Y}_n) = f_{\mu_0, \Sigma_0}(x_0) \prod_{t=1}^n f_{\Phi, Q}(x_t | x_{t-1}) \prod_{t=1}^n f_R(y_t | x_t). \quad (6.64)$$

Under the Gaussian assumption and ignoring constants, the complete data likelihood, (6.64), can be written as

$$\begin{aligned} -2 \ln L_{X,Y}(\boldsymbol{\theta}) &= \ln |\Sigma_0| + (x_0 - \mu_0)' \Sigma_0^{-1} (x_0 - \mu_0) \\ &\quad + n \ln |Q| + \sum_{t=1}^n (x_t - \Phi x_{t-1})' Q^{-1} (x_t - \Phi x_{t-1}) \\ &\quad + n \ln |R| + \sum_{t=1}^n (y_t - A_t x_t)' R^{-1} (y_t - A_t x_t). \end{aligned} \quad (6.65)$$

Thus, in view of (6.65), if we did have the complete data, we could then use the results from multivariate normal theory to easily obtain the MLEs of $\boldsymbol{\theta}$. We do not have the complete data; however, the EM algorithm gives us an iterative method for finding the MLEs of $\boldsymbol{\theta}$ based on the incomplete data, $\mathcal{Y}_n$, by successively maximizing the conditional expectation of the complete data likelihood. To implement the EM algorithm, we write, at iteration j , ($j = 1, 2, \dots$),

$$Q(\boldsymbol{\theta} \mid \boldsymbol{\theta}^{(j-1)}) = E \left\{ -2 \ln L_{X,Y}(\boldsymbol{\theta}) \mid \mathcal{Y}_n, \boldsymbol{\theta}^{(j-1)} \right\}. \quad (6.66)$$

Calculation of (6.66) is the expectation step. Of course, given the current value of the parameters, $\boldsymbol{\theta}^{(j-1)}$, we can use **Property 6.2** to obtain the desired conditional expectations as smoothers. This property yields

$$\begin{aligned} Q(\boldsymbol{\theta} \mid \boldsymbol{\theta}^{(j-1)}) &= \ln |\Sigma_0| + \text{tr} \left\{ \Sigma_0^{-1} [P_0^n + (x_0^n - \mu_0)(x_0^n - \mu_0)'] \right\} \\ &\quad + n \ln |Q| + \text{tr} \left\{ Q^{-1} [S_{11} - S_{10} \Phi' - \Phi S_{10}' + \Phi S_{00} \Phi'] \right\} \\ &\quad + n \ln |R| \\ &\quad + \text{tr} \left\{ R^{-1} \sum_{t=1}^n [(y_t - A_t x_t^n)(y_t - A_t x_t^n)' + A_t P_t^n A_t'] \right\}, \end{aligned} \quad (6.67)$$

where

$$S_{11} = \sum_{t=1}^n (x_t^n x_t^{n'} + P_t^n), \quad (6.68)$$

$$S_{10} = \sum_{t=1}^n (x_t^n x_{t-1}^{n'} + P_{t,t-1}^n), \quad (6.69)$$

and

$$S_{00} = \sum_{t=1}^n (x_{t-1}^n x_{t-1}^{n'} + P_{t-1}^n). \quad (6.70)$$

In (6.67)–(6.70), the smoothers are calculated under the current value of the parameters $\Theta^{(j-1)}$; for simplicity, we have not explicitly displayed this fact.

Minimizing (6.67) with respect to the parameters, at iteration j , constitutes the maximization step, and is analogous to the usual multivariate regression approach, which yields the updated estimates

$$\Phi^{(j)} = S_{10} S_{00}^{-1}, \quad (6.71)$$

$$Q^{(j)} = n^{-1} (S_{11} - S_{10} S_{00}^{-1} S_{10}'), \quad (6.72)$$

and

$$R^{(j)} = n^{-1} \sum_{t=1}^n [(y_t - A_t x_t^n)(y_t - A_t x_t^n)' + A_t P_t^n A_t']. \quad (6.73)$$

The updates for the initial mean and variance–covariance matrix are

$$\mu_0^{(j)} = x_0^n \quad \text{and} \quad \Sigma_0^{(j)} = P_0^n \quad (6.74)$$

obtained from minimizing (6.67).

The overall procedure can be regarded as simply alternating between the Kalman filtering and smoothing recursions and the multivariate normal maximum likelihood estimators, as given by (6.71)–(6.74). Convergence results for the EM algorithm under general conditions can be found in Wu (1983). We summarize the iterative procedure as follows.

- (i) Initialize the procedure by selecting starting values for the parameters $\Theta^{(0)} = \{\mu_0, \Sigma_0, \Phi, Q, R\}$.

On iteration j , ($j = 1, 2, \dots$):

- (ii) Compute the incomplete-data likelihood, $-\ln L_Y(\Theta^{(j-1)})$; see equation (6.63).
- (iii) Perform the E-Step. Use **Property 6.1**, **Property 6.2**, and **Property 6.3** to obtain the smoothed values x_t^n, P_t^n and $P_{t,t-1}^n$, for $t = 1, \dots, n$, using the parameters $\Theta^{(j-1)}$. Use the smoothed values to calculate S_{11}, S_{10}, S_{00} given in (6.68)–(6.70).
- (iv) Perform the M-Step. Update the estimates, μ_0, Σ_0, Φ, Q , and R using (6.71)–(6.74), to obtain $\Theta^{(j)}$.
- (v) Repeat Steps (ii) – (iv) to convergence.

Example 6.8 EM Algorithm for Example 6.3

Using the same data generated in Example 6.6, we performed an EM algorithm estimation of the parameters ϕ , σ_w^2 and σ_v^2 as well as the initial parameters μ_0 and Σ_0 using the script `EM0`. The convergence rate of the EM algorithm compared with the Newton–Raphson procedure is slow. In this example, with convergence being claimed when the relative change in the log likelihood is less than .00001; convergence was attained after 41 iterations. The final estimates, along with their standard errors (in parentheses), were

$$\hat{\phi} = .81 (.08), \quad \hat{\sigma}_w = .87 (.17), \quad \hat{\sigma}_v = .84 (.14),$$

with $\hat{\mu}_0 = -1.98$ and $\hat{\Sigma}_0 = .03$. Evaluation of the standard errors used a call to `fdHess` in the `nlme` R package to evaluate the Hessian at the final estimates. The `nlme` package must be loaded prior to the call to `fdHess`.

```
library(nlme) # loads package nlme
# Generate data (same as Example 6.6)
set.seed(999); num = 100; N = num+1
x = arima.sim(n=N, list(ar = .8, sd=1))
y = ts(x[-1] + rnorm(num,0,1))
# Initial Estimates
u = ts.intersect(y, lag(y,-1), lag(y,-2))
varu = var(u); coru = cor(u)
phi = coru[1,3]/coru[1,2]
q = (1-phi^2)*varu[1,2]/phi
r = varu[1,1] - q/(1-phi^2)
cr = sqrt(r); cq = sqrt(q); mu0 = 0; Sigma0 = 2.8
# EM procedure - output not shown
(em = EM0(num, y, 1, mu0, Sigma0, phi, cq, cr, 75, .00001))
# Standard Errors (this uses nlme)
phi = em$Phi; cq = chol(em$Q); cr = chol(em$R)
mu0 = em$mu0; Sigma0 = em$Sigma0
para = c(phi, cq, cr)
Linn = function(para){ # to evaluate likelihood at estimates
  kf = Kfilter0(num, y, 1, mu0, Sigma0, para[1], para[2], para[3])
  return(kf$like) }
emhess = fdHess(para, function(para) Linn(para))
SE = sqrt(diag(solve(emhess$Hessian)))
# Display Summary of Estimation
estimate = c(para, em$mu0, em$Sigma0); SE = c(SE, NA, NA)
u = cbind(estimate, SE)
rownames(u) = c("phi", "sigw", "sigv", "mu0", "Sigma0"); u
```

	estimate	SE
phi	0.80639903	0.07986272
sigw	0.86442634	0.16719703
sigv	0.84276381	0.13805072
mu0	-1.96010956	NA
Sigma0	0.03638596	NA

STEADY STATE AND ASYMPTOTIC DISTRIBUTION OF THE MLES

The asymptotic distribution of estimators of the model parameters, say, $\hat{\Theta}_n$, is studied in very general terms in Douc, Moulines, and Stoffer (2014, Chapter 13). Earlier treatments can be found in Caines (1988, Chapters 7 and 8), and in Hannan and Deistler

(1988, Chapter 4). In these references, the consistency and asymptotic normality of the estimators are established under general conditions. An essential condition is the stability of the filter. Stability of the filter assures that, for large t , the innovations ϵ_t are basically copies of each other (that is, independent and identically distributed) with a stable covariance matrix Σ that does not depend on t and that, asymptotically, the innovations contain all of the information about the unknown parameters. Although it is not necessary, for simplicity, we shall assume here that $A_t \equiv A$ for all t . Details on departures from this assumption can be found in Jazwinski (1970, Sections 7.6 and 7.8). We also drop the inputs and use the model in the form of (6.1) and (6.2).

For stability of the filter, we assume the eigenvalues of Φ are less than one in absolute value; this assumption can be weakened (for example, see Harvey, 1991, Section 4.3), but we retain it for simplicity. This assumption is enough to ensure the stability of the filter in that, as $t \rightarrow \infty$, the filter error covariance matrix P_t^f converges to P , the steady-state error covariance matrix, the gain matrix K_t converges to K , the steady-state gain matrix, from which it follows that the innovation variance–covariance matrix Σ_t converges to Σ , the steady-state variance–covariance matrix of the stable innovations; details can be found in Jazwinski (1970, Sections 7.6 and 7.8) and Anderson and Moore (1979, Section 4.4). In particular, the steady-state filter error covariance matrix, P , satisfies the Riccati equation:

$$P = \Phi[P - PA'(APA' + R)^{-1}AP]\Phi' + Q;$$

the steady-state gain matrix satisfies $K = PA'[APA' + R]^{-1}$. In **Example 6.5** (see **Table 6.1**), for all practical purposes, stability was reached by the fourth observation.

When the process is in steady-state, we may consider x_{t+1}^t as the steady-state predictor and interpret it as $x_{t+1}^t = E(x_{t+1} | y_t, y_{t-1}, \dots)$. As can be seen from (6.19) and (6.21), the steady-state predictor can be written as

$$\begin{aligned} x_{t+1}^t &= \Phi[I - KA]x_t^{t-1} + \Phi Ky_t \\ &= \Phi x_t^{t-1} + \Phi K \epsilon_t, \end{aligned} \tag{6.75}$$

where ϵ_t is the steady-state innovation process given by

$$\epsilon_t = y_t - E(y_t | y_{t-1}, y_{t-2}, \dots).$$

In the Gaussian case, $\epsilon_t \sim \text{iid } N(0, \Sigma)$, where $\Sigma = APA' + R$. In steady-state, the observations can be written as

$$y_t = Ax_t^{t-1} + \epsilon_t. \tag{6.76}$$

Together, (6.75) and (6.76) make up the steady-state innovations form of the dynamic linear model.

In the following property, we assume the Gaussian state-space model (6.1) and (6.2), is time invariant, i.e., $A_t \equiv A$, the eigenvalues of Φ are within the unit circle and the model has the smallest possible dimension (see Hannan and Deistler, 1988,

Section 2.3 for details). We denote the true parameters by θ_0 , and we assume the dimension of θ_0 is the dimension of the parameter space. Although it is not necessary to assume w_t and v_t are Gaussian, certain additional conditions would have to apply and adjustments to the asymptotic covariance matrix would have to be made; see Douc et al. (2014, Chapter 13).

Property 6.4 Asymptotic Distribution of the Estimators

Under general conditions, let $\hat{\theta}_n$ be the estimator of θ_0 obtained by maximizing the innovations likelihood, $L_Y(\theta)$, as given in (6.63). Then, as $n \rightarrow \infty$,

$$\sqrt{n} (\hat{\theta}_n - \theta_0) \xrightarrow{d} N [0, \mathcal{I}(\theta_0)^{-1}],$$

where $\mathcal{I}(\theta)$ is the asymptotic information matrix given by

$$\mathcal{I}(\theta) = \lim_{n \rightarrow \infty} n^{-1} E \left[-\partial^2 \ln L_Y(\theta) / \partial \theta \partial \theta' \right].$$

For a Newton procedure, the Hessian matrix (as described in Example 6.6) at the time of convergence can be used as an estimate of $n\mathcal{I}(\theta_0)$ to obtain estimates of the standard errors. In the case of the EM algorithm, no derivatives are calculated, but we may include a numerical evaluation of the Hessian matrix at the time of convergence to obtain estimated standard errors. Also, extensions of the EM algorithm exist, such as the SEM algorithm (Meng and Rubin, 1991), that include a procedure for the estimation of standard errors. In the examples of this section, the estimated standard errors were obtained from the numerical Hessian matrix of $-\ln L_Y(\hat{\theta})$, where $\hat{\theta}$ is the vector of parameters estimates at the time of convergence.

6.4 Missing Data Modifications

An attractive feature available within the state-space framework is its ability to treat time series that have been observed irregularly over time. For example, Jones (1980) used the state-space representation to fit ARMA models to series with missing observations, and Palma and Chan (1997) used the model for estimation and forecasting of ARFIMA series with missing observations. Shumway and Stoffer (1982) described the modifications necessary to fit multivariate state-space models via the EM algorithm when data are missing. We will discuss the procedure in detail in this section. Throughout this section, for notational simplicity, we assume the model is of the form (6.1) and (6.2).

Suppose, at a given time t , we define the partition of the $q \times 1$ observation vector $y_t = (y_t^{(1)'}, y_t^{(2)'})'$, where the first $q_{1t} \times 1$ component is observed and the second $q_{2t} \times 1$ component is unobserved, $q_{1t} + q_{2t} = q$. Then, write the partitioned observation equation

$$\begin{pmatrix} y_t^{(1)} \\ y_t^{(2)} \end{pmatrix} = \begin{bmatrix} A_t^{(1)} \\ A_t^{(2)} \end{bmatrix} x_t + \begin{pmatrix} v_t^{(1)} \\ v_t^{(2)} \end{pmatrix}, \quad (6.77)$$

where $A_t^{(1)}$ and $A_t^{(2)}$ are, respectively, the $q_{1t} \times p$ and $q_{2t} \times p$ partitioned observation matrices, and

$$\text{cov} \begin{pmatrix} v_t^{(1)} \\ v_t^{(2)} \end{pmatrix} = \begin{bmatrix} R_{11t} & R_{12t} \\ R_{21t} & R_{22t} \end{bmatrix} \quad (6.78)$$

denotes the covariance matrix of the measurement errors between the observed and unobserved parts.

In the missing data case where $y_t^{(2)}$ is not observed, we may modify the observation equation in the DLM, (6.1)–(6.2), so that the model is

$$x_t = \Phi x_{t-1} + w_t \quad \text{and} \quad y_t^{(1)} = A_t^{(1)} x_t + v_t^{(1)}, \quad (6.79)$$

where now, the observation equation is q_{1t} -dimensional at time t . In this case, it follows directly from [Corollary 6.1](#) that the filter equations hold with the appropriate notational substitutions. If there are no observations at time t , then set the gain matrix, K_t , to the $p \times q$ zero matrix in [Property 6.1](#), in which case $x_t^f = x_t^{t-1}$ and $P_t^f = P_t^{t-1}$.

Rather than deal with varying observational dimensions, it is computationally easier to modify the model by zeroing out certain components and retaining a q -dimensional observation equation throughout. In particular, [Corollary 6.1](#) holds for the missing data case if, at update t , we substitute

$$y_{(t)} = \begin{pmatrix} y_t^{(1)} \\ 0 \end{pmatrix}, \quad A_{(t)} = \begin{bmatrix} A_t^{(1)} \\ 0 \end{bmatrix}, \quad R_{(t)} = \begin{bmatrix} R_{11t} & 0 \\ 0 & I_{22t} \end{bmatrix}, \quad (6.80)$$

for y_t , A_t , and R , respectively, in (6.21)–(6.23), where I_{22t} is the $q_{2t} \times q_{2t}$ identity matrix. With the substitutions (6.80), the innovation values (6.24) and (6.25) will now be of the form

$$\epsilon_{(t)} = \begin{pmatrix} \epsilon_t^{(1)} \\ 0 \end{pmatrix}, \quad \Sigma_{(t)} = \begin{bmatrix} A_t^{(1)} P_t^{t-1} A_t^{(1)'} + R_{11t} & 0 \\ 0 & I_{22t} \end{bmatrix}, \quad (6.81)$$

so that the innovations form of the likelihood given in (6.63) is correct for this case. Hence, with the substitutions in (6.80), maximum likelihood estimation via the innovations likelihood can proceed as in the complete data case.

Once the missing data filtered values have been obtained, Stoffer (1982) also established the smoother values can be processed using [Property 6.2](#) and [Property 6.3](#) with the values obtained from the missing data-filtered values. In the missing data case, the state estimators are denoted

$$x_t^{(s)} = E \left(x_t \mid y_1^{(1)}, \dots, y_s^{(1)} \right), \quad (6.82)$$

with error variance–covariance matrix

$$P_t^{(s)} = E \left\{ \left(x_t - x_t^{(s)} \right) \left(x_t - x_t^{(s)} \right)' \right\}. \quad (6.83)$$

The missing data lag-one smoother covariances will be denoted by $P_{t,t-1}^{(n)}$.

The maximum likelihood estimators in the EM procedure require further modifications for the case of missing data. Now, we consider

$$\mathcal{Y}_n^{(1)} = \{y_1^{(1)}, \dots, y_n^{(1)}\} \quad (6.84)$$

as the incomplete data, and $\mathcal{X}_n, \mathcal{Y}_n$, as defined in (6.64), as the complete data. In this case, the complete data likelihood, (6.64), or equivalently (6.65), is the same, but to implement the E-step, at iteration j , we must calculate

$$\begin{aligned} Q(\theta \mid \theta^{(j-1)}) &= E\{-2 \ln L_{X,Y}(\theta) \mid \mathcal{Y}_n^{(1)}, \theta^{(j-1)}\} \\ &= E_* \left\{ \ln |\Sigma_0| + \text{tr } \Sigma_0^{-1} (x_0 - \mu_0)(x_0 - \mu_0)' \mid \mathcal{Y}_n^{(1)} \right\} \\ &\quad + E_* \left\{ n \ln |Q| + \sum_{t=1}^n \text{tr } [Q^{-1} (x_t - \Phi x_{t-1})(x_t - \Phi x_{t-1})'] \mid \mathcal{Y}_n^{(1)} \right\} \\ &\quad + E_* \left\{ n \ln |R| + \sum_{t=1}^n \text{tr } [R^{-1} (y_t - A_t x_t)(y_t - A_t x_t)'] \mid \mathcal{Y}_n^{(1)} \right\}, \end{aligned} \quad (6.85)$$

where E_* denotes the conditional expectation under $\theta^{(j-1)}$ and tr denotes trace. The first two terms in (6.85) will be like the first two terms of (6.67) with the smoothers x_t^n , P_t^n , and $P_{t,t-1}^n$ replaced by their missing data counterparts, $x_t^{(n)}$, $P_t^{(n)}$, and $P_{t,t-1}^{(n)}$. What changes in the missing data case is the third term of (6.85), where we must evaluate $E_*(y_t^{(2)} \mid \mathcal{Y}_n^{(1)})$ and $E_*(y_t^{(2)} y_t^{(2)'} \mid \mathcal{Y}_n^{(1)})$. In Stoffer (1982), it is shown that

$$\begin{aligned} E_* \{ (y_t - A_t x_t)(y_t - A_t x_t)' \mid \mathcal{Y}_n^{(1)} \} &= \begin{pmatrix} y_t^{(1)} - A_t^{(1)} x_t^{(n)} \\ R_{*21t} R_{*11t}^{-1} (y_t^{(1)} - A_t^{(1)} x_t^{(n)}) \end{pmatrix} \begin{pmatrix} y_t^{(1)} - A_t^{(1)} x_t^{(n)} \\ R_{*21t} R_{*11t}^{-1} (y_t^{(1)} - A_t^{(1)} x_t^{(n)}) \end{pmatrix}' \\ &\quad + \begin{pmatrix} A_t^{(1)} \\ R_{*21t} R_{*11t}^{-1} A_t^{(1)} \end{pmatrix} P_t^{(n)} \begin{pmatrix} A_t^{(1)} \\ R_{*21t} R_{*11t}^{-1} A_t^{(1)} \end{pmatrix}' \\ &\quad + \begin{pmatrix} 0 \\ 0 \ R_{*22t} - R_{*21t} R_{*11t}^{-1} R_{*12t} \end{pmatrix}. \end{aligned} \quad (6.86)$$

In (6.86), the values of R_{*ikt} , for $i, k = 1, 2$, are the current values specified by $\theta^{(j-1)}$. In addition, $x_t^{(n)}$ and $P_t^{(n)}$ are the values obtained by running the smoother under the current parameter estimates specified by $\theta^{(j-1)}$.

In the case in which observed and unobserved components have uncorrelated errors, that is, R_{*12t} is the zero matrix, (6.86) can be simplified to

$$\begin{aligned} E_* \{ (y_t - A_t x_t)(y_t - A_t x_t)' \mid \mathcal{Y}_n^{(1)} \} &= (y_{(t)} - A_{(t)} x_t^{(n)})(y_{(t)} - A_{(t)} x_t^{(n)})' + A_{(t)} P_t^{(n)} A_{(t)}' + \begin{pmatrix} 0 & 0 \\ 0 & R_{*22t} \end{pmatrix}, \end{aligned} \quad (6.87)$$

where $y_{(t)}$ and $A_{(t)}$ are defined in (6.80).

In this simplified case, the missing data M-step looks like the M-step given in (6.68)-(6.74). That is, with

$$S_{(11)} = \sum_{t=1}^n (x_t^{(n)} x_t^{(n)'} + P_t^{(n)}), \quad (6.88)$$

$$S_{(10)} = \sum_{t=1}^n (x_t^{(n)} x_{t-1}^{(n)'} + P_{t,t-1}^{(n)}), \quad (6.89)$$

and

$$S_{(00)} = \sum_{t=1}^n (x_{t-1}^{(n)} x_{t-1}^{(n)'} + P_{t-1}^{(n)}), \quad (6.90)$$

where the smoothers are calculated under the present value of the parameters $\Theta^{(j-1)}$ using the missing data modifications, at iteration j , the *maximization step* is

$$\Phi^{(j)} = S_{(10)} S_{(00)}^{-1}, \quad (6.91)$$

$$Q^{(j)} = n^{-1} (S_{(11)} - S_{(10)} S_{(00)}^{-1} S_{(10)}'), \quad (6.92)$$

and

$$\begin{aligned} R^{(j)} = n^{-1} \sum_{t=1}^n D_t \Big\{ & (y_{(t)} - A_{(t)} x_t^{(n)}) (y_{(t)} - A_{(t)} x_t^{(n)})' \\ & + A_{(t)} P_t^{(n)} A_{(t)}' + \begin{pmatrix} 0 & 0 \\ 0 & R_{22t}^{(j-1)} \end{pmatrix} \Big\} D_t', \end{aligned} \quad (6.93)$$

where D_t is a permutation matrix that reorders the variables at time t in their original order and $y_{(t)}$ and $A_{(t)}$ are defined in (6.80). For example, suppose $q = 3$ and at time t , y_{t2} is missing. Then,

$$y_{(t)} = \begin{pmatrix} y_{t1} \\ y_{t3} \\ 0 \end{pmatrix}, \quad A_{(t)} = \begin{bmatrix} A_{t1} \\ A_{t3} \\ 0' \end{bmatrix}, \quad \text{and} \quad D_t = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix},$$

where A_{ti} is the i th row of A_t and $0'$ is a $1 \times p$ vector of zeros. In (6.93), only R_{11t} gets updated, and R_{22t} at iteration j is simply set to its value from the previous iteration, $j - 1$. Of course, if we cannot assume $R_{12t} = 0$, (6.93) must be changed accordingly using (6.86), but (6.91) and (6.92) remain the same. As before, the parameter estimates for the initial state are updated as

$$\mu_0^{(j)} = x_0^{(n)} \quad \text{and} \quad \Sigma_0^{(j)} = P_0^{(n)}. \quad (6.94)$$

Example 6.9 Longitudinal Biomedical Data

We consider the biomedical data in Example 6.1, which have portions of the three-dimensional vector missing after the 40th day. The maximum likelihood procedure yielded the estimators (code at the end of the example):

```
$Phi
      [,1]  [,2]  [,3]
[1,]  0.984 -0.041 0.0088
[2,]  0.061  0.921 0.0065
[3,] -1.495  2.289 0.7936
```




Fig. 6.5. Smoothed values for various components in the blood parameter tracking problem. The actual data are shown as points, the smoothed values are shown as solid lines, and ± 3 standard error bounds are shown as dashed lines.

```
$Q
      [,1] [,2] [,3]
[1,]  0.0136 -0.0018 0.012
[2,] -0.0018  0.0029 0.018
[3,]  0.0124  0.0179 3.494

$R
      [,1] [,2] [,3]
[1,]  0.007 0.000  0.0
[2,]  0.000 0.017  0.0
[3,]  0.000 0.000  1.1
```

for the transition, state error covariance and observation error covariance matrices, respectively. The coupling between the first and second series is relatively weak, whereas the third series HCT is strongly related to the first two; that is,

$$\hat{x}_{t3} = -1.495x_{t-1,1} + 2.289x_{t-1,2} + .7936x_{t-1,3}.$$

Hence, the HCT is negatively correlated with white blood count (WBC) and positively correlated with platelet count (PLT). Byproducts of the procedure are estimated trajectories for all three longitudinal series and their respective prediction

intervals. In particular, [Figure 6.5](#) shows the data as points, the estimated smoothed values $\hat{x}_t^{(n)}$ as solid lines, and error bounds, $\hat{x}_t^{(n)} \pm 2\sqrt{\hat{P}_t^{(n)}}$ as dotted lines, for critical post-transplant platelet count.

In the following R code we use the script [EM1](#). In this case the observation matrices A_t are either the identity or the zero matrix because either all the series are observed or none of them are observed.

```
y = cbind(WBC, PLT, HCT)
num = nrow(y)
A = array(0, dim=c(3,3,num))
# make array of obs matrices
for(k in 1:num) if (y[k,1] > 0) A[, ,k]= diag(1,3)
# Initial values
mu0 = matrix(0, 3, 1)
Sigma0 = diag(c(.1, .1, 1), 3)
Phi = diag(1, 3)
cQ = diag(c(.1, .1, 1), 3)
cR = diag(c(.1, .1, 1), 3)
# EM procedure - some output previously shown
(em = EM1(num, y, A, mu0, Sigma0, Phi, cQ, cR, 100, .001))
# Graph smoother
ks = Ksmooth1(num, y, A, em$mu0, em$Sigma0, em$Phi, 0, 0, chol(em$Q),
              chol(em$R), 0)
y1s = ks$xs[1,]; y2s = ks$xs[2,]; y3s = ks$xs[3,]
p1 = 2*sqrt(ks$Ps[1,1,])
p2 = 2*sqrt(ks$Ps[2,2,])
p3 = 2*sqrt(ks$Ps[3,3,])
par(mfrow=c(3,1), mar=c(4,4,1,1)+.2)
plot(WBC, type="p", pch=19, ylim=c(1,5), xlab="day")
lines(y1s); lines(y1s+p1, lty=2, col=4); lines(y1s-p1, lty=2, col=4)
plot(PLT, type="p", ylim=c(3,6), pch=19, xlab="day")
lines(y2s); lines(y2s+p2, lty=2, col=4); lines(y2s-p2, lty=2, col=4)
plot(HCT, type="p", pch=19, ylim=c(20,40), xlab="day")
lines(y3s); lines(y3s+p3, lty=2, col=4); lines(y3s-p3, lty=2, col=4)
```

6.5 Smoothing Splines and the Kalman Smoother

There is a direct connection between smoothing splines; e.g., Eubank (1993), Green (1993), or Wahba (1990) and state space models. The basic idea of smoothing splines (recall [Example 2.14](#)) in discrete time is we suppose that data y_t are generated by $y_t = \mu_t + \epsilon_t$ for $t = 1, \dots, n$, where μ_t is a smooth function of t , and ϵ_t is white noise. In cubic smoothing with knots at the time points t , μ_t is estimated by minimizing

$$\sum_{t=1}^n [y_t - \mu_t]^2 + \lambda \sum_{t=1}^n (\nabla^2 \mu_t)^2 \quad (6.95)$$

with respect to μ_t , where $\lambda > 0$ is a smoothing parameter. The parameter λ controls the degree of smoothness, with larger values yielding smoother estimates. For example, if $\lambda = 0$, then the minimizer is the data itself $\hat{\mu}_t = y_t$; consequently, the estimate will not be smooth. If $\lambda = \infty$, then the only way to minimize (6.95) is to choose the

second term to be zero, i.e., $\nabla^2 \mu_t = 0$, in which case it is of the form $\mu_t = \alpha + \beta t$, and we are in the setting of linear regression.² Hence, the choice of $\lambda > 0$ is seen as a trade-off between fitting a line that goes through all the data points and linear regression.

Now, consider the model given by

$$\nabla^2 \mu_t = w_t \quad \text{and} \quad y_t = \mu_t + v_t, \quad (6.96)$$

where w_t and v_t are independent white noise processes with $\text{var}(w_t) = \sigma_w^2$ and $\text{var}(v_t) = \sigma_v^2$. Rewrite (6.96) as

$$\begin{pmatrix} \mu_t \\ \mu_{t-1} \end{pmatrix} = \begin{bmatrix} 2 & -1 \\ 1 & 0 \end{bmatrix} \begin{pmatrix} \mu_{t-1} \\ \mu_{t-2} \end{pmatrix} + \begin{bmatrix} 1 \\ 0 \end{bmatrix} w_t \quad \text{and} \quad y_t = \begin{bmatrix} 1 & 0 \end{bmatrix} \begin{pmatrix} \mu_t \\ \mu_{t-1} \end{pmatrix} + v_t, \quad (6.97)$$

so that the state vector is $x_t = (\mu_t, \mu_{t-1})'$. It is clear then that (6.96) specifies a state space model.

Note that the model is similar to the local level model discussed in [Example 6.5](#). In particular, the state process could be written as $\mu_t = \mu_{t-1} + \eta_t$, where $\eta_t = \eta_{t-1} + w_t$. An example of such a trajectory can be seen in [Figure 6.6](#); note that the generated data in [Figure 6.6](#) look like the global temperature data in [Figure 1.2](#).

Next, we examine the problem of estimating the states, x_t , when the model parameters, $\theta = \{\sigma_w^2, \sigma_v^2\}$, are specified. For ease, we assume x_0 is fixed. Then using the notation surrounding equations (6.64)–(6.65), the goal is to find the MLE of $\mathcal{X}_n = \{x_1, \dots, x_n\}$ given $\mathcal{Y}_n = \{y_1, \dots, y_n\}$; i.e., to maximize $\log f_\theta(\mathcal{X}_n | \mathcal{Y}_n)$ with respect to the states. Because of the Gaussianity, the maximum (or mode) of the distribution is when the states are estimated by x_t^n , the conditional means. These values are, of course, the smoothers obtained via [Property 6.2](#).

But $\log f_\theta(\mathcal{X}_n | \mathcal{Y}_n) = \log f_\theta(\mathcal{X}_n, \mathcal{Y}_n) - \log f_\theta(\mathcal{Y}_n)$, so maximizing the complete data likelihood, $\log f_\theta(\mathcal{X}_n, \mathcal{Y}_n)$ with respect to $\mathcal{X}_n$, is an equivalent problem. Writing (6.65) in the notation of (6.96), we have,

$$-2 \log f_\theta(\mathcal{X}_n, \mathcal{Y}_n) \propto \sigma_w^{-2} \sum_{t=1}^n (\nabla^2 \mu_t)^2 + \sigma_v^{-2} \sum_{t=1}^n (y_t - \mu_t)^2, \quad (6.98)$$

where we have kept only the terms involving the states, μ_t . If we set $\lambda = \sigma_v^2 / \sigma_w^2$, we can write

$$-2 \log f_\theta(\mathcal{X}_n, \mathcal{Y}_n) \propto \lambda \sum_{t=1}^n (\nabla^2 \mu_t)^2 + \sum_{t=1}^n (y_t - \mu_t)^2, \quad (6.99)$$

so that maximizing $\log f_\theta(\mathcal{X}_n, \mathcal{Y}_n)$ with respect to the states is equivalent to minimizing (6.99), which is the original problem stated in (6.95).

In the general state space setting, we would estimate σ_w^2 and σ_v^2 via maximum likelihood as described in [§6.3](#), and then obtain the smoothed state values by running [Property 6.2](#) with the estimated variances, say $\hat{\sigma}_w^2$ and $\hat{\sigma}_v^2$. In this case, the estimated value of the smoothing parameter would be given by $\hat{\lambda} = \hat{\sigma}_v^2 / \hat{\sigma}_w^2$.

² That the unique general solution to $\nabla^2 \mu_t = 0$ is of the form $\mu_t = \alpha + \beta t$ follows from difference equation theory; e.g., see Mickens (1990).




Fig. 6.6. Display for **Example 6.10**: Simulated state process, μ_t and observations y_t from the model (6.96) with $n = 50$, $\sigma_w = .1$ and $\sigma_v = 1$. Estimated smoother (dashed lines): $\hat{\mu}_{t|n}$ and corresponding 95% confidence band. Gcv smoothing spline (thin solid line).

Example 6.10 Smoothing Splines

In this example, we generated the signal, or state process, μ_t and observations y_t from the model (6.96) with $n = 50$, $\sigma_w = .1$ and $\sigma_v = 1$. The state is displayed in **Figure 6.6** as a thick solid line, and the observations are displayed as points. We then estimated σ_w and σ_v using Newton-Raphson techniques and obtained $\hat{\sigma}_w = .08$ and $\hat{\sigma}_v = .94$. We then used **Property 6.2** to generate the estimated smoothers, say, $\hat{\mu}_{t|n}$, and those values are displayed in **Figure 6.6** as a thick dashed line along with a corresponding 95% (pointwise) confidence band as thin dashed lines. Finally, we used the R function `smooth.spline` to fit a smoothing spline to the data based on the method of generalized cross-validation (gcv). The fitted spline is displayed in **Figure 6.6** as a thin solid line, which is nearly identical to $\hat{\mu}_{t|n}$.

The R code to reproduce **Figure 6.6** is given below.

```
set.seed(123)
num = 50
w = rnorm(num, 0, .1)
x = cumsum(cumsum(w))
y = x + rnorm(num, 0, 1)
plot.ts(x, ylab="", lwd=2, ylim=c(-1,8))
lines(y, type='o', col='#808080')
## State Space ##
Phi = matrix(c(2,1,-1,0), 2)
A = matrix(c(1,0), 1)
mu0 = matrix(0, 2)
Sigma0 = diag(1, 2)
Linn=function(para){
  sigw = para[1]
  sigv = para[2]
  cQ = diag(c(sigw,0))
  kf = Kfilter0(num,y,A,mu0,Sigma0,Phi,cQ,sigv)
  return(kf$like)
}
## Estimation ##
init.par=c(.1,1)
(est = optim(init.par, Linn, NULL, method="BFGS", hessian=TRUE,
  control=list(trace=1,REPORT=1)))
SE = sqrt(diag(solve(est$hessian)))
# output
```

```

estimate = est$par
u = cbind(estimate, SE)
rownames(u)= c("sigw","sigv"); u
# Smooth
sigw = est$par[1]
cQ = diag(c(sigw,0))
sigv = est$par[2]
ks = Ksmooth0(num,y,A,mu0,Sigma0,Phi,cQ,sigv)
xsmoo = ts(ks$xs[1,1,])
psmoo = ts(ks$Ps[1,1,])
upp = xsmoo+2*sqrt(psmoo)
low =xsmoo-2*sqrt(psmoo)
lines(xsmoo, col=4, lty=2, lwd=3)
lines(upp, col=4,lty=2)
lines(low, col=4, lty=2)
lines(smooth.spline(y),lty=1,col=2) # R smoothing spline function
legend("topleft", c("Observations","State"), pch=c(1,-1), lty=c(1,1),
      lwd=c(1,2), col=c('#808080',1), )
legend("bottomright", c("Smoother", "GCV Spline"), lty=c(2,1), lwd=c(3,1),
      col=c(4,2) )

```

6.6 Structural Models: Signal Extraction and Forecasting

In order to develop computing techniques for handling a versatile cross section of possible models, it is necessary to restrict the state-space model somewhat, and we consider one possible class of specializations in this section. The components of the model are taken as linear processes that can be adapted to represent fixed and disturbed trends and periodicities as well as classical autoregressions. The observed series is regarded as being a sum of component signal series. To illustrate the possibilities, we consider an example that shows how to fit a sum of trend, seasonal, and irregular components to the quarterly earnings data that we have considered before.

Example 6.11 Johnson & Johnson Quarterly Earnings

Consider the quarterly earnings series from the U.S. company Johnson & Johnson as given in [Figure 1.1](#). The series is highly nonstationary, and there is both a trend signal that is gradually increasing over time and a seasonal component that cycles every four quarters or once per year. The seasonal component is getting larger over time as well. Transforming into logarithms or even taking the n th root does not seem to make the series stationary, as there is a slight bend to the transformed curve. Suppose, however, we consider the series to be the sum of a trend component, a seasonal component, and a white noise. That is, let the observed series be expressed as

$$y_t = T_t + S_t + v_t, \quad (6.100)$$

where T_t is trend and S_t is the seasonal component. Suppose we allow the trend to increase exponentially; that is,

$$T_t = \phi T_{t-1} + w_{t1}, \quad (6.101)$$

where the coefficient $\phi > 1$ characterizes the increase. Let the seasonal component be modeled as

$$S_t + S_{t-1} + S_{t-2} + S_{t-3} = w_{t2}, \quad (6.102)$$

which corresponds to assuming the seasonal component is expected to sum to zero over a complete period or four quarters. To express this model in state-space form, let $x_t = (T_t, S_t, S_{t-1}, S_{t-2})'$ be the state vector so the observation equation (6.2) can be written as

$$y_t = \begin{pmatrix} 1 & 1 & 0 & 0 \end{pmatrix} \begin{pmatrix} T_t \\ S_t \\ S_{t-1} \\ S_{t-2} \end{pmatrix} + v_t,$$

with the state equation written as

$$\begin{pmatrix} T_t \\ S_t \\ S_{t-1} \\ S_{t-2} \end{pmatrix} = \begin{pmatrix} \phi & 0 & 0 & 0 \\ 0 & -1 & -1 & -1 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} T_{t-1} \\ S_{t-1} \\ S_{t-2} \\ S_{t-3} \end{pmatrix} + \begin{pmatrix} w_{t1} \\ w_{t2} \\ 0 \\ 0 \end{pmatrix},$$

where $R = r_{11}$ and

$$Q = \begin{pmatrix} q_{11} & 0 & 0 & 0 \\ 0 & q_{22} & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}.$$

The model reduces to state-space form, (6.1) and (6.2), with $p = 4$ and $q = 1$. The parameters to be estimated are r_{11} , the noise variance in the measurement equations, q_{11} and q_{22} , the model variances corresponding to the trend and seasonal components and ϕ , the transition parameter that models the growth rate. Growth is about 3% per year, and we began with $\phi = 1.03$. The initial mean was fixed at $\mu_0 = (.7, 0, 0, 0)'$, with uncertainty modeled by the diagonal covariance matrix with $\Sigma_{0ii} = .04$, for $i = 1, \dots, 4$. Initial state covariance values were taken as $q_{11} = .01$, $q_{22} = .01$. The measurement error covariance was started at $r_{11} = .25$.

After about 20 iterations of a Newton–Raphson, the transition parameter estimate was $\hat{\phi} = 1.035$, corresponding to exponential growth with inflation at about 3.5% per year. The measurement uncertainty was small at $\sqrt{\hat{r}_{11}} = .0005$, compared with the model uncertainties $\sqrt{\hat{q}_{11}} = .1397$ and $\sqrt{\hat{q}_{22}} = .2209$. Figure 6.7 shows the smoothed trend estimate and the exponentially increasing seasonal components. We may also consider forecasting the Johnson & Johnson series, and the result of a 12-quarter forecast is shown in Figure 6.8 as basically an extension of the latter part of the observed data.

This example uses the `Kfilter0` and `Ksmooth0` scripts as follows.

```
num = length(jj)
A = cbind(1,1,0,0)
# Function to Calculate Likelihood
Linn=function(para){
  Phi = diag(0,4); Phi[1,1] = para[1]
```




Fig. 6.7. Estimated trend component, T_t^n , and seasonal component, S_t^n , of the Johnson and Johnson quarterly earnings series. Gray areas are three root MSE bounds.

```

Phi[2,]=c(0,-1,-1,-1); Phi[3,]=c(0,1,0,0); Phi[4,]=c(0,0,1,0)
cQ1 = para[2]; cQ2 = para[3]      # sqrt q11 and sqrt q22
cQ=diag(0,4); cQ[1,1]=cQ1; cQ[2,2]=cQ2
cR = para[4]                      # sqrt r11
kf = Kfilter0(num,jj,A,mu0,Sigma0,Phi,cQ,cR)
return(kf$like)
}
# Initial Parameters
mu0 = c(.7,0,0,0); Sigma0 = diag(.04,4)
init.par = c(1.03,.1,.1,.5)      # Phi[1,1], the 2 Qs and R
# Estimation
est = optim(init.par, Linn, NULL, method="BFGS", hessian=TRUE,
            control=list(trace=1,REPORT=1))
SE = sqrt(diag(solve(est$hessian)))
u = cbind(estimate=est$par,SE)
rownames(u)=c("Phi11","sigw1","sigw2","sigv"); u
# Smooth
Phi = diag(0,4); Phi[1,1] = est$par[1]
Phi[2,]=c(0,-1,-1,-1); Phi[3,]=c(0,1,0,0); Phi[4,]=c(0,0,1,0)
cQ1 = est$par[2]; cQ2 = est$par[3]
cQ = diag(0,4); cQ[1,1]=cQ1; cQ[2,2]=cQ2
cR = est$par[4]
ks = Ksmooth0(num,jj,A,mu0,Sigma0,Phi,cQ,cR)
# Plot
Tsm = ts(ks$xs[1,,], start=1960, freq=4)
Ssm = ts(ks$xs[2,,], start=1960, freq=4)
p1 = 3*sqrt(ks$Ps[1,1,]); p2 = 3*sqrt(ks$Ps[2,2,])
par(mfrow=c(2,1), mar=c(3,3,2,1), mgp=c(1.6,.6,0))

```




Fig. 6.8. A 12-quarter forecast for the Johnson & Johnson quarterly earnings series. The forecasts are shown as a continuation of the data (points connected by a solid line). The gray area represents three root MSPE bounds.

```
plot(Tsm, main="Trend Component", ylab="Trend")
xx=c(time(jj), rev(time(jj)))
yy=c(Tsm-pl, rev(Tsm+pl))
polygon(xx, yy, border=NA, col=gray(.4, alpha = .2))
plot(Ssm, main="Seasonal Component", ylim=c(-5,4), ylab="Season")
xx=c(time(jj), rev(time(jj)))
yy=c(Ssm-p2, rev(Ssm+p2))
polygon(xx, yy, border=NA, col=gray(.4, alpha = .2))
# Forecast
n.ahead=12; y = ts(append(jj, rep(0,n.ahead)), start=1960, freq=4)
rmspe = rep(0,n.ahead); x00 = ks$xf[, ,num]; P00 = ks$Pf[, ,num]
Q=t(cQ)%*%cQ; R=t(cR)%*%(cR)
for (m in 1:n.ahead){
  xp = Phi%*%x00; Pp = Phi%*%P00%*%t(Phi)+Q
  sig = A%*%Pp%*%t(A)+R; K = Pp%*%t(A)%*%(1/sig)
  x00 = xp; P00 = Pp-K%*%A%*%Pp
  y[num+m] = A%*%xp; rmspe[m] = sqrt(sig)
}
dev.new()
par(mar=c(3,3,1,1), mgp=c(1.6,.6,0))
plot(y, type="o", main="", ylab="", ylim=c(5,30), xlim=c(1975,1984))
upp = ts(y[(num+1):(num+n.ahead)]+3*rmspe, start=1981, freq=4)
low = ts(y[(num+1):(num+n.ahead)]-3*rmspe, start=1981, freq=4)
xx=c(time(low), rev(time(upp)))
yy=c(low, rev(upp))
polygon(xx, yy, border=NA, col=gray(.4, alpha = .2))
abline(v=1980.75, lty=3)
```

Note that the Cholesky decomposition of Q does not exist here, however, the diagonal form allows us to use standard deviations for the first two diagonal elements of cQ . This technicality can be avoided using a form of the model that we present in the next section.

6.7 State-Space Models with Correlated Errors

Sometimes it is advantageous to write the state-space model in a slightly different way, as is done by numerous authors; for example, Anderson and Moore (1979) and Hannan and Deistler (1988). Here, we write the state-space model as

$$x_{t+1} = \Phi x_t + \Upsilon u_{t+1} + \Theta w_t \quad t = 0, 1, \dots, n \quad (6.103)$$

$$y_t = A_t x_t + \Gamma u_t + v_t \quad t = 1, \dots, n \quad (6.104)$$

where, in the state equation, $x_0 \sim N_p(\mu_0, \Sigma_0)$, Φ is $p \times p$, and Υ is $p \times r$, Θ is $p \times m$ and $w_t \sim \text{iid } N_m(0, Q)$. In the observation equation, A_t is $q \times p$ and Γ is $q \times r$, and $v_t \sim \text{iid } N_q(0, R)$. In this model, while w_t and v_t are still white noise series (both independent of x_0), we also allow the state noise and observation noise to be correlated at time t ; that is,

$$\text{cov}(w_s, v_t) = S \delta_s^t, \quad (6.105)$$

where δ_s^t is Kronecker's delta; note that S is an $m \times q$ matrix. The major difference between this form of the model and the one specified by (6.3)–(6.4) is that this model starts the state noise process at $t = 0$ in order to ease the notation related to the concurrent covariance between w_t and v_t . Also, the inclusion of the matrix Θ allows us to avoid using a singular state noise process as was done in Example 6.11.

To obtain the innovations, $\epsilon_t = y_t - A_t x_t^{t-1} - \Gamma u_t$, and the innovation variance $\Sigma_t = A_t P_t^{t-1} A_t' + R$, in this case, we need the one-step-ahead state predictions. Of course, the filtered estimates will also be of interest, and they will be needed for smoothing. Property 6.2 (the smoother) as displayed in §6.2 still holds. The following property generates the predictor x_{t+1}^t from the past predictor x_t^{t-1} when the noise terms are correlated and exhibits the filter update.

Property 6.5 The Kalman Filter with Correlated Noise

For the state-space model specified in (6.103) and (6.104), with initial conditions x_1^0 and P_1^0 , for $t = 1, \dots, n$,

$$x_{t+1}^t = \Phi x_t^{t-1} + \Upsilon u_{t+1} + K_t \epsilon_t \quad (6.106)$$

$$P_{t+1}^t = \Phi P_t^{t-1} \Phi' + \Theta Q \Theta' - K_t \Sigma_t K_t' \quad (6.107)$$

where $\epsilon_t = y_t - A_t x_t^{t-1} - \Gamma u_t$ and the gain matrix is given by

$$K_t = [\Phi P_t^{t-1} A_t' + \Theta S][A_t P_t^{t-1} A_t' + R]^{-1}. \quad (6.108)$$

The filter values are given by

$$x_t^t = x_t^{t-1} + P_t^{t-1} A_t' [A_t P_t^{t-1} A_t' + R]^{-1} \epsilon_{t+1}, \quad (6.109)$$

$$P_t^t = P_t^{t-1} - P_t^{t-1} A_t' [A_t P_t^{t-1} A_t' + R]^{-1} A_t P_t^{t-1}. \quad (6.110)$$

The derivation of **Property 6.5** is similar to the derivation of the Kalman filter in **Property 6.1 (Problem 6.18)**; we note that the gain matrix K_t differs in the two properties. The filter values, (6.109)–(6.110), are symbolically identical to (6.19) and (6.20). To initialize the filter, we note that

$$x_1^0 = E(x_1) = \Phi\mu_0 + \Upsilon u_1, \quad \text{and} \quad P_1^0 = \text{var}(x_1) = \Phi\Sigma_0\Phi' + \Theta Q\Theta'.$$

In the next two subsections, we show how to use the model (6.103)–(6.104) for fitting ARMAX models and for fitting (multivariate) regression models with autocorrelated errors. To put it succinctly, for ARMAX models, the inputs enter in the state equation and for regression with autocorrelated errors, the inputs enter in the observation equation. It is, of course, possible to combine the two models and we give an example of this at the end of the section.

6.7.1 ARMAX Models

Consider a k -dimensional ARMAX model given by

$$y_t = \Upsilon u_t + \sum_{j=1}^p \Phi_j y_{t-j} + \sum_{k=1}^q \Theta_k v_{t-k} + v_t. \quad (6.111)$$

The observations y_t are a k -dimensional vector process, the Φ s and Θ s are $k \times k$ matrices, Υ is $k \times r$, u_t is the $r \times 1$ input, and v_t is a $k \times 1$ white noise process; in fact, (6.111) and (5.90) are identical models, but here, we have written the observations as y_t . We now have the following property.

Property 6.6 A State-Space Form of ARMAX

For $p \geq q$, let

$$F = \begin{bmatrix} \Phi_1 & I & 0 & \cdots & 0 \\ \Phi_2 & 0 & I & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \Phi_{p-1} & 0 & 0 & \cdots & I \\ \Phi_p & 0 & 0 & \cdots & 0 \end{bmatrix} \quad G = \begin{bmatrix} \Theta_1 + \Phi_1 \\ \vdots \\ \Theta_q + \Phi_q \\ \Phi_{q+1} \\ \vdots \\ \Phi_p \end{bmatrix} \quad H = \begin{bmatrix} \Upsilon \\ 0 \\ \vdots \\ 0 \end{bmatrix} \quad (6.112)$$

where F is $kp \times kp$, G is $kp \times k$, and H is $kp \times r$. Then, the state-space model given by

$$x_{t+1} = Fx_t + Hu_{t+1} + Gv_t, \quad (6.113)$$

$$y_t = Ax_t + v_t, \quad (6.114)$$

where $A = [I, 0, \dots, 0]$ is $k \times kp$ and I is the $k \times k$ identity matrix, implies the ARMAX model (6.111). If $p < q$, set $\Phi_{p+1} = \dots = \Phi_q = 0$, in which case $p = q$ and (6.113)–(6.114) still apply. Note that the state process is kp -dimensional, whereas the observations are k -dimensional.

This form of the model is somewhat different than the form suggested in §6.1, equations (6.6)–(6.8). For example, in (6.8), by setting A_t equal to the $p \times p$ identity matrix (for all t) and setting $R = 0$ implies the data y_t in (6.8) follow a VAR(m) process. In doing so, however, we do not make use of the ability to allow for correlated state and observation error, so a singularity is introduced into the system in the form of $R = 0$. The method in Property 6.6 avoids that problem, and points out the fact that the same model can take many forms. We do not prove Property 6.6 directly, but the following example should suggest how to establish the general result.

Example 6.12 Univariate ARMAX(1, 1) in State-Space Form

Consider the univariate ARMAX(1, 1) model

$$y_t = \alpha_t + \phi y_{t-1} + \theta v_{t-1} + v_t,$$

where $\alpha_t = \mathcal{Y}u_t$ to ease the notation. For a simple example, if $\mathcal{Y} = (\beta_0, \beta_1)$ and $u_t = (1, t)'$, the model for y_t would be ARMA(1,1) with linear trend, $y_t = \beta_0 + \beta_1 t + \phi y_{t-1} + \theta v_{t-1} + v_t$. Using Property 6.6, we can write the model as

$$x_{t+1} = \phi x_t + \alpha_{t+1} + (\theta + \phi)v_t, \quad (6.115)$$

and

$$y_t = x_t + v_t. \quad (6.116)$$

In this case, (6.115) is the state equation with $w_t \equiv v_t$ and (6.116) is the observation equation. Consequently, $\text{cov}(w_t, v_t) = \text{var}(v_t) = R$, and $\text{cov}(w_t, v_s) = 0$ when $s \neq t$, so Property 6.5 would apply. To verify (6.115) and (6.116) specify an ARMAX(1, 1) model, we have

$$\begin{aligned} y_t &= x_t + v_t && \text{from (6.116)} \\ &= \phi x_{t-1} + \alpha_t + (\theta + \phi)v_{t-1} + v_t && \text{from (6.115)} \\ &= \alpha_t + \phi(x_{t-1} + v_{t-1}) + \theta v_{t-1} + v_t && \text{rearrange terms} \\ &= \alpha_t + \phi y_{t-1} + \theta v_{t-1} + v_t, && \text{from (6.116).} \end{aligned}$$

Together, Property 6.5 and Property 6.6 can be used to accomplish maximum likelihood estimation as described in §6.3 for ARMAX models. The ARMAX model is only a special case of the model (6.103)–(6.104), which is quite rich, as will be discovered in the next subsection.

6.7.2 Multivariate Regression with Autocorrelated Errors

In regression with autocorrelated errors, we are interested in fitting the regression model

$$y_t = \Gamma u_t + \varepsilon_t \quad (6.117)$$

to a $k \times 1$ vector process, y_t , with r regressors $u_t = (u_{t1}, \dots, u_{tr})'$ where ε_t is vector ARMA(p, q) and Γ is a $k \times r$ matrix of regression parameters. We note that the regressors do not have to vary with time (e.g., $u_{t1} \equiv 1$ includes a constant in the regression) and that the case $k = 1$ was treated in §3.9.

To put the model in state-space form, we simply notice that $\varepsilon_t = y_t - \Gamma u_t$ is a k -dimensional ARMA(p, q) process. Thus, if we set $H = 0$ in (6.113), and include Γu_t in (6.114), we obtain

$$x_{t+1} = Fx_t + Gv_t, \quad (6.118)$$

$$y_t = \Gamma u_t + Ax_t + v_t, \quad (6.119)$$

where the model matrices A , F , and G are defined in [Property 6.6](#). The fact that (6.118)–(6.119) is multivariate regression with autocorrelated errors follows directly from [Property 6.6](#) by noticing that together, $x_{t+1} = Fx_t + Gv_t$ and $\varepsilon_t = Ax_t + v_t$ imply $\varepsilon_t = y_t - \Gamma u_t$ is vector ARMA(p, q).

As in the case of ARMAX models, regression with autocorrelated errors is a special case of the state-space model, and the results of [Property 6.5](#) can be used to obtain the innovations form of the likelihood for parameter estimation.

Example 6.13 Mortality, Temperature and Pollution

This example combines both techniques of [§6.7.1](#) and [§6.7.2](#). We will fit an ARMAX model to the detrended mortality series [cmort](#). The detrending part of the example constitutes the regression with autocorrelated errors.

Here, we let M_t denote the weekly cardiovascular mortality series, T_t as the corresponding temperature series [tempr](#), and P_t as the corresponding particulate series. A preliminary analysis suggests the following considerations (no output is shown):

- An AR(2) model fits well to detrended M_t :
`fit1 = sarima(cmort, 2, 0, 0, xreg=time(cmort))`
- The CCF between the mortality residuals, the temperature series and the particulates series, shows a strong correlation with temperature lagged one week (T_{t-1}), concurrent particulate level (P_t) and the particulate level about one month prior (P_{t-4}).
`acf(cbind(dmort <- resid(fit1$fit), tempr, part))`
`lag2.plot(tempr, dmort, 8)`
`lag2.plot(part, dmort, 8)`

From these results, we decided to fit the ARMAX model

$$\tilde{M}_t = \phi_1 \tilde{M}_{t-1} + \phi_2 \tilde{M}_{t-2} + \beta_1 T_{t-1} + \beta_2 P_t + \beta_3 P_{t-4} + v_t \quad (6.120)$$

to the detrended mortality series, $\tilde{M}_t = M_t - (\alpha + \beta_4 t)$, where $v_t \sim \text{iid } N(0, \sigma_v^2)$. To write the model in state-space form using [Property 6.6](#), let

$$x_{t+1} = \Phi x_t + \Upsilon u_{t+1} + \Theta v_t \quad t = 0, 1, \dots, n$$

$$y_t = \alpha + Ax_t + \Gamma u_t + v_t \quad t = 1, \dots, n$$

with

$$\Phi = \begin{bmatrix} \phi_1 & 1 \\ \phi_2 & 0 \end{bmatrix} \quad \Upsilon = \begin{bmatrix} \beta_1 & \beta_2 & \beta_3 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \quad \Theta = \begin{bmatrix} \phi_1 \\ \phi_2 \end{bmatrix}$$

$A = \begin{bmatrix} 1 & 0 \end{bmatrix}$, $\Gamma = \begin{bmatrix} 0 & 0 & 0 & \beta_4 \end{bmatrix}$, $u_t = (T_{t-1}, P_t, P_{t-4}, t)'$, $y_t = M_t$. We could have included α in Γ ; however, to reduce the number of parameters to be estimated numerically, we centered M_t by its sample mean and removed the constant α from the model. Note that the state process is bivariate and the observation process is univariate.

Some additional data analysis notes are: (1) t is centered by its mean; (2) P_t and P_{t-4} are highly correlated, so orthogonalizing these two inputs would be advantageous (although we did not do it here); (3) T_t and T_t^2 , as in Chapter 2, are not needed in the model when T_{t-1} is included; (4) initial values of the parameters are taken from the preliminary investigation.

The estimates and standard errors are displayed along with the following R code; investigation of the residuals shows that the model fits well.

```
dm = cmort - mean(cmort)           # center mortality
trend = time(cmort) - mean(time(cmort)) # center time
ded = ts.intersect(dm, u1=lag(tempr,-1), u2=part, u3=lag(part,-4), u4=trend,
  dframe=TRUE)
y = ded$dm; input = cbind(ded$u1, ded$u2, ded$u3, ded$u4)
num = length(y); A = array(c(1,0), dim = c(1,2,num))
# Function to Calculate Likelihood
Linn=function(para){
  phi1=para[1]; phi2=para[2]; cR=para[3]
  b1=para[4]; b2=para[5]; b3=para[6]; b4=para[7]
  mu0 = matrix(c(0,0), 2, 1); Sigma0 = diag(100, 2)
  Phi = matrix(c(phi1, phi2, 1, 0), 2, 2)
  Theta = matrix(c(phi1, phi2), 2, 2)
  Ups = matrix(c(b1, 0, b2, 0, b3, 0, 0, 0), 2, 4)
  Gam = matrix(c(0, 0, 0, b4), 1, 4); cQ = cR; S = cR^2
  kf = Kfilter2(num, y, A, mu0, Sigma0, Phi, Ups, Gam, Theta, cQ, cR, S,
    input)
  return(kf$like) }
# Estimation
phi1=.4; phi2=.4; cR=5; b1=-.1; b2=.1; b3=.1; b4=-1.5
init.par = c(phi1, phi2, cR, b1, b2, b3, b4) # initial parameters
L = c(0,0,0,-1,0,0,-2); U = c(1,1,10,0,1,1,0) # used in optim
est = optim(init.par, Linn, NULL, method="L-BFGS-B", lower=L, upper=U,
  hessian=TRUE, control=list(trace=1,REPORT=1))
SE = sqrt(diag(solve(est$hessian)))
# Results
u = cbind(estimate=est$par, SE)
rownames(u)=c("phi1", "phi2", "sigv", "TL1", "P", "PL4", "trnd")
round(u,3)
      estimate      SE
phi1      0.315  0.037
phi2      0.318  0.038
sigv      5.062  0.160
TL1      -0.119  0.011 (beta 1)
P          0.119  0.017 (beta 2)
PL4        0.067  0.018 (beta 3)
trnd      -1.339  0.220 (beta 4)
```

The residuals can be obtained by running `Kfilter2` again at the final estimates; they are returned as `innov`. For example,

```
phi1=est$par[1]; phi2=est$par[2]; cR=est$par[3]
```

```

b1=est$par[4]; b2=est$par[5]; b3=est$par[6]; b4=est$par[7]
mu0 = matrix(c(0,0), 2, 1); Sigma0 = diag(100, 2)
Phi = matrix(c(phi1, phi2, 1, 0), 2, 2)
Theta = matrix(c(phi1, phi2), 2)
Ups = matrix(c(b1, 0, b2, 0, b3, 0, 0, 0), 2, 4)
Gam = matrix(c(0, 0, 0, b4), 1, 4); cQ = cR; S = cR^2
kf = Kfilter2(num,y,A,mu0,Sigma0,Phi,Ups,Gam,Theta,cQ,cR,S,input)
res = ts(as.vector(kf$innov), start=1970, freq=52)
plot(res); acf(res,156)

```

6.8 Bootstrapping State-Space Models

Although in §6.3 we discussed the fact that under general conditions (which we assume to hold in this section) the MLEs of the parameters of a DLM are consistent and asymptotically normal, time series data are often of short or moderate length. Several researchers have found evidence that samples must be fairly large before asymptotic results are applicable (Dent and Min, 1978; Ansley and Newbold, 1980). Moreover, as we discussed in Example 3.35, problems occur if the parameters are near the boundary of the parameter space. In this section, we discuss an algorithm for bootstrapping state-space models; this algorithm and its justification, including the non-Gaussian case, along with numerous examples, can be found in Stoffer and Wall (1991) and in Stoffer and Wall (2004). In view of §6.7, anything we do or say here about DLMs applies equally to ARMAX models.

Using the DLM given by (6.103)–(6.105) and Property 6.5, we write the innovations form of the filter as

$$\epsilon_t = y_t - A_t x_t^{t-1} - \Gamma u_t, \quad (6.121)$$

$$\Sigma_t = A_t P_t^{t-1} A_t' + R, \quad (6.122)$$

$$K_t = [\Phi P_t^{t-1} A_t' + \Theta S] \Sigma_t^{-1}, \quad (6.123)$$

$$x_{t+1}^t = \Phi x_t^{t-1} + \Gamma u_{t+1} + K_t \epsilon_t, \quad (6.124)$$

$$P_{t+1}^t = \Phi P_t^{t-1} \Phi' + \Theta Q \Theta' - K_t \Sigma_t K_t'. \quad (6.125)$$

This form of the filter is just a rearrangement of the filter given in Property 6.5.

In addition, we can rewrite the model to obtain its innovations form,

$$x_{t+1}^t = \Phi x_t^{t-1} + \Gamma u_{t+1} + K_t \epsilon_t, \quad (6.126)$$

$$y_t = A_t x_t^{t-1} + \Gamma u_t + \epsilon_t. \quad (6.127)$$

This form of the model is a rewriting of (6.121) and (6.124), and it accommodates the bootstrapping algorithm.

As discussed in Example 6.5, although the innovations ϵ_t are uncorrelated, initially, Σ_t can be vastly different for different time points t . Thus, in a resampling procedure, we can either ignore the first few values of ϵ_t until Σ_t stabilizes or we can work with the standardized innovations

$$e_t = \Sigma_t^{-1/2} \epsilon_t, \quad (6.128)$$

so we are guaranteed these innovations have, at least, the same first two moments. In (6.128), $\Sigma_t^{1/2}$ denotes the unique square root matrix of Σ_t defined by $\Sigma_t^{1/2} \Sigma_t^{1/2} = \Sigma_t$. In what follows, we base the bootstrap procedure on the standardized innovations, but we stress the fact that, even in this case, ignoring startup values might be necessary, as noted by Stoffer & Wall (1991).

The model coefficients and the correlation structure of the model are uniquely parameterized by a $k \times 1$ parameter vector θ_0 ; that is, $\Phi = \Phi(\theta_0)$, $\gamma = \gamma(\theta_0)$, $Q = Q(\theta_0)$, $A_t = A_t(\theta_0)$, $\Gamma = \Gamma(\theta_0)$, and $R = R(\theta_0)$. Recall the innovations form of the Gaussian likelihood (ignoring a constant) is

$$\begin{aligned} -2 \ln L_Y(\theta) &= \sum_{t=1}^n \left[\ln |\Sigma_t(\theta)| + \epsilon_t(\theta)' \Sigma_t(\theta)^{-1} \epsilon_t(\theta) \right] \\ &= \sum_{t=1}^n \left[\ln |\Sigma_t(\theta)| + e_t(\theta)' e_t(\theta) \right]. \end{aligned} \quad (6.129)$$

We stress the fact that it is not necessary for the model to be Gaussian to consider (6.129) as the criterion function to be used for parameter estimation.

Let $\hat{\theta}$ denote the MLE of θ_0 , that is, $\hat{\theta} = \operatorname{argmax}_{\theta} L_Y(\theta)$, obtained by the methods discussed in §6.3. Let $\epsilon_t(\hat{\theta})$ and $\Sigma_t(\hat{\theta})$ be the innovation values obtained by running the filter, (6.121)–(6.125), under $\hat{\theta}$. Once this has been done, the bootstrap procedure is accomplished by the following steps.

- (i) Construct the standardized innovations

$$e_t(\hat{\theta}) = \Sigma_t^{-1/2}(\hat{\theta}) \epsilon_t(\hat{\theta}).$$

- (ii) Sample, with replacement, n times from the set $\{e_1(\hat{\theta}), \dots, e_n(\hat{\theta})\}$ to obtain $\{e_1^*(\hat{\theta}), \dots, e_n^*(\hat{\theta})\}$, a bootstrap sample of standardized innovations.
 (iii) Construct a bootstrap data set $\{y_1^*, \dots, y_n^*\}$ as follows. Define the $(p+q) \times 1$ vector $\xi_t = (x_{t+1}', y_t')'$. Stacking (6.126) and (6.127) results in a vector first-order equation for ξ_t given by

$$\xi_t = F_t \xi_{t-1} + G u_t + H_t e_t, \quad (6.130)$$

where

$$F_t = \begin{bmatrix} \Phi & 0 \\ A_t & 0 \end{bmatrix}, \quad G = \begin{bmatrix} \gamma \\ \Gamma \end{bmatrix}, \quad H_t = \begin{bmatrix} K_t \Sigma_t^{1/2} \\ \Sigma_t^{1/2} \end{bmatrix}.$$

Thus, to construct the bootstrap data set, solve (6.130) using $e_t^*(\hat{\theta})$ in place of e_t . The exogenous variables u_t and the initial conditions of the Kalman filter remain fixed at their given values, and the parameter vector is held fixed at $\hat{\theta}$.

- (iv) Using the bootstrap data set $\{y_t^*; t = 1, \dots, n\}$, construct a likelihood, $L_{Y^*}(\theta)$, and obtain the MLE of θ , say, $\hat{\theta}^*$.




Fig. 6.9. Quarterly interest rate for Treasury bills (dashed line) and quarterly inflation rate (solid line) in the Consumer Price Index.

- (v) Repeat steps 2 through 4, a large number, B , of times, obtaining a bootstrapped set of parameter estimates $\{\hat{\theta}_b^*; b = 1, \dots, B\}$. The finite sample distribution of $\hat{\theta} - \theta_0$ may be approximated by the distribution of $\hat{\theta}_b^* - \hat{\theta}$, $b = 1, \dots, B$.

In the next example, we discuss the case of a linear regression model, but where the regression coefficients are stochastic and allowed to vary with time. The state-space model provides a convenient setting for the analysis of such models.

Example 6.14 Stochastic Regression

Figure 6.9 shows the quarterly inflation rate (solid line), y_t , in the Consumer Price Index and the quarterly interest rate recorded for Treasury bills (dashed line), z_t , from the first quarter of 1953 through the second quarter of 1980, $n = 110$ observations. These data are taken from Newbold and Bos (1985).

In this example, we consider one analysis that was discussed in Newbold and Bos (1985, pp. 61-73), that focused on the first 50 observations and where quarterly inflation was modeled as being stochastically related to quarterly interest rate,

$$y_t = \alpha + \beta_t z_t + v_t,$$

where α is a fixed constant, β_t is a stochastic regression coefficient, and v_t is white noise with variance σ_v^2 . The stochastic regression term, which comprises the state variable, is specified by a first-order autoregression,

$$(\beta_t - b) = \phi(\beta_{t-1} - b) + w_t,$$

where b is a constant, and w_t is white noise with variance σ_w^2 . The noise processes, v_t and w_t , are assumed to be uncorrelated.

Using the notation of the state-space model (6.103) and (6.104), we have in the state equation, $x_t = \beta_t$, $\Phi = \phi$, $u_t \equiv 1$, $\Upsilon = (1 - \phi)b$, $Q = \sigma_w^2$, and in the

Table 6.2. Comparison of Standard Errors

Parameter	Asymptotic		Bootstrap
	MLE	Standard Error	Standard Error
ϕ	.865	.223	.463
α	-.686	.487	.557
b	.788	.226	.821
σ_w	.115	.107	.216
σ_v	1.135	.147	.340

observation equation, $A_t = z_t$, $\Gamma = \alpha$, $R = \sigma_v^2$, and $S = 0$. The parameter vector is $\Theta = (\phi, \alpha, b, \sigma_w, \sigma_v)'$. The results of the Newton–Raphson estimation procedure are listed in Table 6.2. Also shown in the Table 6.2 are the corresponding standard errors obtained from $B = 500$ runs of the bootstrap. These standard errors are simply the standard deviations of the bootstrapped estimates, that is, the square root of $\sum_{b=1}^B (\Theta_{ib}^* - \bar{\Theta}_i^*)^2 / (B - 1)$, where Θ_{ib} represents the i th parameter, $i = 1, \dots, 5$, and $\bar{\Theta}_i^* = \sum_{b=1}^B \Theta_{ib}^* / B$.

The asymptotic standard errors listed in Table 6.2 are typically much smaller than those obtained from the bootstrap. For most of the cases, the bootstrapped standard errors are at least 50% larger than the corresponding asymptotic value. Also, asymptotic theory prescribes the use of normal theory when dealing with the parameter estimates. The bootstrap, however, allows us to investigate the small sample distribution of the estimators and, hence, provides more insight into the data analysis.

For example, Figure 6.10 shows the bootstrap distribution of the estimator of ϕ in the upper left-hand corner. This distribution is highly skewed with values concentrated around .8, but with a long tail to the left. Some quantiles are -.09 (5%), .11 (10%), .34 (25%), .73 (50%), .86 (75%), .96 (90%), .98 (95%), and they can be used to obtain confidence intervals. For example, a 90% confidence interval for ϕ would be approximated by $(-.09, .96)$. This interval is ridiculously wide and includes 0 as a plausible value of ϕ ; we will interpret this after we discuss the results of the estimation of σ_w .

Figure 6.10 shows the bootstrap distribution of $\hat{\sigma}_w$ in the lower right-hand corner. The distribution is concentrated at two locations, one at approximately $\hat{\sigma}_w = .25$ (which is the median of the distribution of values away from 0) and the other at $\hat{\sigma}_w = 0$. The cases in which $\hat{\sigma}_w \approx 0$ correspond to deterministic state dynamics. When $\sigma_w = 0$ and $|\phi| < 1$, then $\beta_t \approx b$ for large t , so the approximately 25% of the cases in which $\hat{\sigma}_w \approx 0$ suggest a fixed state, or constant coefficient model. The cases in which $\hat{\sigma}_w$ is away from zero would suggest a truly stochastic regression parameter. To investigate this matter further, the off-diagonals of Figure 6.10 show the joint bootstrapped estimates, $(\hat{\phi}, \hat{\sigma}_w)$, for positive values of $\hat{\phi}^*$. The joint distribution suggests $\hat{\sigma}_w > 0$ corresponds to $\hat{\phi} \approx 0$. When $\phi = 0$, the state dynamics are given by $\beta_t = b + w_t$. If, in addition, σ_w is small relative to b , the system is nearly deterministic; that is, $\beta_t \approx b$. Considering these results, the




Fig. 6.10. Bootstrap distribution, $B = 500$, of (i) the estimator of ϕ (upper left), (ii) the estimator of σ_w (lower right), and (iii) the estimators jointly (off-diagonals). Only the values corresponding to $\hat{\phi}^* \geq 0$ are shown.

bootstrap analysis leads us to conclude the dynamics of the data are best described in terms of a fixed regression effect.

The following R code was used for this example. We note that the first line of the code sets the relative tolerance for determining convergence for the numerical optimization. *Using the current default setting may result in a long run time of the algorithm* and we suggest the value be decreased on slower machines or for demonstration purposes. Also, decreasing the number of bootstrap replicates will improve computation time; for example, setting `tol=.001` and `nboot=200` yields reasonable results. In this example, we fix the first three values of the data for the resampling scheme.

```
tol = sqrt(.Machine$double.eps) # convergence tolerance !!!
nboot = 500                     # number of bootstrap replicates !!!
y = window(qinfl, c(1953,1), c(1965,2)) # inflation
z = window(qintr, c(1953,1), c(1965,2)) # interest
num = length(y); input = matrix(1, num, 1)
A = array(z, dim=c(1,1,num))
# Function to Calculate Likelihood
Linn=function(para){
  phi=para[1]; alpha=para[2]; b=para[3]; Ups=(1-phi)*b
  cQ=para[4]; cR=para[5]
  kf=Kfilter2(num,y,A,mu0,Sigma0,phi,Ups,alpha,1,cQ,cR,0,input)
  return(kf$like) }
# Parameter Estimation
mu0=1; Sigma0=.01; phi=.84; alpha=-.77; b=.85; cQ=.12; cR=1.1
```

```

init.par = c(phi, alpha, b, cQ, cR) # initial parameters
est = optim(init.par, Linn, NULL, method="BFGS", hessian=TRUE,
  control=list(trace=1,REPORT=1,reltol=tol))
SE = sqrt(diag(solve(est$hessian)))
phi = est$par[1]; alpha=est$par[2]; b=est$par[3]; Ups=(1-phi)*b
cQ=est$par[4]; cR=est$par[5]
rbind(estimate=est$par, SE)

```

	phi	alpha	b	sigma_w	sigma_v
estimate	0.8653348	-0.6855891	0.7879308	0.1145682	1.1353139
SE	0.2231382	0.4865775	0.2255649	0.1071838	0.1472067

```

# BEGIN BOOTSTRAP
Linn2=function(para){ # likelihood for bootstrapped data
  phi=para[1]; alpha=para[2]; b=para[3]; Ups=(1-phi)*b
  cQ=para[4]; cR=para[5]
  kf=Kfilter2(num,y.star,A,mu0,Sigma0,phi,Ups,alpha,1,cQ,cR,0,input)
  return(kf$like) }
# Run the filter at the estimates
kf=Kfilter2(num,y,A,mu0,Sigma0,phi,Ups,alpha,1,cQ,cR,0,input)
# Pull out necessary values from the filter and initialize
xp=kf$xp; innov=kf$innov; sig=kf$sig; K=kf$K; e=innov/sqrt(sig)
e.star=e; y.star=y; xp.star=xp; k=4:50
para.star = matrix(0, nboot, 5) # to store estimates
init.par=c(.84,-.77,.85,.12,1.1)
for (i in 1:nboot){cat("iteration:", i, "\n")
  e.star[k] = sample(e[k], replace=TRUE)
  for (j in k){
    xp.star[j] = phi*xp.star[j-1]+Ups+K[j]*sqrt(sig[j])*e.star[j] }
  y.star[k] = z[k]*xp.star[k]+alpha+sqrt(sig[k])*e.star[k]
  est.star = optim(init.par, Linn2, NULL, method="BFGS",
    control=list(reltol=tol))
  para.star[i,] = cbind(est.star$par[1], est.star$par[2], est.star$par[3],
    abs(est.star$par[4]), abs(est.star$par[5]))}
rmse = rep(NA,5) # compute bootstrapped SEs (Table 6.2)
for(i in 1:5){
  rmse[i]=sqrt(sum((para.star[,i]-est$par[i])^2)/nboot)
  cat(i, rmse[i],"\n") }
  1 0.46294 2 0.55698 3 0.82148 4 0.21595 5 0.34011
# Plot for phi vs sigma_w
phi = para.star[,1]; sigw = abs(para.star[,4])
phi = ifelse(phi<0, NA, phi) # any phi < 0 not plotted
panel.hist <- function(x, ...){
  usr <- par("usr"); on.exit(par(usr))
  par(usr = c(usr[1:2], 0, 1.5) )
  h <- hist(x, plot = FALSE)
  breaks <- h$breaks; nB <- length(breaks)
  y <- h$counts; y <- y/max(y)
  rect(breaks[-nB], 0, breaks[-1], y, ...)}
u = cbind(phi, sigw); colnames(u) = c("f","s")
pairs(u, cex=1.5, pch=1, diag.panel=panel.hist, cex.labels=1.5,
  font.labels=5)

```

6.9 Dynamic Linear Models with Switching

The problem of modeling changes in regimes for vector-valued time series has been of interest in many different fields. In §5.5, we explored the idea that the dynamics of the system of interest might change over the course of time. In Example 5.6, we saw that pneumonia and influenza mortality rates behave differently when a flu epidemic occurs than when no epidemic occurs. As another example, some authors (for example, Hamilton, 1989, or McCulloch and Tsay, 1993) have explored the possibility the dynamics of the quarterly U.S. GNP series (say, y_t) analyzed in Example 3.35 might be different during expansion ($\nabla \log y_t > 0$) than during contraction ($\nabla \log y_t < 0$). In this section, we will concentrate on the method presented in Shumway and Stoffer (1991). One way of modeling change in an evolving time series is by assuming the dynamics of some underlying model changes discontinuously at certain undetermined points in time. Our starting point is the DLM given by (6.1) and (6.2), namely,

$$x_t = \Phi x_{t-1} + w_t, \quad (6.131)$$

to describe the $p \times 1$ state dynamics, and

$$y_t = A_t x_t + v_t \quad (6.132)$$

to describe the $q \times 1$ observation dynamics. Recall w_t and v_t are Gaussian white noise sequences with $\text{var}(w_t) = Q$, $\text{var}(v_t) = R$, and $\text{cov}(w_t, v_s) = 0$ for all s and t .

Generalizations of (6.131) and (6.132) to include the possibility of changes occurring over time have been approached by allowing changes in the error covariances (Harrison and Stevens, 1976, Gordon and Smith, 1988, 1990) or by assigning mixture distributions to the observation errors v_t (Peña and Guttman, 1988). Approximations to filtering were derived in all of the aforementioned articles. An application to monitoring renal transplants was described in Smith and West (1983) and in Gordon and Smith (1990). Changes can also be modeled in the classical regression case by allowing switches in the design matrix, as in Quandt (1972).

Switching via a stationary Markov chain with independent observations has been developed by Lindgren (1978) and Goldfeld and Quandt (1973). In the Markov chain approach, we declare the dynamics of the system at time t are generated by one of m possible regimes evolving according to a Markov chain over time. As a simple example, suppose the dynamics of a univariate time series, y_t , is generated by either the model (1) $y_t = \beta_1 y_{t-1} + w_t$ or the model (2) $y_t = \beta_2 y_{t-1} + w_t$. We will write the model as $y_t = \phi_t y_{t-1} + w_t$ such that $\Pr(\phi_t = \beta_j) = \pi_j$, $j = 1, 2$, $\pi_1 + \pi_2 = 1$, and with the Markov property

$$\Pr(\phi_t = \beta_j \mid \phi_{t-1} = \beta_i, \phi_{t-2} = \beta_{i_2}, \dots) = \Pr(\phi_t = \beta_j \mid \phi_{t-1} = \beta_i) = \pi_{ij},$$

for $i, j = 1, 2$ (and $i_2, \dots = 1, 2$). As previously mentioned, Markov switching for dependent data has been applied by Hamilton (1989) to detect changes between positive and negative growth periods in the economy. Applications to speech recognition have been considered by Juang and Rabiner (1985). The case in which the particular

regime is unknown to the observer comes under the heading of hidden Markov models, and the techniques related to analyzing these models are summarized in Rabiner and Juang (1986). An application of the idea of switching to the tracking of multiple targets has been considered in Bar-Shalom (1978), who obtained approximations to Kalman filtering in terms of weighted averages of the innovations.

Example 6.15 Tracking Multiple Targets

The approach of Shumway and Stoffer (1991) was motivated primarily by the problem of tracking a large number of moving targets using a vector y_t of sensors. In this problem, we do not know at any given point in time which target any given sensor has detected. Hence, it is the structure of the measurement matrix A_t in (6.132) that is changing, and not the dynamics of the signal x_t or the noises, w_t or v_t . As an example, consider a 3×1 vector of satellite measurements $y_t = (y_{t1}, y_{t2}, y_{t3})'$ that are observations on some combination of a 3×1 vector of targets or signals, $x_t = (x_{t1}, x_{t2}, x_{t3})'$. For the measurement matrix

$$A_t = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

for example, the first sensor, y_{t1} , observes the second target, x_{t2} ; the second sensor, y_{t2} , observes the first target, x_{t1} ; and the third sensor, y_{t3} , observes the third target, x_{t3} . All possible detection configurations will define a set of possible values for A_t , say, $\{M_1, M_2, \dots, M_m\}$, as a collection of plausible measurement matrices.

Example 6.16 Modeling Economic Change

As another example of the switching model presented in this section, consider the case in which the dynamics of the linear model changes suddenly over the history of a given realization. For example, Lam (1990) has given the following generalization of Hamilton's (1989) model for detecting positive and negative growth periods in the economy. Suppose the data are generated by

$$y_t = z_t + n_t, \quad (6.133)$$

where z_t is an autoregressive series and n_t is a random walk with a drift that switches between two values α_0 and $\alpha_0 + \alpha_1$. That is,

$$n_t = n_{t-1} + \alpha_0 + \alpha_1 S_t, \quad (6.134)$$

with $S_t = 0$ or 1, depending on whether the system is in state 1 or state 2. For the purpose of illustration, suppose

$$z_t = \phi_1 z_{t-1} + \phi_2 z_{t-2} + w_t \quad (6.135)$$

is an AR(2) series with $\text{var}(w_t) = \sigma_w^2$. Lam (1990) wrote (6.133) in a differenced form

$$\nabla y_t = z_t - z_{t-1} + \alpha_0 + \alpha_1 S_t, \quad (6.136)$$

which we may take as the observation equation (6.132) with state vector

$$x_t = (z_t, z_{t-1}, \alpha_0, \alpha_1)' \quad (6.137)$$

and

$$M_1 = [1, -1, 1, 0] \quad \text{and} \quad M_2 = [1, -1, 1, 1] \quad (6.138)$$

determining the two possible economic conditions. The state equation, (6.131), is of the form

$$\begin{pmatrix} z_t \\ z_{t-1} \\ \alpha_0 \\ \alpha_1 \end{pmatrix} = \begin{bmatrix} \phi_1 & \phi_2 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{pmatrix} z_{t-1} \\ z_{t-2} \\ \alpha_0 \\ \alpha_1 \end{pmatrix} + \begin{pmatrix} w_t \\ 0 \\ 0 \\ 0 \end{pmatrix}. \quad (6.139)$$

The observation equation, (6.136), can be written as

$$\nabla y_t = A_t x_t + v_t, \quad (6.140)$$

where we have included the possibility of observational noise, and where $\Pr(A_t = M_1) = 1 - \Pr(A_t = M_2)$, with M_1 and M_2 given in (6.138).

To incorporate a reasonable switching structure for the measurement matrix into the DLM that is compatible with both practical situations previously described, we assume that the m possible configurations are states in a nonstationary, independent process defined by the time-varying probabilities

$$\pi_j(t) = \Pr(A_t = M_j), \quad (6.141)$$

for $j = 1, \dots, m$ and $t = 1, 2, \dots, n$. Important information about the current state of the measurement process is given by the filtered probabilities of being in state j , defined as the conditional probabilities

$$\pi_j(t|t) = \Pr(A_t = M_j | Y_t), \quad (6.142)$$

which also vary as a function of time. In (6.142), we have used the notation $Y_s = \{y_1, \dots, y_s\}$. The filtered probabilities (6.142) give the time-varying estimates of the probability of being in state j given the data to time t .

It will be important for us to obtain estimators of the configuration probabilities, $\pi_j(t|t)$, the predicted and filtered state estimators, x_t^{t-1} and x_t^t , and the corresponding error covariance matrices P_t^{t-1} and P_t^t . Of course, the predictor and filter estimators will depend on the parameters, Θ , of the DLM. In many situations, the parameters will be unknown and we will have to estimate them. Our focus will be on maximum likelihood estimation, but other authors have taken a Bayesian approach that assigns priors to the parameters, and then seeks posterior distributions of the model parameters; see, for example, Gordon and Smith (1990), Peña and Guttman (1988), or McCulloch and Tsay (1993).

We now establish the recursions for the filters associated with the state x_t and the switching process, A_t . As discussed in §6.3, the filters are also an essential part of the maximum likelihood procedure. The predictors, $x_t^{t-1} = E(x_t | \mathcal{Y}_{t-1})$, and filters, $x_t^t = E(x_t | \mathcal{Y}_t)$, and their associated error variance–covariance matrices, P_t^{t-1} and P_t^t , are given by

$$x_t^{t-1} = \Phi x_{t-1}^{t-1}, \quad (6.143)$$

$$P_t^{t-1} = \Phi P_{t-1}^{t-1} \Phi' + Q, \quad (6.144)$$

$$x_t^t = x_t^{t-1} + \sum_{j=1}^m \pi_j(t|t) K_{tj} \epsilon_{tj}, \quad (6.145)$$

$$P_t^t = \sum_{j=1}^m \pi_j(t|t) (I - K_{tj} M_j) P_t^{t-1}, \quad (6.146)$$

$$K_{tj} = P_t^{t-1} M_j' \Sigma_{tj}^{-1}, \quad (6.147)$$

where the innovation values in (6.145) and (6.147) are

$$\epsilon_{tj} = y_t - M_j x_t^{t-1}, \quad (6.148)$$

$$\Sigma_{tj} = M_j P_t^{t-1} M_j' + R, \quad (6.149)$$

for $j = 1, \dots, m$.

Equations (6.143)–(6.147) exhibit the filter values as weighted linear combinations of the m innovation values, (6.148)–(6.149), corresponding to each of the possible measurement matrices. The equations are similar to the approximations introduced by Bar-Shalom and Tse (1975), by Gordon and Smith (1990), and Peña and Guttman (1988).

To verify (6.145), let the indicator $I(A_t = M_j) = 1$ when $A_t = M_j$, and zero otherwise. Then, using (6.21),

$$\begin{aligned} x_t^t &= E(x_t | \mathcal{Y}_t) = E[E(x_t | Y_t, A_t) | \mathcal{Y}_t] \\ &= E\left\{ \sum_{j=1}^m E(x_t | Y_t, A_t = M_j) I(A_t = M_j) | \mathcal{Y}_t \right\} \\ &= E\left\{ \sum_{j=1}^m [x_t^{t-1} + K_{tj}(y_t - M_j x_t^{t-1})] I(A_t = M_j) | \mathcal{Y}_t \right\} \\ &= \sum_{j=1}^m \pi_j(t|t) [x_t^{t-1} + K_{tj}(y_t - M_j x_t^{t-1})], \end{aligned}$$

where K_{tj} is given by (6.147). Equation (6.146) is derived in a similar fashion; the other relationships, (6.143), (6.144), and (6.147), follow from straightforward applications of the Kalman filter results given in **Property 6.1**.

Next, we derive the filters $\pi_j(t|t)$. Let $f_j(t|t-1)$ denote the conditional density of y_t given the past $y_1, \dots, y_{t-1}$, and $A_t = M_j$, for $j = 1, \dots, m$. Then,

$$\pi_j(t|t) = \frac{\pi_j(t)f_j(t|t-1)}{\sum_{k=1}^m \pi_k(t)f_k(t|t-1)}, \quad (6.150)$$

where we assume the distribution $\pi_j(t)$, for $j = 1, \dots, m$ has been specified before observing $y_1, \dots, y_t$ (details follow as in [Example 6.17](#) below). If the investigator has no reason to prefer one state over another at time t , the choice of uniform priors, $\pi_j(t) = m^{-1}$, for $j = 1, \dots, m$, will suffice. Smoothness can be introduced by letting

$$\pi_j(t) = \sum_{i=1}^m \pi_i(t-1|t-1)\pi_{ij}, \quad (6.151)$$

where the non-negative weights π_{ij} are chosen so $\sum_{i=1}^m \pi_{ij} = 1$. If the A_t process was Markov with transition probabilities π_{ij} , then (6.151) would be the update for the filter probability, as shown in the next example.

Example 6.17 Hidden Markov Chain Model

If $\{A_t\}$ is a hidden Markov chain with stationary transition probabilities $\pi_{ij} = \Pr(A_t = M_j | A_{t-1} = M_i)$, for $i, j = 1, \dots, m$, letting $p(\cdot)$ denote a generic probability function, we have

$$\begin{aligned} \pi_j(t|t) &= \frac{p(A_t = M_j, y_t, \mathcal{Y}_{t-1})}{p(y_t, \mathcal{Y}_{t-1})} \\ &= \frac{p(Y_{t-1})p(A_t = M_j | \mathcal{Y}_{t-1})p(y_t | A_t = M_j, \mathcal{Y}_{t-1})}{p(Y_{t-1})p(y_t | \mathcal{Y}_{t-1})} \\ &= \frac{\pi_j(t|t-1)f_j(t|t-1)}{\sum_{k=1}^m \pi_k(t|t-1)f_k(t|t-1)}. \end{aligned} \quad (6.152)$$

In the Markov case, the conditional probabilities

$$\pi_j(t|t-1) = \Pr(A_t = M_j | \mathcal{Y}_{t-1})$$

in (6.152) replace the unconditional probabilities, $\pi_j(t) = \Pr(A_t = M_j)$, in (6.150).

To evaluate (6.152), we must be able to calculate $\pi_j(t|t-1)$ and $f_j(t|t-1)$. We will discuss the calculation of $f_j(t|t-1)$ after this example. To derive $\pi_j(t|t-1)$, note,

$$\begin{aligned} \pi_j(t|t-1) &= \Pr(A_t = M_j | \mathcal{Y}_{t-1}) \\ &= \sum_{i=1}^m \Pr(A_t = M_j, A_{t-1} = M_i | \mathcal{Y}_{t-1}) \\ &= \sum_{i=1}^m \Pr(A_t = M_j | A_{t-1} = M_i) \Pr(A_{t-1} = M_i | \mathcal{Y}_{t-1}) \\ &= \sum_{i=1}^m \pi_{ij} \pi_i(t-1|t-1). \end{aligned} \quad (6.153)$$

Expression (6.151) comes from equation (6.153), where, as previously noted, we replace $\pi_j(t|t-1)$ by $\pi_j(t)$.

The difficulty in extending the approach here to the Markov case is the dependence among the y_t , which makes it necessary to enumerate over all possible histories to derive the filtering equations. This problem will be evident when we derive the conditional density $f_j(t|t-1)$. Equation (6.151) has $\pi_j(t)$ as a function of the past observations, $\mathcal{Y}_{t-1}$, which is inconsistent with our model assumption. Nevertheless, this seems to be a reasonable compromise that allows the data to modify the probabilities $\pi_j(t)$, without having to develop a highly computer-intensive technique.

As previously suggested, the computation of $f_j(t|t-1)$, without some approximations, is highly computer-intensive. To evaluate $f_j(t|t-1)$, consider the event

$$A_1 = M_{j_1}, \dots, A_{t-1} = M_{j_{t-1}}, \quad (6.154)$$

for $j_i = 1, \dots, m$, and $i = 1, \dots, t-1$, which specifies a specific set of measurement matrices through the past; we will write this event as $A_{(t-1)} = M_{(\ell)}$. Because m^{t-1} possible outcomes exist for $A_1, \dots, A_{t-1}$, the index ℓ runs through $\ell = 1, \dots, m^{t-1}$. Using this notation, we may write

$$\begin{aligned} f_j(t|t-1) &= \sum_{\ell=1}^{m^{t-1}} \Pr\{A_{(t-1)} = M_{(\ell)} \mid \mathcal{Y}_{t-1}\} f(y_t \mid \mathcal{Y}_{t-1}, A_t = M_j, A_{(t-1)} = M_{(\ell)}) \\ &\equiv \sum_{\ell=1}^{m^{t-1}} \alpha(\ell) N(y_t \mid \mu_{tj}(\ell), \Sigma_{tj}(\ell)), \quad j = 1, \dots, m, \end{aligned} \quad (6.155)$$

where the notation $N(\cdot \mid b, B)$ represents the normal density with mean vector b and variance-covariance matrix B . That is, $f_j(t|t-1)$ is a mixture of normals with non-negative weights $\alpha(\ell) = \Pr\{A_{(t-1)} = M_{(\ell)} \mid \mathcal{Y}_{t-1}\}$ such that $\sum_{\ell} \alpha(\ell) = 1$, and with each normal distribution having mean vector

$$\mu_{tj}(\ell) = M_j x_t^{t-1}(\ell) = M_j E[x_t \mid \mathcal{Y}_{t-1}, A_{(t-1)} = M_{(\ell)}] \quad (6.156)$$

and covariance matrix

$$\Sigma_{tj}(\ell) = M_j P_t^{t-1}(\ell) M_j' + R. \quad (6.157)$$

This result follows because the conditional distribution of y_t in (6.155) is identical to the fixed measurement matrix case presented in §4.2. The values in (6.156) and (6.157), and hence the densities, $f_j(t|t-1)$, for $j = 1, \dots, m$, can be obtained directly from the Kalman filter, **Property 6.1**, with the measurement matrices $A_{(t-1)}$ fixed at $M_{(\ell)}$.

Although $f_j(t|t-1)$ is given explicitly in (6.155), its evaluation is highly computer intensive. For example, with $m = 2$ states and $n = 20$ observations, we have to filter over $2 + 2^2 + \dots + 2^{20}$ possible sample paths; note, $2^{20} = 1,048,576$. One remedy is to trim (remove), at each t , highly improbable sample paths; that is, remove events in (6.154) with extremely small probability of occurring, and then evaluate $f_j(t|t-1)$ as if the trimmed sample paths could not have occurred. Another alternative, as suggested by Gordon and Smith (1990) and Shumway and Stoffer

(1991), is to approximate $f_j(t|t-1)$ using the closest (in the sense of Kulback–Leibler distance) normal distribution. In this case, the approximation leads to choosing normal distribution with the same mean and variance associated with $f_j(t|t-1)$; that is, we approximate $f_j(t|t-1)$ by a normal with mean $M_j x_t^{t-1}$ and variance Σ_{tj} given in (6.149).

To develop a procedure for maximum likelihood estimation, the joint density of the data is

$$\begin{aligned} f(y_1, \dots, y_n) &= \prod_{t=1}^n f(y_t | \mathcal{Y}_{t-1}) \\ &= \prod_{t=1}^n \sum_{j=1}^m \Pr(A_t = M_j | \mathcal{Y}_{t-1}) f(y_t | A_t = M_j, \mathcal{Y}_{t-1}), \end{aligned}$$

and hence, the likelihood can be written as

$$\ln L_Y(\Theta) = \sum_{t=1}^n \ln \left(\sum_{j=1}^m \pi_j(t) f_j(t|t-1) \right). \quad (6.158)$$

For the hidden Markov model, $\pi_j(t)$ would be replaced by $\pi_j(t|t-1)$. In (6.158), we will use the normal approximation to $f_j(t|t-1)$. That is, henceforth, we will consider $f_j(t|t-1)$ as the normal, $N(M_j x_t^{t-1}, \Sigma_{tj})$, density, where x_t^{t-1} is given in (6.143) and Σ_{tj} is given in (6.149). We may consider maximizing (6.158) directly as a function of the parameters $\Theta = \{\mu_0, \Phi, Q, R\}$ using a Newton method, or we may consider applying the EM algorithm to the complete data likelihood.

To apply the EM algorithm as in §6.3, we call $x_0, x_1, \dots, x_n, A_1, \dots, A_n$, and $y_1, \dots, y_n$, the complete data, with likelihood given by

$$\begin{aligned} -2 \ln L_{X,A,Y}(\Theta) &= \ln |\Sigma_0| + (x_0 - \mu_0)' \Sigma_0^{-1} (x_0 - \mu_0) \\ &\quad + n \ln |Q| + \sum_{t=1}^n (x_t - \Phi x_{t-1})' Q^{-1} (x_t - \Phi x_{t-1}) \\ &\quad - 2 \sum_{t=1}^n \sum_{j=1}^m I(A_t = M_j) \ln \pi_j(t) + n \ln |R| \\ &\quad + \sum_{t=1}^n \sum_{j=1}^m I(A_t = M_j) (y_t - A_t x_t)' R^{-1} (y_t - A_t x_t). \end{aligned} \quad (6.159)$$

As discussed in §6.3, we require the minimization of the conditional expectation

$$Q(\Theta | \Theta^{(k-1)}) = E \left\{ -2 \ln L_{X,A,Y}(\Theta) \mid \mathcal{Y}_n, \Theta^{(k-1)} \right\}, \quad (6.160)$$

with respect to Θ at each iteration, $k = 1, 2, \dots$. The calculation and maximization of (6.160) is similar to the case of (6.66). In particular, with

$$\pi_j(t|n) = E[I(A_t = M_j) \mid \mathcal{Y}_n], \quad (6.161)$$

we obtain on iteration k ,

$$\pi_j^{(k)}(t) = \pi_j(t|n), \quad (6.162)$$

$$\mu_0^{(k)} = x_0^n, \quad (6.163)$$

$$\Phi^{(k)} = S_{10}S_{00}^{-1}, \quad (6.164)$$

$$Q^{(k)} = n^{-1} (S_{11} - S_{10}S_{00}^{-1}S_{10}'), \quad (6.165)$$

and

$$R^{(k)} = n^{-1} \sum_{t=1}^n \sum_{j=1}^m \pi_j(t|n) [(y_t - M_j x_t^n)(y_t - M_j x_t^n)' + M_j P_t^n M_j']. \quad (6.166)$$

where S_{11}, S_{10}, S_{00} are given in (6.68)-(6.70). As before, at iteration k , the filters and the smoothers are calculated using the current values of the parameters, $\Theta^{(k-1)}$, and Σ_0 is held fixed. Filtering is accomplished by using (6.143)-(6.147). Smoothing is derived in a similar manner to the derivation of the filter, and one is led to the smoother given in Property 6.2 and Property 6.3, with one exception, the initial smoother covariance, (6.56), is now

$$P_{n,n-1}^n = \sum_{j=1}^m \pi_j(n|n)(I - K_{tj}M_j)\Phi P_{n-1}^{n-1}. \quad (6.167)$$

Unfortunately, the computation of $\pi_j(t|n)$ is excessively complicated, and requires integrating over mixtures of normal distributions. Shumway and Stoffer (1991) suggest approximating the smoother $\pi_j(t|n)$ by the filter $\pi_j(t|t)$, and find the approximation works well.

Example 6.18 Analysis of the Influenza Data

We use the results of this section to analyze the U.S. monthly pneumonia and influenza mortality data presented in §5.4, Figure 5.7. Letting y_t denote the mortality caused by pneumonia and influenza at month t , we model y_t in terms of a structural component model coupled with a hidden Markov process that determines whether a flu epidemic exists.

The model consists of three structural components. The first component, x_{t1} , is an AR(2) process chosen to represent the periodic (seasonal) component of the data,

$$x_{t1} = \alpha_1 x_{t-1,1} + \alpha_2 x_{t-2,1} + w_{t1}, \quad (6.168)$$

where w_{t1} is white noise, with $\text{var}(w_{t1}) = \sigma_1^2$. The second component, x_{t2} , is an AR(1) process with a nonzero constant term, which is chosen to represent the sharp rise in the data during an epidemic,

$$x_{t2} = \beta_0 + \beta_1 x_{t-1,2} + w_{t2}, \quad (6.169)$$

where w_{t2} is white noise, with $\text{var}(w_{t2}) = \sigma_2^2$. The third component, x_{t3} , is a fixed trend component given by,

$$x_{t3} = x_{t-1,3} + w_{t3}, \quad (6.170)$$

where $\text{var}(w_{t3}) = 0$. The case in which $\text{var}(w_{t3}) > 0$, which corresponds to a stochastic trend (random walk), was tried here, but the estimation became unstable, and lead to us fitting a fixed, rather than stochastic, trend. Thus, in the final model, the trend component satisfies $\nabla x_{t3} = 0$; recall in [Example 5.6](#) the data were also differenced once before fitting the model.

Throughout the years, periods of normal influenza mortality (state 1) are modeled as

$$y_t = x_{t1} + x_{t3} + v_t, \quad (6.171)$$

where the measurement error, v_t , is white noise with $\text{var}(v_t) = \sigma_v^2$. When an epidemic occurs (state 2), mortality is modeled as

$$y_t = x_{t1} + x_{t2} + x_{t3} + v_t. \quad (6.172)$$

The model specified in (6.168)–(6.172) can be written in the general state-space form. The state equation is

$$\begin{pmatrix} x_{t1} \\ x_{t-1,1} \\ x_{t2} \\ x_{t3} \end{pmatrix} = \begin{bmatrix} \alpha_1 & \alpha_2 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & \beta_1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{pmatrix} x_{t-1,1} \\ x_{t-2,1} \\ x_{t-1,2} \\ x_{t-1,3} \end{pmatrix} + \begin{pmatrix} 0 \\ 0 \\ \beta_0 \\ 0 \end{pmatrix} + \begin{pmatrix} w_{t1} \\ 0 \\ w_{t2} \\ 0 \end{pmatrix}. \quad (6.173)$$

Of course, (6.173) can be written in the standard state-equation form as

$$x_t = \Phi x_{t-1} + \Upsilon u_t + w_t, \quad (6.174)$$

where $x_t = (x_{t1}, x_{t-1,1}, x_{t2}, x_{t3})'$, $\Upsilon = (0, 0, \beta_0, 0)'$, $u_t \equiv 1$, and Φ is a 4×4 matrix with σ_1^2 as the (1,1)-element, σ_2^2 as the (3,3)-element, and the remaining elements set equal to zero. The observation equation is

$$y_t = A_t x_t + v_t, \quad (6.175)$$

where A_t is 1×4 , and v_t is white noise with $\text{var}(v_t) = R = \sigma_v^2$. We assume all components of variance w_{t1} , w_{t2} , and v_t are uncorrelated.

As discussed in (6.171) and (6.172), A_t can take one of two possible forms

$$\begin{aligned} A_t = M_1 &= [1, 0, 0, 1] && \text{no epidemic,} \\ A_t = M_2 &= [1, 0, 1, 1] && \text{epidemic,} \end{aligned}$$

corresponding to the two possible states of (1) no flu epidemic and (2) flu epidemic, such that $\Pr(A_t = M_1) = 1 - \Pr(A_t = M_2)$. In this example, we will assume A_t is a hidden Markov chain, and hence we use the updating equations given in [Example 6.17](#), (6.152) and (6.153), with transition probabilities $\pi_{11} = \pi_{22} = .75$ (and, thus, $\pi_{12} = \pi_{21} = .25$).

Parameter estimation was accomplished using a quasi-Newton–Raphson procedure to maximize the approximate log likelihood given in (6.158), with initial

Table 6.3. Estimation Results for Influenza Data

Parameter	Initial Model Estimates	Final Model Estimates
α_1	1.422 (.100)	1.406 (.079)
α_2	-.634 (.089)	-.622 (.069)
β_0	.276 (.056)	.210 (.025)
β_1	-.312 (.218)	—
σ_1	.023 (.003)	.023 (.005)
σ_2	.108 (.017)	.112 (.017)
σ_v	.002 (.009)	—

Estimated standard errors in parentheses

values of $\pi_1(1|0) = \pi_2(1|0) = .5$. **Table 6.3** shows the results of the estimation procedure. On the initial fit, two estimates are not significant, namely, $\hat{\beta}_1$ and $\hat{\sigma}_v$. When $\sigma_v^2 = 0$, there is no measurement error, and the variability in data is explained solely by the variance components of the state system, namely, σ_1^2 and σ_2^2 . The case in which $\beta_1 = 0$ corresponds to a simple level shift during a flu epidemic. In the final model, with β_1 and σ_v^2 removed, the estimated level shift ($\hat{\beta}_0$) corresponds to an increase in mortality by about .2 per 1000 during a flu epidemic. The estimates for the final model are also listed in **Table 6.3**.

Figure 6.11(a) shows a plot of the data, y_t , for the ten-year period of 1969-1978 as well as an indicator that takes the value of 1 if the estimated approximate conditional probability exceeds .5, i.e., $\hat{\pi}_2(t|t-1) > .5$. The estimated prediction probabilities do a reasonable job of predicting a flu epidemic, although the peak in 1972 is missed.

Figure 6.11(b) shows the estimated filtered values (that is, filtering is done using the parameter estimates) of the three components of the model, x_{t1}^t , x_{t2}^t , and x_{t3}^t . Except for initial instability (which is not shown), $\hat{x}_{t1}^t$ represents the seasonal (cyclic) aspect of the data, $\hat{x}_{t2}^t$ represents the spikes during a flu epidemic, and $\hat{x}_{t3}^t$ represents the slow decline in flu mortality over the ten-year period of 1969-1978.

One-month-ahead prediction, say, $\hat{y}_t^{t-1}$, is obtained as

$$\hat{y}_t^{t-1} = M_1 \hat{x}_{t1}^{t-1} \quad \text{if} \quad \hat{\pi}_1(t|t-1) > \hat{\pi}_2(t|t-1),$$

$$\hat{y}_t^{t-1} = M_2 \hat{x}_{t1}^{t-1} \quad \text{if} \quad \hat{\pi}_1(t|t-1) \leq \hat{\pi}_2(t|t-1).$$

Of course, $\hat{x}_{t1}^{t-1}$ is the estimated state prediction, obtained via the filter presented in (6.143)-(6.147) (with the addition of the constant term in the model) using the estimated parameters. The results are shown in **Figure 6.11(c)**. The precision of the forecasts can be measured by the innovation variances, Σ_{t1} when no epidemic is predicted, and Σ_{t2} when an epidemic is predicted. These values become stable quickly, and when no epidemic is predicted, the estimated standard error of the prediction is approximately .02 (this is the square root of Σ_{t1} for t large); when a flu epidemic is predicted, the estimated standard error of the prediction is approximately .11.




Fig. 6.11. (a) Influenza data, y_t , (line–points) and a prediction indicator (0 or 1) that an epidemic occurs in month t given the data up to month $t - 1$ (dashed line). (b) The three filtered structural components of influenza mortality: $\hat{x}_{t1}^t$ (cyclic trace), $\hat{x}_{t2}^t$ (spiked trace), and $\hat{x}_{t3}^t$ (negative linear trace). (c) One-month-ahead predictions shown as upper and lower limits $\hat{y}_t^{t-1} \pm 2\sqrt{\hat{P}_t^{t-1}}$ (dashed lines), of the number of pneumonia and influenza deaths, and y_t (points).

The results of this analysis are impressive given the small number of parameters and the degree of approximation that was made to obtain a computationally simple method for fitting a complex model. In particular, **Figure 6.11(a)** shows that the model is never fooled as to when a flu epidemic will occur. This result is particularly impressive, given that, for example, in 1971, it appeared as though an epidemic was about to begin, but it never was realized, and the model predicted no flu epidemic that year. As seen in **Figure 6.11(c)**, the predicted mortality tends to be underestimated during the peaks, but the true values are typically within one standard error of the predicted value. Further evidence of the strength of this technique can be found in the example given in Shumway and Stoffer (1991).

The R code for the final model estimation is as follows.

```
y = as.matrix(flu); num = length(y); nstate = 4;
```

```

M1 = as.matrix(cbind(1,0,0,1)) # obs matrix normal
M2 = as.matrix(cbind(1,0,1,1)) # obs matrix flu epi
prob = matrix(0,num,1); yp = y # to store pi2(t/t-1) & y(t/t-1)
xfilter = array(0, dim=c(nstate,1,num)) # to store x(t/t)
# Function to Calculate Likelihood
Linn = function(para){
  alpha1=para[1]; alpha2=para[2]; beta0=para[3]
  sQ1=para[4]; sQ2=para[5]; like=0
  xf=matrix(0, nstate, 1) # x filter
  xp=matrix(0, nstate, 1) # x pred
  Pf=diag(.1, nstate) # filter cov
  Pp=diag(.1, nstate) # pred cov
  pi11 <- .75 -> pi22; pi12 <- .25 -> pi21; pif1 <- .5 -> pif2
  phi=matrix(0,nstate,nstate)
  phi[1,1]=alpha1; phi[1,2]=alpha2; phi[2,1]=1; phi[4,4]=1
  Ups = as.matrix(rbind(0,0,beta0,0))
  Q = matrix(0,nstate,nstate)
  Q[1,1]=sQ1^2; Q[3,3]=sQ2^2; R=0 # R=0 in final model
  # begin filtering #
  for(i in 1:num){
    xp = phi%*%xf + Ups; Pp = phi%*%Pf%*%t(phi) + Q
    sig1 = as.numeric(M1%*%Pp%*%t(M1) + R)
    sig2 = as.numeric(M2%*%Pp%*%t(M2) + R)
    k1 = Pp%*%t(M1)/sig1; k2 = Pp%*%t(M2)/sig2
    e1 = y[i]-M1%*%xp; e2 = y[i]-M2%*%xp
    pip1 = pif1*pi11 + pif2*pi21; pip2 = pif1*pi12 + pif2*pi22;
    den1 = (1/sqrt(sig1))*exp(-.5*e1^2/sig1);
    den2 = (1/sqrt(sig2))*exp(-.5*e2^2/sig2);
    denom = pip1*den1 + pip2*den2;
    pif1 = pip1*den1/denom; pif2 = pip2*den2/denom;
    pif1=as.numeric(pif1); pif2=as.numeric(pif2)
    e1=as.numeric(e1); e2=as.numeric(e2)
    xf = xp + pif1*k1*e1 + pif2*k2*e2
    eye = diag(1, nstate)
    Pf = pif1*(eye-k1%*%M1)%*%Pp + pif2*(eye-k2%*%M2)%*%Pp
    like = like - log(pip1*den1 + pip2*den2)
    prob[i]<-pip2; xfilter[,i]<-xf; innov.sig<-c(sig1,sig2)
    yp[i]<-ifelse(pip1 > pip2, M1%*%xp, M2%*%xp) }
  return(like) }
# Estimation
alpha1=1.4; alpha2=-.5; beta0=.3; sQ1=.1; sQ2=.1
init.par = c(alpha1, alpha2, beta0, sQ1, sQ2)
(est = optim(init.par, Linn, NULL, method="BFGS", hessian=TRUE,
  control=list(trace=1,REPORT=1)))
SE = sqrt(diag(solve(est$hessian)))
u = cbind(estimate=est$par, SE)
rownames(u)=c("alpha1","alpha2","beta0","sQ1","sQ2"); u
      estimate      SE
alpha1  1.40570967 0.078587727
alpha2 -0.62198715 0.068733109
beta0   0.21049042 0.024625302
sQ1     0.02310306 0.001635291
sQ2     0.11217287 0.016684663
# Graphics
predepi = ifelse(prob<.5,0,1); k = 6:length(y)
Time = time(flu)[k]

```

```

par(mfrow=c(3,1), mar=c(2,3,1,1)+.1, cex=.9)
plot(Time, y[k], type="o", ylim=c(0,1),ylab="")
  lines(Time, predepi[k], lty="dashed", lwd=1.2)
  text(1979,.95,"(a)")
plot(Time, xfilter[1,,k], type="l", ylim=c(-.1,.4), ylab="")
  lines(Time, xfilter[3,,k]); lines(Time, xfilter[4,,k])
  text(1979,.35,"(b)")
plot(Time, y[k], type="p", pch=1, ylim=c(.1,.9),ylab="")
  prde1 = 2*sqrt(innov.sig[1]); prde2 = 2*sqrt(innov.sig[2])
  prde = ifelse(predepi[k]<.5, prde1,prde2)
  lines(Time, yp[k]+prde, lty=2, lwd=1.5)
  lines(Time, yp[k]-prde, lty=2, lwd=1.5)
  text(1979,.85,"(c)")

```

6.10 Stochastic Volatility

Recently, there has been considerable interest in stochastic volatility models. These models are similar to the ARCH models presented in Chapter 5, but they add a stochastic noise term to the equation for σ_t . Recall from §5.4 that a GARCH(1, 1) model for a return, which we denote here by r_t , is given by

$$r_t = \sigma_t \epsilon_t \quad (6.176)$$

$$\sigma_t^2 = \alpha_0 + \alpha_1 r_{t-1}^2 + \beta_1 \sigma_{t-1}^2, \quad (6.177)$$

where ϵ_t is Gaussian white noise. If we define

$$h_t = \log \sigma_t^2 \quad \text{and} \quad y_t = \log r_t^2,$$

then (6.176) can be written as

$$y_t = h_t + \log \epsilon_t^2. \quad (6.178)$$

Equation (6.178) is considered the observation equation, and the stochastic variance h_t is considered to be an unobserved state process. Instead of (6.177), however, the model assumes the volatility process follows, in its basic form, an autoregression,

$$h_t = \phi_0 + \phi_1 h_{t-1} + w_t, \quad (6.179)$$

where w_t is white Gaussian noise with variance σ_w^2 .

Together, (6.178) and (6.179) make up the stochastic volatility model due to Taylor (1982). If ϵ_t^2 had a log-normal distribution, (6.178)–(6.179) would form a Gaussian state-space model, and we could then use standard DLM results to fit the model to data. Unfortunately, $y_t = \log r_t^2$ is rarely normal, so we typically keep the ARCH normality assumption on ϵ_t ; in which case, $\log \epsilon_t^2$ is distributed as the log of a chi-squared random variable with one degree of freedom. This density is given by

$$f(x) = \frac{1}{\sqrt{2\pi}} \exp \left\{ -\frac{1}{2} (e^x - x) \right\} \quad -\infty < x < \infty, \quad (6.180)$$

and its mean and variance are -1.27 and $\pi^2/2$, respectively; the density (6.180) is highly skewed with a long tail on the left (see Figure 6.13).

Various approaches to the fitting of stochastic volatility models have been examined; these methods include a wide range of assumptions on the observational noise process. A good summary of the proposed techniques, both Bayesian (via MCMC) and non-Bayesian approaches (such as quasi-maximum likelihood estimation and the EM algorithm), can be found in Jacquier et al. (1994), and Shephard (1996). Simulation methods for classical inference applied to stochastic volatility models are discussed in Danielson (1994) and Sandmann and Koopman (1998).

Kim, Shephard and Chib (1998) proposed modeling the log of a chi-squared random variable by a mixture of seven normals to approximate the first four moments of the observational error distribution; the mixture is fixed and no additional model parameters are added by using this technique. The basic model assumption that ϵ_t is Gaussian is unrealistic for most applications. In an effort to keep matters simple but more general (in that we allow the observational error dynamics to depend on parameters that will be fitted), our method of fitting stochastic volatility models is to retain the Gaussian state equation (6.179), but to write the observation equation, with $y_t = \log r_t^2$, as

$$y_t = \alpha + h_t + \eta_t, \quad (6.181)$$

where η_t is white noise, whose distribution is a mixture of two normals, one centered at zero. In particular, we write

$$\eta_t = I_t z_{t0} + (1 - I_t) z_{t1}, \quad (6.182)$$

where I_t is an iid Bernoulli process, $\Pr\{I_t = 0\} = \pi_0$, $\Pr\{I_t = 1\} = \pi_1$ ($\pi_0 + \pi_1 = 1$), $z_{t0} \sim \text{iid } N(0, \sigma_0^2)$, and $z_{t1} \sim \text{iid } N(\mu_1, \sigma_1^2)$.

The advantage to this model is that it is easy to fit because it uses normality. In fact, the model equations (6.179) and (6.181)-(6.182) are similar to those presented in Peña and Guttman (1988), who used the idea to obtain a robust Kalman filter, and, as previously mentioned, in Kim et al. (1998). The material presented in §6.9 applies here, and in particular, the filtering equations for this model are

$$h_{t+1}^t = \phi_0 + \phi_1 h_t^{t-1} + \sum_{j=0}^1 \pi_{tj} K_{tj} \epsilon_{tj}, \quad (6.183)$$

$$P_{t+1}^t = \phi_1^2 P_t^{t-1} + \sigma_w^2 - \sum_{j=0}^1 \pi_{tj} K_{tj}^2 \Sigma_{tj}, \quad (6.184)$$

$$\epsilon_{t0} = y_t - \alpha - h_t^{t-1}, \quad \epsilon_{t1} = y_t - \alpha - h_t^{t-1} - \mu_1, \quad (6.185)$$

$$\Sigma_{t0} = P_t^{t-1} + \sigma_0^2, \quad \Sigma_{t1} = P_t^{t-1} + \sigma_1^2, \quad (6.186)$$

$$K_{t0} = \phi_1 P_t^{t-1} / \Sigma_{t0}, \quad K_{t1} = \phi_1 P_t^{t-1} / \Sigma_{t1}. \quad (6.187)$$

To complete the filtering, we must be able to assess the probabilities $\pi_{t1} = \Pr(I_t = 1 \mid y_1, \dots, y_t)$, for $t = 1, \dots, n$; of course, $\pi_{t0} = 1 - \pi_{t1}$. Let $f_j(t \mid t-1)$ denote the conditional density of y_t given the past $y_1, \dots, y_{t-1}$, and $I_t = j$ ($j = 0, 1$). Then,

$$\pi_{t1} = \frac{\pi_1 f_1(t | t-1)}{\pi_0 f_0(t | t-1) + \pi_1 f_1(t | t-1)}, \quad (6.188)$$

where we assume the distribution π_j , for $j = 0, 1$ has been specified *a priori*. If the investigator has no reason to prefer one state over another the choice of uniform priors, $\pi_1 = 1/2$, will suffice. Unfortunately, it is computationally difficult to obtain the exact values of $f_j(t | t-1)$; although we can give an explicit expression of $f_j(t | t-1)$, the actual computation of the conditional density is prohibitive. A viable approximation, however, is to choose $f_j(t | t-1)$ to be the normal density, $N(h_t^{t-1} + \mu_j, \Sigma_{tj})$, for $j = 0, 1$ and $\mu_0 = 0$; see §6.9 for details.

The innovations filter given in (6.183)–(6.188) can be derived from the Kalman filter by a simple conditioning argument; e.g., to derive (6.183), write

$$\begin{aligned} E(h_{t+1} | y_1, \dots, y_t) &= \sum_{j=0}^1 E(h_{t+1} | y_1, \dots, y_t, I_t = j) \Pr(I_t = j | y_1, \dots, y_t) \\ &= \sum_{j=0}^1 (\phi_0 + \phi_1 h_t^{t-1} + K_{tj} \epsilon_{tj}) \pi_{tj} \\ &= \phi_0 + \phi_1 h_t^{t-1} + \sum_{j=0}^1 \pi_{tj} K_{tj} \epsilon_{tj}. \end{aligned}$$

Estimation of the parameters, $\Theta = (\phi_0, \phi_1, \sigma_0^2, \mu_1, \sigma_1^2, \sigma_w^2)'$, is accomplished via MLE based on the likelihood given by

$$\ln L_Y(\Theta) = \sum_{t=1}^n \ln \left(\sum_{j=0}^1 \pi_j f_j(t | t-1) \right), \quad (6.189)$$

where the density $f_j(t | t-1)$ is approximated by the normal density, $N(h_t^{t-1} + \mu_j, \sigma_j^2)$, previously mentioned. We may consider maximizing (6.189) directly as a function of the parameters Θ using a Newton method, or we may consider applying the EM algorithm to the complete data likelihood.

Example 6.19 Analysis of the New York Stock Exchange Returns

The top of Figure 6.12 shows the log of the squares of returns, $y_t = \log r_t^2$, for 200 of the 2000 daily observations of the NYSE previously displayed in Figure 1.4. Model (6.179) and (6.181)–(6.182), with π_1 fixed at .5, was fit to the data using a quasi-Newton–Raphson method to maximize (6.189). The results are given in Table 6.4. Figure 6.13 compares the density of the log of a χ_1^2 with the fitted normal mixture; we note the data indicate a substantial amount of probability in the upper tail that the log- χ_1^2 distribution misses.

Finally, the bottom of Figure 6.12 shows y_t for $800 \leq t \leq 1000$, which includes the crash of October 19, 1987, with $y_t^{t-1} = \hat{\alpha} + h_t^{t-1}$ superimposed on the graph; compare with Figure 5.6. Also displayed are error bounds.

The R code when ϕ_0 is included in the model is as follows.




Fig. 6.12. Two hundred observations of $y_t = \log r_t^2$, for $801 \leq t \leq 1000$, where r_t is the daily return of the NYSE (top); the crash of October 19, 1987 occurs at $t = 938$. Corresponding one-step-ahead predicted log volatility, $\log \sigma_t^2$, with ± 2 standard prediction errors (bottom).


```

y = log(nyse^2)
num = length(y)
# Initial Parameters
phi0=0; phi1=.95; sQ=.2; alpha=mean(y); sR0=1; mu1=-3; sR1=2
init.par = c(phi0, phi1, sQ, alpha, sR0, mu1, sR1)
# Innovations Likelihood
Linn = function(para){
  phi0=para[1]; phi1=para[2]; sQ=para[3]; alpha=para[4]
  sR0=para[5]; mu1=para[6]; sR1=para[7]
  sv = SVfilter(num,y,phi0,phi1,sQ,alpha,sR0,mu1,sR1)
  return(sv$like) }
# Estimation
(est = optim(init.par, Linn, NULL, method="BFGS", hessian=TRUE,
  control=list(trace=1,REPORT=1)))
SE = sqrt(diag(solve(est$hessian)))
u = cbind(estimated=est$par, SE)
rownames(u)=c("phi0","phi1","sQ","alpha","sigv0","mu1","sigv1"); u
# Graphics (need filters at the estimated parameters)
phi0=est$par[1]; phi1=est$par[2]; sQ=est$par[3]; alpha=est$par[4]
sR0=est$par[5]; mu1=est$par[6]; sR1=est$par[7]
sv = SVfilter(num,y,phi0,phi1,sQ,alpha,sR0,mu1,sR1)
# densities plot (f is chi-sq, fm is fitted mixture)
x = seq(-15,6,by=.01)
f = exp(-.5*(exp(x)-x))/(sqrt(2*pi))
f0 = exp(-.5*(x^2)/sR0^2)/(sR0*sqrt(2*pi))
f1 = exp(-.5*(x-mu1)^2/sR1^2)/(sR1*sqrt(2*pi))
fm = (f0+f1)/2
plot(x, f, type="l"); lines(x, fm, lty=2,lwd=2)
dev.new(); par(mfrow=c(2,1)); Time=801:1000

```

Table 6.4. Estimation Results for the NYSE Fit

Parameter	Estimate	Estimated
		Standard Error
ϕ_0	-.006	.016
ϕ_1	.988	.007
σ_w	.091	.027
α	-9.613	1.269
σ_0	1.220	.065
μ_1	-2.292	.205
σ_1	2.683	.105

**Fig. 6.13.** Density of the log of a χ_1^2 as given by (6.180) (solid line) and the fitted normal mixture (dashed line) from the NYSE example.

```

plot(Time, y[Time], type="l", main="log(Squared NYSE Returns)")
plot(Time, sv$xp[Time], type="l", main="Predicted log-Volatility",
      ylim=c(-1.5, 1.8), ylab="", xlab="")
lines(Time, sv$xp[Time]+2*sqrt(sv$Pp[Time]), lty="dashed")
lines(Time, sv$xp[Time]-2*sqrt(sv$Pp[Time]), lty="dashed")

```

It is possible to use the bootstrap procedure described in §6.8 for the stochastic volatility model, with some minor changes. The following procedure was described in Stoffer and Wall (2004). We develop a vector first-order equation, as was done in (6.130). First, using (6.185), and noting that $y_t = \pi_{t0}y_t + \pi_{t1}y_t$, we may write

$$y_t = \alpha + h_t^{t-1} + \pi_{t0}\epsilon_{t0} + \pi_{t1}(\epsilon_{t1} + \mu_1). \quad (6.190)$$

Consider the standardized innovations

$$e_{tj} = \Sigma_{tj}^{-1/2} \epsilon_{tj}, \quad j = 0, 1, \quad (6.191)$$

and define the 2×1 vector

$$e_t = \begin{bmatrix} e_{t0} \\ e_{t1} \end{bmatrix}.$$




Fig. 6.14. Log of the squared residuals from an MA(2) fit on GNP growth rate.

Also, define the 2×1 vector

$$\xi_t = \begin{bmatrix} h_{t+1}^t \\ y_t \end{bmatrix}.$$

Combining (6.183) and (6.190) results in a vector first-order equation for ξ_t given by

$$\xi_t = F\xi_{t-1} + G_t + H_t e_t, \quad (6.192)$$

where

$$F = \begin{bmatrix} \phi_1 & 0 \\ 1 & 0 \end{bmatrix}, \quad G_t = \begin{bmatrix} \phi_0 \\ \alpha + \pi_{t1}\mu_1 \end{bmatrix}, \quad H_t = \begin{bmatrix} \pi_{t0}K_{t0}\Sigma_{t0}^{1/2} & \pi_{t1}K_{t1}\Sigma_{t1}^{1/2} \\ \pi_{t0}\Sigma_{t0}^{1/2} & \pi_{t1}\Sigma_{t1}^{1/2} \end{bmatrix}.$$

Hence, the steps in bootstrapping for this case are the same as steps 1 through 5 described in §5.7, but with (6.130) replaced by the following first-order equation:

$$\xi_t^* = F(\hat{\Theta})\xi_{t-1}^* + G_t(\hat{\Theta}; \hat{\pi}_{t1}) + H_t(\hat{\Theta}; \hat{\pi}_{t1})e_t^*, \quad (6.193)$$

where $\hat{\Theta} = (\hat{\phi}_0, \hat{\phi}_1, \hat{\sigma}_0^2, \hat{\alpha}, \hat{\mu}_1, \hat{\sigma}_1^2, \hat{\sigma}_w^2)'$ is the MLE of Θ , and $\hat{\pi}_{t1}$ is estimated via (6.188), replacing $f_1(t | t-1)$ and $f_0(t | t-1)$ by their respective estimated normal densities ($\hat{\pi}_{t0} = 1 - \hat{\pi}_{t1}$).

Example 6.20 Analysis of the U.S. GNP Growth Rate

In Example 5.4, we fit an ARCH model to the U.S. GNP growth rate. In this example, we will fit a stochastic volatility model to the residuals from the MA(2) fit on the growth rate (see Example 3.38). Figure 6.14 shows the log of the squared residuals, say y_t , from the MA(2) fit on the U.S. GNP series. The stochastic volatility model (6.178)–(6.182) was then fit to y_t . 6.5 shows the MLEs of the model parameters along with their asymptotic SEs assuming the model is correct. Also displayed in Table 6.5 are the means and SEs of $B = 500$ bootstrapped samples. There is some amount of agreement between the asymptotic values and the bootstrapped values. The interest here, however, is not so much in the SEs, but in the actual




Fig. 6.15. Bootstrap histogram and asymptotic distribution of $\hat{\phi}_1$ for the U.S. GNP example.

sampling distribution of the estimates. For example, [Figure 6.15](#) compares the bootstrap histogram and asymptotic normal distribution of $\hat{\phi}_1$. In this case, the bootstrap distribution exhibits positive kurtosis and skewness which is missed by the assumption of asymptotic normality.

The R code for this example is as follows. We held ϕ_0 at 0 for this analysis because it was not significantly different from 0 in an initial analysis.

```
n.boot = 500 # number of bootstrap replicates
tol = sqrt(Machine$double.eps) # convergence tolerance
gnpgr = diff(log(gnp))
fit = arima(gnpgr, order=c(1,0,0))
y = as.matrix(log(resid(fit)^2))
num = length(y)
plot.ts(y, ylab="")
# Initial Parameters
phi1 = .9; sQ = .5; alpha = mean(y); sR0 = 1; mu1 = -3; sR1 = 2.5
init.par = c(phi1, sQ, alpha, sR0, mu1, sR1)
# Innovations Likelihood
Linn=function(para){
  phi1 = para[1]; sQ = para[2]; alpha = para[3]
  sR0 = para[4]; mu1 = para[5]; sR1 = para[6]
  sv = SVfilter(num, y, 0, phi1, sQ, alpha, sR0, mu1, sR1)
  return(sv$like) }
# Estimation
(est = optim(init.par, Linn, NULL, method="BFGS", hessian=TRUE,
  control=list(trace=1,REPORT=1)))
SE = sqrt(diag(solve(est$hessian)))
u = rbind(estimated=est$par, SE)
colnames(u)=c("phi1","sQ","alpha","sig0","mu1","sig1"); u
      phi1    sQ  alpha  sig0    mu1  sig1
estimated 0.8790 0.3878 -9.6624 0.8325 -2.3412 2.4516
SE         0.1061 0.2172 0.3386 0.2034 0.4952 0.2927

# Bootstrap
para.star = matrix(0, n.boot, 6) # to store parameter estimates
```

Table 6.5. Estimates and Standard Errors for GNP Example

Parameter	MLE	Asymptotic	Bootstrap [†]
		SE	SE
ϕ_1	.879	.106	.074
σ_w	.388	.217	.428
α	-9.662	.339	1.536
σ_0	.833	.203	.389
μ_1	-2.341	.495	.437
σ_1	2.452	.293	.330

[†] Based on 500 bootstrapped samples.

```

Linn2 = function(para){  # calculate likelihood
  phi1 = para[1]; sQ = para[2]; alpha = para[3]
  sR0 = para[4]; mu1 = para[5]; sR1 = para[6]
  sv = SVfilter(num, y.star, 0, phi1, sQ, alpha, sR0, mu1, sR1)
  return(sv$like) }
for (jb in 1:n.boot){
  cat("iteration:", jb, "\n")
  phi1 = est$par[1]; sQ = est$par[2]; alpha = est$par[3]
  sR0 = est$par[4]; mu1 = est$par[5]; sR1 = est$par[6]
  Q = sQ^2; R0 = sR0^2; R1 = sR1^2
  sv = SVfilter(num, y, 0, phi1, sQ, alpha, sR0, mu1, sR1)
  sig0 = sv$Pp+R0; sig1 = sv$Pp+R1;
  K0 = sv$Pp/sig0; K1 = sv$Pp/sig1
  inn0 = y-sv$xp-alpha; inn1 = y-sv$xp-mu1-alpha
  den1 = (1/sqrt(sig1))*exp(-.5*inn1^2/sig1)
  den0 = (1/sqrt(sig0))*exp(-.5*inn0^2/sig0)
  fpi1 = den1/(den0+den1)
  # start resampling at t=4
  e0 = inn0/sqrt(sig0); e1 = inn1/sqrt(sig1)
  indx = sample(4:num, replace=TRUE)
  sinn = cbind(c(e0[1:3], e0[indx]), c(e1[1:3], e1[indx]))
  eF = matrix(c(phi1, 1, 0, 0), 2, 2)
  xi = cbind(sv$xp,y) # initialize
  for (i in 4:num){ # generate boot sample
    G = matrix(c(0, alpha+fpi1[i]*mu1), 2, 1)
    h21 = (1-fpi1[i])*sqrt(sig0[i]); h11 = h21*K0[i]
    h22 = fpi1[i]*sqrt(sig1[i]); h12 = h22*K1[i]
    H = matrix(c(h11,h21,h12,h22),2,2)
    xi[i,] = t(eF%*%as.matrix(xi[i-1,],2) + G + H%*%as.matrix(sinn[i,],2))}
# Estimates from boot data
y.star = xi[,2]
phi1=.9; sQ=.5; alpha=mean(y.star); sR0=1; mu1=-3; sR1=2.5
init.par = c(phi1, sQ, alpha, sR0, mu1, sR1) # same as for data
est.star = optim(init.par, Linn2, NULL, method="BFGS",
  control=list(reltol=tol))
para.star[jb,] = cbind(est.star$par[1], abs(est.star$par[2]),
  est.star$par[3], abs(est.star$par[4]), est.star$par[5],
  abs(est.star$par[6])) }
# Some summary statistics and graphics
rmse = rep(NA,6) # SEs from the bootstrap
for(i in 1:6){

```

```

rmse[i] = sqrt(sum((para.star[,i]-est$par[i])^2)/n.boot)
cat(i, rmse[i], "\n") }
dev.new(); phi = para.star[,1]
hist(phi, 15, prob=TRUE, main="", xlim=c(.4,1.2), xlab="")
u = seq(.4, 1.2, by=.01)
lines(u,dnorm(u, mean=.8790267, sd=.1061884), lty="dashed", lwd=2)

```

6.11 Bayesian Analysis of State Space Models

We now consider some Bayesian approaches to fitting linear Gaussian state space models via Markov chain Monte Carlo (MCMC) methods. We assume that the model is given by (6.1)–(6.2); inputs are allowed in the model, but we do not display them for the sake of brevity. In this case, Frühwirth-Schnatter (1994) and Carter and Kohn (1994) established the MCMC procedure that we will discuss here. A comprehensive text that we highly recommend for this case is Petris et al. (2009) and the corresponding R package `dlm`. For nonlinear and non-Gaussian models, the reader is referred to Douc, Moulines, & Stoffer (2014). Throughout this section, we have n observations denoted by $\mathcal{Y}_n = \{y_1, \dots, y_n\}$, whereas the states are denoted as $\mathcal{X}_n = \{x_0, x_1, \dots, x_n\}$, with x_0 being the initial state.

MCMC methods refer to Monte Carlo integration methods that use a Markovian updating scheme to sample from intractable posterior distributions. The most common MCMC method is the Gibbs sampler, which is essentially a modification of the Metropolis algorithm (Metropolis et al., 1953) developed by Hastings (1970) in the statistical setting and by Geman and Geman (1984) in the context of image restoration. Later, Tanner and Wong (1987) used the ideas in their substitution sampling approach, and Gelfand and Smith (1990) developed the Gibbs sampler for a wide class of parametric models. The basic strategy is to use conditional distributions to set up a Markov chain to obtain samples from a joint distribution. The following simple case demonstrates this idea.

Example 6.21 Gibbs Sampling for the Bivariate Normal

Suppose we wish to obtain samples from a bivariate normal distribution,

$$\begin{pmatrix} X \\ Y \end{pmatrix} \sim N \left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix} \right],$$

where $|\rho| < 1$, but we can only generate samples from a univariate normal.

- The univariate conditionals are

$$(X | Y = y) \sim N(\rho y, 1 - \rho^2) \quad \text{and} \quad (Y | X = x) \sim N(\rho x, 1 - \rho^2),$$

and we can simulate from these distributions.

- Construct a Markov chain: Pick $X^{(0)} = x_0$, and then iterate the process $X^{(0)} = x_0 \mapsto Y^{(0)} \mapsto X^{(1)} \mapsto Y^{(1)} \mapsto \dots \mapsto X^{(k)} \mapsto Y^{(k)} \mapsto \dots$, where

$$\begin{aligned} (Y^{(k)} | X^{(k)} = x_k) &\sim N(\rho x_k, 1 - \rho^2) \\ (X^{(k)} | Y^{(k-1)} = y_{k-1}) &\sim N(\rho y_{k-1}, 1 - \rho^2). \end{aligned}$$

- The joint distribution of $(X^{(k)}, Y^{(k)})$ is

$$\begin{pmatrix} X^{(k)} \\ Y^{(k)} \end{pmatrix} \sim N \left[\begin{pmatrix} \rho^{2k} x_0 \\ \rho^{2k+1} x_0 \end{pmatrix}, \begin{pmatrix} 1 - \rho^{4k} & \rho(1 - \rho^{4k}) \\ \rho(1 - \rho^{4k}) & 1 - \rho^{4k+2} \end{pmatrix} \right].$$

- Thus, for any starting value, x_0 , $(X^{(k)}, Y^{(k)}) \rightarrow_d (X, Y)$ as $k \rightarrow \infty$; the rate depends on ρ . Then one would run the chain and throw away the initial n_0 sampled values (burnin) and retain the rest.

For state space models, the main objective is to obtain the posterior density of the parameters $p(\Theta | \mathcal{Y}_n)$ or $p(\mathcal{X}_n | \mathcal{Y}_n)$ if the states are meaningful. For example, the states do not have any meaning for an ARMA model, but they are important for a stochastic volatility model. It is generally easier to get samples from the full posterior $p(\Theta, \mathcal{X}_n | \mathcal{Y}_n)$ and then marginalize (“average”) to obtain $p(\Theta | \mathcal{Y}_n)$ or $p(\mathcal{X}_n | \mathcal{Y}_n)$. As previously mentioned, the most popular method is to run a full Gibbs sampler, alternating between sampling model parameters and latent state sequences from their respective full conditional distributions.

Procedure 6.1 Gibbs Sampler for State Space Models

- (i) Draw $\Theta' \sim p(\Theta | \mathcal{X}_n, \mathcal{Y}_n)$
- (ii) Draw $\mathcal{X}'_n \sim p(\mathcal{X}_n | \Theta', \mathcal{Y}_n)$

Procedure 6.1-(i) is generally much easier because it conditions on the complete data $\{\mathcal{X}_n, \mathcal{Y}_n\}$, which we saw in §6.3 can simplify the problem. **Procedure 6.1-(ii)** amounts to sampling from the joint smoothing distribution of the latent state sequence and is generally difficult. For linear Gaussian models, however, both parts of **Procedure 6.1** are relatively easy to perform.

To accomplish **Procedure 6.1-(i)**, note that

$$p(\Theta | \mathcal{X}_n, \mathcal{Y}_n) \propto \pi(\Theta) p(x_0 | \Theta) \prod_{t=1}^n p(x_t | x_{t-1}, \Theta) p(y_t | x_t, \Theta) \quad (6.194)$$

where $\pi(\Theta)$ is the prior on the parameters. The prior often depends on “hyperparameters” that add another level to the hierarchy. For simplicity, these hyperparameters are assumed to be known. The parameters are typically conditionally independent with distributions from standard parametric families (at least as long as the prior distribution is conjugate relative to the Bayesian model specification). For non-conjugate models, one option is to replace **Procedure 6.1-(i)** with a Metropolis-Hastings step, which is feasible since the complete data density $p(\Theta, \mathcal{X}_n, \mathcal{Y}_n)$ can be evaluated pointwise.

For example, in the univariate model

$$x_t = \phi x_{t-1} + w_t \quad \text{and} \quad y_t = x_t + v_t$$

where $w_t \sim \text{iid } N(0, \sigma_w^2)$ independent of $v_t \sim \text{iid } N(0, \sigma_v^2)$, we can use the normal and inverse gamma (IG) distributions for priors. In this case, the priors on the variance

components are chosen from a conjugate family, that is, $\sigma_w^2 \sim \text{IG}(a_0/2, b_0/2)$ independent of $\sigma_v^2 \sim \text{IG}(c_0/2, d_0/2)$, where IG denotes the inverse (reciprocal) gamma distribution. Then, for example, if the prior on ϕ is Gaussian, $\phi \sim \text{N}(\mu_\phi, \sigma_\phi^2)$, then $\phi \mid \sigma_w, \mathcal{X}_n, \mathcal{Y}_n \sim \text{N}(Bb, B)$, where

$$B^{-1} = \frac{1}{\sigma_\phi^2} + \frac{1}{\sigma_w^2} \sum_{t=1}^n x_{t-1}^2, \quad b = \frac{\mu_\phi}{\sigma_\phi^2} + \frac{1}{\sigma_w^2} \sum_{t=1}^n x_t x_{t-1}.$$

and

$$\sigma_w^2 \mid \phi, \mathcal{X}_n, \mathcal{Y}_n \sim \text{IG}\left(\frac{1}{2}(a_0 + n), \frac{1}{2}\{b_0 + \sum_{t=1}^n [x_t - \phi x_{t-1}]^2\}\right);$$

$$\sigma_v^2 \mid \mathcal{X}_n, \mathcal{Y}_n \sim \text{IG}\left(\frac{1}{2}(c_0 + n), \frac{1}{2}\{c_0 + \sum_{t=1}^n [y_t - x_t]^2\}\right).$$

For **Procedure 6.1-(ii)**, the goal is to sample the entire set of state vectors, $\mathcal{X}_n$, from the posterior density $p(\mathcal{X}_n \mid \Theta, \mathcal{Y}_n)$, where Θ is a fixed set of parameters obtained from the previous step. We will write the posterior as $p_\Theta(\mathcal{X}_n \mid \mathcal{Y}_n)$ to save space. Because of the Markov structure, we can write,

$$p_\Theta(\mathcal{X}_n \mid \mathcal{Y}_n) = p_\Theta(x_n \mid \mathcal{Y}_n) p_\Theta(x_{n-1} \mid x_n, \mathcal{Y}_{n-1}) \cdots p_\Theta(x_0 \mid x_1). \quad (6.195)$$

In view of (6.195), it is possible to sample the entire set of state vectors, $\mathcal{X}_n$, by sequentially simulating the individual states backward. This process yields a simulation method that Frühwirth-Schnatter (1994) called the forward-filtering, backward-sampling (FFBS) algorithm. From (6.195), we see that we must obtain the densities

$$p_\Theta(x_t \mid x_{t+1}, \mathcal{Y}_t) \propto p_\Theta(x_t \mid \mathcal{Y}_t) p_\Theta(x_{t+1} \mid x_t).$$

In particular, we know that $x_t \mid \mathcal{Y}_t \sim \text{N}_p^\Theta(x_t^t, P_t^t)$ and $x_{t+1} \mid x_t \sim \text{N}_p^\Theta(\Phi x_t, Q)$. And because the processes are Gaussian, we need only obtain the conditional means and variances, say, $m_t = E_\Theta(x_t \mid \mathcal{Y}_t, x_{t+1})$ and $V_t = \text{var}_\Theta(x_t \mid \mathcal{Y}_t, x_{t+1})$. In particular,

$$m_t = x_t^t + J_t(x_{t+1} - x_{t+1}^t) \quad \text{and} \quad V_t = P_t^t - J_t P_{t+1}^t J_t', \quad (6.196)$$

for $t = n-1, n-2, \dots, 0$, where J_t is defined in (6.48). We note that m_t has already been derived in (6.49). To derive m_t and V_t using standard normal theory, use a strategy similar to the derivation of the filter in **Property 6.1**. That is,

$$\begin{pmatrix} x_t \\ x_{t+1} \end{pmatrix} \mid \mathcal{Y}_t \sim \text{N}\left(\begin{bmatrix} x_t^t \\ x_{t+1}^t \end{bmatrix}, \begin{bmatrix} P_t^t & P_t^t \Phi' \\ \Phi P_t^t & P_{t+1}^t \end{bmatrix}\right);$$

now use (B.9), (B.10), and the definition of J_t in (6.48). Also, recall the proof of **Property 6.3** wherein we noted the off-diagonal $P_{t+1,t}^t = \Phi P_t^t$.

Hence, given Θ , the algorithm is to first sample x_n from a $\text{N}_p^\Theta(x_n^n, P_n^n)$, where x_n^n and P_n^n are obtained from the Kalman filter, **Property 6.1**, and then sample x_t from a $\text{N}_p^\Theta(m_t, V_t)$, for $t = n-1, n-2, \dots, 0$, where the conditioning value of x_{t+1} is the value previously sampled.




Fig. 6.16. Display for [Example 6.22](#): Top: Generated states, x_t and data y_t . Contours of the likelihood (solid line) of the data and sampled posterior values as points. Bottom: Marginal sampled posteriors and posterior means (vertical lines) of each variance component. The true values are $\sigma_v^2 = 1$ and $\sigma_w^2 = .5$.

Example 6.22 Local Level Model

In this example, we consider the local level model previously discussed in [Example 6.4](#). Here, we consider the model

$$y_t = x_t + v_t \quad \text{and} \quad x_t = x_{t-1} + w_t$$

where $v_t \sim \text{iid } N(0, \sigma_v^2 = 1)$ independent of $w_t \sim \text{iid } N(0, \sigma_w^2 = .5)$. This is the univariate model we just discussed, but where $\phi = 1$. In this case, we used IG priors for each of the variance components.

For the prior distributions, all parameters (a_0, b_0, c_0, d_0) were set to .02. We generated 1010 samples, using the first 10 as burn-in. [Figure 6.16](#) displays the simulated data and states, the contours of the likelihood of the data, the sampled posterior values as points, and the marginal sampled posteriors of each variance component along with the posterior means. [Figure 6.17](#) compares the actual smoother x_t^n with the posterior mean of the sampled smoothed values. In addition, a pointwise 95% credible interval is displayed as a filled area.

The following code was used in this example.

```
##-- Notation --##
#           y(t) = x(t) + v(t);    v(t) ~ iid N(0,V)
#           x(t) = x(t-1) + w(t);  w(t) ~ iid N(0,W)
# priors:  x(0) ~ N(m0,C0); V ~ IG(a,b); W ~ IG(c,d)
# FFBS:    x(t|t) ~ N(m,C); x(t/n) ~ N(mm,CC); x(t/t+1) ~ N(a,R)
##--
ffbs = function(y,V,W,m0,C0){
  n = length(y); a = rep(0,n); R = rep(0,n)
```




Fig. 6.17. Display for [Example 6.22](#): True smoother, x_t^n , the data y_t , and the posterior mean of the sampled smoother values; the filled in area shows 2.5% to 97.5%-tiles of the draws.

```

m = rep(0,n); C = rep(0,n); B = rep(0,n-1)
H = rep(0,n-1); mm = rep(0,n); CC = rep(0,n)
x = rep(0,n); llike = 0.0
for (t in 1:n){
  if(t==1){
    a[1] = m0; R[1] = C0 + W
  }else{
    a[t] = m[t-1]; R[t] = C[t-1] + W }
  f = a[t]
  Q = R[t] + V
  A = R[t]/Q
  m[t] = a[t]+A*(y[t]-f)
  C[t] = R[t]-Q*A**2
  B[t-1] = C[t-1]/R[t]
  H[t-1] = C[t-1]-R[t]*B[t-1]**2
  llike = llike + dnorm(y[t],f,sqrt(Q),log=TRUE)
}
mm[n] = m[n]; CC[n] = C[n]
x[n] = rnorm(1,m[n],sqrt(C[n]))
for (t in (n-1):1){
  mm[t] = m[t] + C[t]/R[t+1]*(mm[t+1]-a[t+1])
  CC[t] = C[t] - (C[t]^2)/(R[t+1]^2)*(R[t+1]-CC[t+1])
  x[t] = rnorm(1,m[t]+B[t]*(x[t+1]-a[t+1]),sqrt(H[t])) }
return(list(x=x,m=m,C=C,mm=mm,CC=CC,llike=llike))
}
# Simulate states and data
# -----
set.seed(1); W = 0.5; V = 1.0
n = 100; m0 = 0.0; C0 = 10.0; x0 = 0
w = rnorm(n,0,sqrt(W))
v = rnorm(n,0,sqrt(V))
x = y = rep(0,n)
x[1] = x0 + w[1]

```

```

y[1] = x[1] + v[1]
for (t in 2:n){
  x[t] = x[t-1] + w[t]
  y[t] = x[t] + v[t] }
### actual smoother (for plotting)
ks = Ksmooth0(num=n, y, A=1, m0, C0, Phi=1, cQ=sqrt(W), cR=sqrt(V))
xsmooth = as.vector(ks$xs)
###
run = ffbs(y,V,W,m0,C0)
m = run$m; C = run$C; mm = run$mm
CC = run$CC; L1 = m-2*C; U1 = m+2*C
L2 = mm-2*CC; U2 = mm+2*CC
N = 50
Vs = seq(0.1,2,length=N)
Ws = seq(0.1,2,length=N)
likes = matrix(0,N,N)
for (i in 1:N){
  for (j in 1:N){
    V = Vs[i]
    W = Ws[j]
    run = ffbs(y,V,W,m0,C0)
    likes[i,j] = run$llike }
}
# Hyperparameters
a = 0.01; b = 0.01; c = 0.01; d = 0.01
# MCMC step
set.seed(90210)
burn = 10; M = 1000
niter = burn + M
V1 = V; W1 = W
draws = NULL
all_draws = NULL
for (iter in 1:niter){
  run = ffbs(y,V1,W1,m0,C0)
  x = run$x
  V1 = 1/rgamma(1,a+n/2,b+sum((y-x)^2)/2)
  W1 = 1/rgamma(1,c+(n-1)/2,d+sum(diff(x)^2)/2)
  draws = rbind(draws,c(V1,W1,x))
}
all_draws = draws[,1:2]
q025 = function(x){quantile(x,0.025)}
q975 = function(x){quantile(x,0.975)}
draws = draws[(burn+1):(niter),]
xs = draws[,3:(n+2)]
lx = apply(xs,2,q025)
mx = apply(xs,2,mean)
ux = apply(xs,2,q975)
## plot of the data
par(mfrow=c(2,2), mgp=c(1.6,.6,0), mar=c(3,3.2,1,1))
ts.plot(ts(x), ts(y), ylab='', col=c(1,8), lwd=2)
points(y)
legend(0,11,legend=c("x(t)", "y(t)"),lty=1,col=c(1,8),lwd=2,bty="n",pch=c(-1,1))
contour(Vs, Ws, exp(likes), xlab=expression(sigma[v]^2),
        ylab=expression(sigma[w]^2), drawlabels=FALSE, ylim=c(0,1.2))
points(draws[,1:2], pch=16, col=rgb(.9,0,0,0.3),cex=.7)
hist(draws[,1], ylab="Density",main="", xlab=expression(sigma[v]^2))

```

```

abline(v=mean(draws[,1]), col=3, lwd=3)
hist(draws[,2],main="", ylab="Density", xlab=expression(sigma[w]^2))
abline(v=mean(draws[,2]), col=3, lwd=3)
## plot states
dev.new()
par(mgp=c(1.6,.6,0), mar=c(3,3.2,1,1))
ts.plot(cbind(lx,ux), ylab="", main="", lty=2, col= rgb(0,.6,.3,alpha=.2),
        ylim=c(min(y),max(y)))
lines(mx, lwd=2,lty=1 ,col= rgb(0,.3,0,alpha=.5))
lines(xsmooth, lwd=2, lty=2, col= rgb(1,0,0,alpha=.75))
xx=c(1:100, 100:1)
yy=c(lx, rev(ux))
polygon(xx, yy, border=NA, col= rgb(0,.6,.3,alpha=.1),lty=2)
lines(y, type='o', pch=20, col=rgb(.4,.4,.4,alpha=1) )
legend('topleft', c('true smoother', 'data', 'posterior mean'),
        lty=c(2,1,1), lwd=c(2,1,2), pch=c(-1,20,-1),
col= c(rgb(1,0,0,alpha=.5), col=rgb(.4,.4,.4,alpha=1) , col=
      rgb(0,.3,0,alpha=.5)) )

```

Next, we consider a more complicated model.

Example 6.23 Structural Model

Consider the Johnson & Johnson quarterly earnings per share series that was discussed in [Example 6.11](#). Recall that the model is

$$y_t = \begin{pmatrix} 1 & 1 & 0 & 0 \end{pmatrix} x_t + v_t,$$

$$x_t = \begin{pmatrix} T_t \\ S_t \\ S_{t-1} \\ S_{t-2} \end{pmatrix} = \begin{pmatrix} \phi & 0 & 0 & 0 \\ 0 & -1 & -1 & -1 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} T_{t-1} \\ S_{t-1} \\ S_{t-2} \\ S_{t-3} \end{pmatrix} + \begin{pmatrix} w_{t1} \\ w_{t2} \\ 0 \\ 0 \end{pmatrix}$$

where $R = \sigma_v^2$ and

$$Q = \begin{pmatrix} \sigma_{w,11}^2 & 0 & 0 & 0 \\ 0 & \sigma_{w,22}^2 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}.$$

The parameters to be estimated are the transition parameter associated with the growth rate, $\phi > 1$, the observation noise variance, σ_v^2 , and the state noise variances associated with the trend and the seasonal components, $\sigma_{w,11}^2$ and $\sigma_{w,22}^2$, respectively.

In this case, sampling from $p(\mathcal{X}_n \mid \Theta, \mathcal{Y}_n)$ follows directly from (6.195)–(6.196). Next, we discuss how to sample from $p(\Theta \mid \mathcal{X}_n, \mathcal{Y}_n)$. For the transition parameter, write $\phi = 1 + \beta$, where $0 < \beta \ll 1$; recall that in [Example 6.11](#), ϕ was estimated to be 1.035, which indicated a growth rate, β , of 3.5%. Note that the trend component may be rewritten as

$$\nabla T_t = T_t - T_{t-1} = \beta T_{t-1} + w_{t1}.$$




Fig. 6.18. Parameter estimation results for **Example 6.23**. The top row displays the traces of 1000 draws after burn-in. The middle row displays the ACF of the traces. The sampled posteriors are displayed in the last row (the mean is marked by a solid vertical line).

Consequently, conditional on the states, the parameter β is the slope in the linear regression (through the origin) of ∇T_t on T_{t-1} , for $t = 1, \dots, n$, and w_{t1} is the




Fig. 6.19. **Example 6.23** smoother estimates of trend (T_t) and season (S_t) along with corresponding 99% credible intervals.

error. As is typical, we put a Normal–Inverse Gamma (IG) prior on $(\beta, \sigma_{w,11}^2)$, i.e., $\beta \mid \sigma_{w,11}^2 \sim N(b_0, \sigma_{w,11}^2 B_0)$ and $\sigma_{w,11}^2 \sim \text{IG}(n_0/2, n_0 s_0^2/2)$, with known hyperparameters b_0, B_0, n_0, s_0^2 .

We also used IG priors for the other two variance components, σ_v^2 and $\sigma_{w,22}^2$. In this case, if the prior $\sigma_v^2 \sim \text{IG}(n_0/2, n_0 s_0^2/2)$, then the posterior $\sigma_v^2 \mid \mathcal{X}_n, \mathcal{Y}_n \sim \text{IG}(n_v/2, n_v s_v^2/2)$, where $n_v = n_0 + n$, and $n_v s_v^2 = n_0 s_0^2 + \sum_{t=1}^n (Y_t - T_t - S_t)^2$. Similarly, if the prior $\sigma_{w,22}^2 \sim \text{IG}(n_0/2, n_0 s_0^2/2)$, then the posterior $\sigma_{w,22}^2 \mid \mathcal{X}_n, \mathcal{Y}_n \sim \text{IG}(n_w/2, n_w s_w^2/2)$, where $n_w = n_0 + (n - 3)$, and $n_w s_w^2 = n_0 s_0^2 + \sum_{t=1}^{n-3} (S_t - S_{t-1} - S_{t-2} - S_{t-3})^2$.

Figure 6.18 displays the results of the posterior estimates of the parameters. The top row of the figure displays the traces of 1000 draws, after a burn-in of 100, with a step size of 10 (i.e., every 10th sampled value is retained). The middle row of the figure displays the ACF of the traces, and the sampled posteriors are displayed in the last row of the figure. The results of this analysis are comparable to the results obtained in Example 6.11; the posterior mean and median for ϕ indicates a 3.7% growth rate in the Johnson & Johnson quarterly earnings over this time period.

Figure 6.19 displays the smoothers of trend (T_t) and season (S_t) along with 99% credible intervals. Again, these results are comparable to the results obtained in Example 6.11.

The R code for this example is as follows:

```
require(plyr) # used to view progress (install it if you don't have it)
y = jj
### setup - model and initial parameters
set.seed(90210)
n = length(y)
F = c(1,1,0,0) # this is A
G = diag(0,4) # G is Phi
G[1,1] = 1.03
G[2,]=c(0,-1,-1,-1); G[3,]=c(0,1,0,0); G[4,]=c(0,0,1,0)
a1 = rbind(.7,0,0,0) # this is mu0
R1 = diag(.04,4) # this is Sigma0
V = .1
W11 = .1
W22 = .1
##-- FFBS --##
ffbs = function(y,F,G,V,W11,W22,a1,R1){
  n = length(y)
  Ws = diag(c(W11,W22,1,1)) # this is Q with 1s as a device only
  iW = diag(1/diag(Ws),4)
  a = matrix(0,n,4) # this is m_t
  R = array(0,c(n,4,4)) # this is V_t
  m = matrix(0,n,4)
  C = array(0,c(n,4,4))
  a[1,] = a1[1,]
  R[1,,] = R1
  f = t(F)%*%a[1,]
  Q = t(F)%*%R[1,,]%*%F+V
  A = R[1,,]%*%F/Q[1,1]
  m[1,] = a[1,]+A%*%(y[1]-f)
  C[1,,] = R[1,,]-A%*%t(A)*Q[1,1]
```

```

for (t in 2:n){
  a[t,] = G%%m[t-1,]
  R[t,,] = G%%C[t-1,,]%%t(G) + Ws
  f      = t(F)%%a[t,]
  Q      = t(F)%%R[t,,]%%F+V
  A      = R[t,,]%%F/Q[1,1]
  m[t,]  = a[t,]+A%%(y[t]-f)
  C[t,,] = R[t,,]-A%%t(A)*Q[1,1]
}
xb      = matrix(0,n,4)
xb[n,]  = m[n,] + t(chol(C[n,,]))%%rnorm(4)
for (t in (n-1):1){
  iC = solve(C[t,,])
  CCC = solve(t(G)%%iW%%G+iC)
  mmm = CCC%%(t(G)%%iW%%xb[t+1,]+iC%%m[t,])
  xb[t,] = mmm + t(chol(CCC))%%rnorm(4)
}
return(xb)
}

##-- Prior hyperparameters --##
# b0 = 0 # mean for beta = phi -1
# B0 = Inf # var for beta (non-informative => use OLS for sampling beta)
n0 = 10 # use same for all- the prior is 1/Gamma(n0/2, n0*s20_/2)
s20v = .001 # for V
s20w = .05 # for Ws
##-- MCMC scheme --##
set.seed(90210)
burnin = 100
step = 10
M = 1000
niter = burnin+step*M
pars = matrix(0,niter,4)
xbs = array(0,c(niter,n,4))
pr <- progress_text() # displays progress
pr$init(niter)
for (iter in 1:niter){
  xb = ffbs(y,F,G,V,W11,W22,a1,R1)
  u = xb[,1]
  yu = diff(u); xu = u[-n] # for phihat and se(phihat)
  regu = lm(yu~0+xu) # est of beta = phi-1
  phies = as.vector(coef(summary(regu)))[1:2] + c(1,0) # phi estimate and SE
  dft = df.residual(regu)
  G[1,1] = phies[1] + rt(1,dft)*phies[2] # use a t
  V = 1/rgamma(1, (n0+n)/2, (n0*s20v/2) + sum((y-xb[,1]-xb[,2])^2)/2)
  W11 = 1/rgamma(1, (n0+n-1)/2, (n0*s20w/2) +
    sum((xb[-1,1]-phies[1]*xb[-n,1])^2)/2)
  W22 = 1/rgamma(1, (n0+ n-3)/2, (n0*s20w/2) + sum((xb[4:n,2] +
    xb[3:(n-1),2]+ xb[2:(n-2),2] +xb[1:(n-3),2])^2)/2)
  xbs[iter,,] = xb
  pars[iter,] = c(G[1,1], sqrt(V), sqrt(W11), sqrt(W22))
  pr$step()
}
# Plot results
ind = seq(burnin+1,niter,by=step)

```



```

names= c(expression(phi), expression(sigma[v]), expression(sigma[w~11]),
          expression(sigma[w~22]))
dev.new(height=5)
par(mfcol=c(3,4), mar=c(2,2,.25,0)+.75, mgp=c(1.6,.6,0), oma=c(0,0,1,0))
for (i in 1:4){
  plot.ts(pars[ind,i],xlab="iterations", ylab="trace", main="")
  mtext(names[i], side=3, line=.5, cex=1)
  acf(pars[ind,i],main="", lag.max=25, xlim=c(1,25), ylim=c(-.4,.4))
  hist(pars[ind,i],main="",xlab="")
  abline(v=mean(pars[ind,i]), lwd=2, col=3)
}
mxb = cbind(apply(xbs[ind,,1],2,mean), apply(xbs[,2],2,mean))
lxb = cbind(apply(xbs[ind,,1],2,quantile,0.005),
            apply(xbs[ind,,2],2,quantile,0.005))
uxb = cbind(apply(xbs[ind,,1],2,quantile,0.995),
            apply(xbs[ind,,2],2,quantile,0.995))
mxb = ts(mxb, start = tsp(jj)[1], freq=4)
lxb = ts(lxb, start = tsp(jj)[1], freq=4)
uxb = ts(uxb, start = tsp(jj)[1], freq=4)
names = c('Trend', 'Season')
dev.new(height=4.5)
par(mfrow=c(2,1), mar=c(2,2,0,0)+.7, mgp=c(1.6,.6,0))
for (i in 1:2){
  L = min(lxb[,i])-.01; U = max(uxb[,i]) +.01
  plot(mxb[,i], ylab=names[i], ylim=c(L,U))
  xx=c(time(jj), rev(time(jj)))
  yy=c(lxb[,i], rev(uxb[,i]))
  polygon(xx, yy, border=NA, col=gray(.4, alpha = .2))
}

```

Problems

Section 6.1

6.1 Consider a system process given by

$$x_t = -.9x_{t-2} + w_t \quad t = 1, \dots, n$$

where $x_0 \sim N(0, \sigma_0^2)$, $x_{-1} \sim N(0, \sigma_1^2)$, and w_t is Gaussian white noise with variance σ_w^2 . The system process is observed with noise, say,

$$y_t = x_t + v_t,$$

where v_t is Gaussian white noise with variance σ_v^2 . Further, suppose x_0 , x_{-1} , $\{w_t\}$ and $\{v_t\}$ are independent.

- Write the system and observation equations in the form of a state space model.
- Find the values of σ_0^2 and σ_1^2 that make the observations, y_t , stationary.
- Generate $n = 100$ observations with $\sigma_w = 1$, $\sigma_v = 1$ and using the values of σ_0^2 and σ_1^2 found in (b). Do a time plot of x_t and of y_t and compare the two processes. Also, compare the sample ACF and PACF of x_t and of y_t .

(d) Repeat (c), but with $\sigma_v = 10$.

6.2 Consider the state-space model presented in **Example 6.3**. Let $x_t^{t-1} = E(x_t \mid y_{t-1}, \dots, y_1)$ and let $P_t^{t-1} = E(x_t - x_t^{t-1})^2$. The innovation sequence or residuals are $\epsilon_t = y_t - y_t^{t-1}$, where $y_t^{t-1} = E(y_t \mid y_{t-1}, \dots, y_1)$. Find $\text{cov}(\epsilon_s, \epsilon_t)$ in terms of x_t^{t-1} and P_t^{t-1} for (i) $s \neq t$ and (ii) $s = t$.

Section 6.2

6.3 Simulate $n = 100$ observations from the following state-space model:

$$x_t = .8x_{t-1} + w_t \quad \text{and} \quad y_t = x_t + v_t$$

where $x_0 \sim N(0, 2.78)$, $w_t \sim \text{iid } N(0, 1)$, and $v_t \sim \text{iid } N(0, 1)$ are all mutually independent. Compute and plot the data, y_t , the one-step-ahead predictors, y_t^{t-1} along with the root mean square prediction errors, $E^{1/2}(y_t - y_t^{t-1})^2$ using **Figure 6.3** as a guide.

6.4 Suppose the vector $z = (x', y')'$, where x ($p \times 1$) and y ($q \times 1$) are jointly distributed with mean vectors μ_x and μ_y and with covariance matrix

$$\text{cov}(z) = \begin{pmatrix} \Sigma_{xx} & \Sigma_{xy} \\ \Sigma_{yx} & \Sigma_{yy} \end{pmatrix}.$$

Consider projecting x on $\mathcal{M} = \overline{\text{sp}}\{1, y\}$, say, $\hat{x} = \mathbf{b} + By$.

(a) Show the orthogonality conditions can be written as

$$E(x - \mathbf{b} - By) = 0,$$

$$E[(x - \mathbf{b} - By)y'] = 0,$$

leading to the solutions

$$\mathbf{b} = \mu_x - B\mu_y \quad \text{and} \quad B = \Sigma_{xy}\Sigma_{yy}^{-1}.$$

(b) Prove the mean square error matrix is

$$MSE = E[(x - \mathbf{b} - By)x'] = \Sigma_{xx} - \Sigma_{xy}\Sigma_{yy}^{-1}\Sigma_{yx}.$$

(c) How can these results be used to justify the claim that, in the absence of normality, **Property 6.1** yields the best linear estimate of the state x_t given the data Y_t , namely, x_t^t , and its corresponding MSE, namely, P_t^t ?

6.5 Projection Theorem Derivation of Property 6.2. Throughout this problem, we use the notation of **Property 6.2** and of the Projection Theorem given in **Appendix B**, where $\mathcal{H}$ is L^2 . If $\mathcal{L}_{k+1} = \overline{\text{sp}}\{y_1, \dots, y_{k+1}\}$, and $\mathcal{V}_{k+1} = \overline{\text{sp}}\{y_{k+1} - y_{k+1}^k\}$, for $k = 0, 1, \dots, n-1$, where y_{k+1}^k is the projection of y_{k+1} on $\mathcal{L}_k$, then, $\mathcal{L}_{k+1} = \mathcal{L}_k \oplus \mathcal{V}_{k+1}$. We assume $P_0^0 > 0$ and $R > 0$.

(a) Show the projection of x_k on $\mathcal{L}_{k+1}$, that is, x_k^{k+1} , is given by

$$x_k^{k+1} = x_k^k + H_{k+1}(y_{k+1} - y_{k+1}^k),$$

where H_{k+1} can be determined by the orthogonality property

$$E \left\{ (x_k - H_{k+1}(y_{k+1} - y_{k+1}^k)) (y_{k+1} - y_{k+1}^k)' \right\} = 0.$$

Show

$$H_{k+1} = P_k^k \Phi' A'_{k+1} [A_{k+1} P_{k+1}^k A'_{k+1} + R]^{-1}.$$

(b) Define $J_k = P_k^k \Phi' [P_{k+1}^k]^{-1}$, and show

$$x_k^{k+1} = x_k^k + J_k(x_{k+1}^{k+1} - x_{k+1}^k).$$

(c) Repeating the process, show

$$x_k^{k+2} = x_k^k + J_k(x_{k+1}^{k+1} - x_{k+1}^k) + H_{k+2}(y_{k+2} - y_{k+2}^k),$$

solving for H_{k+2} . Simplify and show

$$x_k^{k+2} = x_k^k + J_k(x_{k+1}^{k+2} - x_{k+1}^k).$$

(d) Using induction, conclude

$$x_k^n = x_k^k + J_k(x_{k+1}^n - x_{k+1}^k),$$

which yields the smoother with $k = t - 1$.

Section 6.3

6.6 Consider the univariate state-space model given by state conditions $x_0 = w_0$, $x_t = x_{t-1} + w_t$ and observations $y_t = x_t + v_t$, $t = 1, 2, \dots$, where w_t and v_t are independent, Gaussian, white noise processes with $\text{var}(w_t) = \sigma_w^2$ and $\text{var}(v_t) = \sigma_v^2$.

(a) Show that y_t follows an IMA(1,1) model, that is, ∇y_t follows an MA(1) model.

(b) Fit the model specified in part (a) to the logarithm of the glacial varve series and compare the results to those presented in **Example 3.32**.

6.7 Consider the model

$$y_t = x_t + v_t,$$

where v_t is Gaussian white noise with variance σ_v^2 , x_t are independent Gaussian random variables with mean zero and $\text{var}(x_t) = r_t \sigma_x^2$ with x_t independent of v_t , and $r_1, \dots, r_n$ are known constants. Show that applying the EM algorithm to the problem of estimating σ_x^2 and σ_v^2 leads to updates (represented by hats)

$$\hat{\sigma}_x^2 = \frac{1}{n} \sum_{t=1}^n \frac{\sigma_t^2 + \mu_t^2}{r_t} \quad \text{and} \quad \hat{\sigma}_v^2 = \frac{1}{n} \sum_{t=1}^n [(y_t - \mu_t)^2 + \sigma_t^2],$$

where, based on the current estimates (represented by tildes),

$$\mu_t = \frac{r_t \tilde{\sigma}_x^2}{r_t \tilde{\sigma}_x^2 + \tilde{\sigma}_v^2} y_t \quad \text{and} \quad \sigma_t^2 = \frac{r_t \tilde{\sigma}_x^2 \tilde{\sigma}_v^2}{r_t \tilde{\sigma}_x^2 + \tilde{\sigma}_v^2}.$$

6.8 To explore the stability of the filter, consider a univariate state-space model. That is, for $t = 1, 2, \dots$, the observations are $y_t = x_t + v_t$ and the state equation is $x_t = \phi x_{t-1} + w_t$, where $\sigma_w = \sigma_v = 1$ and $|\phi| < 1$. The initial state, x_0 , has zero mean and variance one.

- (a) Exhibit the recursion for P_t^{t-1} in **Property 6.1** in terms of P_{t-1}^{t-2} .
- (b) Use the result of (a) to verify P_t^{t-1} approaches a limit ($t \rightarrow \infty$) P that is the positive solution of $P^2 - \phi^2 P - 1 = 0$.
- (c) With $K = \lim_{t \rightarrow \infty} K_t$ as given in **Property 6.1**, show $|1 - K| < 1$.
- (d) Show, in steady-state, the one-step-ahead predictor, $y_{n+1}^n = E(y_{n+1} \mid y_n, y_{n-1}, \dots)$, of a future observation satisfies

$$y_{n+1}^n = \sum_{j=0}^{\infty} \phi^j K (1 - K)^{j-1} y_{n+1-j}.$$

6.9 In §6.3, we discussed that it is possible to obtain a recursion for the gradient vector, $-\partial \ln L_Y(\Theta)/\partial \Theta$. Assume the model is given by (6.1) and (6.2) and A_t is a known design matrix that does not depend on Θ , in which case **Property 6.1** applies. For the gradient vector, show

$$\begin{aligned} \partial \ln L_Y(\Theta)/\partial \Theta_i = \sum_{t=1}^n \left\{ \epsilon_t' \Sigma_t^{-1} \frac{\partial \epsilon_t}{\partial \Theta_i} - \frac{1}{2} \epsilon_t' \Sigma_t^{-1} \frac{\partial \Sigma_t}{\partial \Theta_i} \Sigma_t^{-1} \epsilon_t \right. \\ \left. + \frac{1}{2} \text{tr} \left(\Sigma_t^{-1} \frac{\partial \Sigma_t}{\partial \Theta_i} \right) \right\}, \end{aligned}$$

where the dependence of the innovation values on Θ is understood. In addition, with the general definition $\partial_i g = \partial g(\Theta)/\partial \Theta_i$, show the following recursions, for $t = 2, \dots, n$ apply:

- (i) $\partial_i \epsilon_t = -A_t \partial_i x_t^{t-1}$,
- (ii) $\partial_i x_t^{t-1} = \partial_i \Phi x_{t-1}^{t-2} + \Phi \partial_i x_{t-1}^{t-2} + \partial_i K_{t-1} \epsilon_{t-1} + K_{t-1} \partial_i \epsilon_{t-1}$,
- (iii) $\partial_i \Sigma_t = A_t \partial_i P_{t-1}^{t-1} A_t' + \partial_i R$,
- (iv) $\partial_i K_t = [\partial_i \Phi P_{t-1}^{t-1} A_t' + \Phi \partial_i P_{t-1}^{t-1} A_t' - K_t \partial_i \Sigma_t] \Sigma_t^{-1}$,
- (v) $\partial_i P_t^{t-1} = \partial_i \Phi P_{t-1}^{t-2} \Phi' + \Phi \partial_i P_{t-1}^{t-2} \Phi' + \Phi P_{t-1}^{t-2} \partial_i \Phi' + \partial_i Q$,
 $- \partial_i K_{t-1} \Sigma_t K_{t-1}' - K_{t-1} \partial_i \Sigma_t K_{t-1}' - K_{t-1} \Sigma_t \partial_i K_{t-1}'$,

using the fact that $P_t^{t-1} = \Phi P_{t-1}^{t-2} \Phi' + Q - K_{t-1} \Sigma_t K_{t-1}'$.

6.10 Continuing with the previous problem, consider the evaluation of the Hessian matrix and the numerical evaluation of the asymptotic variance-covariance matrix of the parameter estimates. The information matrix satisfies

$$E \left\{ -\frac{\partial^2 \ln L_Y(\Theta)}{\partial \Theta \partial \Theta'} \right\} = E \left\{ \left(\frac{\partial \ln L_Y(\Theta)}{\partial \Theta} \right) \left(\frac{\partial \ln L_Y(\Theta)}{\partial \Theta} \right)' \right\};$$

see Anderson (1984, Section 4.4), for example. Show the (i, j) -th element of the information matrix, say, $\mathcal{I}_{ij}(\Theta) = E \left\{ -\partial^2 \ln L_Y(\Theta)/\partial \Theta_i \partial \Theta_j \right\}$, is

$$\begin{aligned} \mathcal{I}_{ij}(\Theta) = \sum_{t=1}^n E \Big\{ & \partial_i \epsilon_t' \Sigma_t^{-1} \partial_j \epsilon_t + \frac{1}{2} \text{tr}(\Sigma_t^{-1} \partial_i \Sigma_t \Sigma_t^{-1} \partial_j \Sigma_t) \\ & + \frac{1}{4} \text{tr}(\Sigma_t^{-1} \partial_i \Sigma_t) \text{tr}(\Sigma_t^{-1} \partial_j \Sigma_t) \Big\}. \end{aligned}$$

Consequently, an approximate Hessian matrix can be obtained from the sample by dropping the expectation, E , in the above result and using only the recursions needed to calculate the gradient vector.

Section 6.4

6.11 As an example of the way the state-space model handles the missing data problem, suppose the first-order autoregressive process

$$x_t = \phi x_{t-1} + w_t$$

has an observation missing at $t = m$, leading to the observations $y_t = A_t x_t$, where $A_t = 1$ for all t , except $t = m$ wherein $A_t = 0$. Assume $x_0 = 0$ with variance $\sigma_w^2/(1 - \phi^2)$, where the variance of w_t is σ_w^2 . Show the Kalman smoother estimators in this case are

$$x_t^n = \begin{cases} \phi y_1 & t = 0, \\ \frac{\phi}{1 + \phi^2} (y_{m-1} + y_{m+1}) & t = m, \\ y, & t \neq 0, m, \end{cases}$$

with mean square covariances determined by

$$P_t^n = \begin{cases} \sigma_w^2 & t = 0, \\ \sigma_w^2/(1 + \phi^2) & t = m, \\ 0 & t \neq 0, m. \end{cases}$$

6.12 The data set [ar1miss](#) is $n = 100$ observations generated from an AR(1) process, $x_t = \phi x_{t-1} + w_t$, with $\phi = .9$ and $\sigma_w = 1$, where 10% of the data has been zeroed out at random. Considering the zeroed out data to be missing data, use the results of [Problem 6.11](#) to estimate the parameters of the model, ϕ and σ_w , using the EM algorithm, and then estimate the missing values.

Section 6.5

6.13 Let y_t represent the global temperature series ([gtemp](#)) shown in [Figure 1.2](#).

- (a) Fit a smoothing spline using `gcv` (the default) to y_t and plot the result superimposed on the data. Repeat the fit using `spar=.7`; the `gcv` method yields `spar=.5` approximately. ([Example 2.14](#) on page 71 may help. Also in R, see the help file [?smooth.spline](#).)

- (b) Write the model $y_t = x_t + v_t$ with $\nabla^2 x_t = w_t$, in state-space form. Fit this state-space model to y_t , and exhibit a time plot the estimated smoother, $\hat{x}_t^n$ and the corresponding error limits, $\hat{x}_t^n \pm 2\sqrt{\hat{P}_t^n}$ superimposed on the data.
- (c) Superimpose all the fits from parts (a) and (b) [include the error bounds] on the data and briefly compare and contrast the results.

Section 6.6

6.14 Redo [Example 6.11](#) on the *logged* Johnson & Johnson quarterly earnings per share.

6.15 Fit a structural model to quarterly unemployment as follows. Use the data in `unemp`, which are monthly. The series can be made quarterly by aggregating and averaging: `y = aggregate(unemp, nfrequency=4, FUN=mean)`, so that `y` is the quarterly average unemployment. Use [Example 6.11](#) as a guide.

Section 6.7

- 6.16** (a) Fit an AR(2) to the recruitment series, R_t in `rec`, and consider a lag-plot of the residuals from the fit versus the SOI series, S_t in `soi`, at various lags, S_{t-h} , for $h = 0, 1, \dots$. Use the lag-plot to argue that S_{t-5} is reasonable to include as an exogenous variable.
- (b) Fit an ARX(2) to R_t using S_{t-5} as an exogenous variable and comment on the results; include an examination of the innovations.

6.17 Use [Property 6.6](#) to complete the following exercises.

- (a) Write a univariate AR(1) model, $y_t = \phi y_{t-1} + v_t$, in state-space form. Verify your answer is indeed an AR(1).
- (b) Repeat (a) for an MA(1) model, $y_t = v_t + \theta v_{t-1}$.
- (c) Write an IMA(1,1) model, $y_t = y_{t-1} + v_t + \theta v_{t-1}$, in state-space form.

6.18 Verify [Property 6.5](#).

6.19 Verify [Property 6.6](#).

Section 6.8

6.20 Repeat the bootstrap analysis of [Example 6.14](#) on the entire three-month Treasury bills and rate of inflation data set of 110 observations. Do the conclusions of [Example 6.14](#)—that the dynamics of the data are best described in terms of a fixed, rather than stochastic, regression—still hold?

Section 6.9

6.21 Fit the switching model described in [Example 6.16](#) to the growth rate of GNP. The data are in [gnp](#) and, in the notation of the example, y_t is log-GNP and ∇y_t is the growth rate. Use the code in [Example 6.18](#) as a guide.

6.22 Argue that a switching model is reasonable in explaining the behavior of the number of sunspots (see [Figure 4.22](#)) and then fit a switching model to the sunspot data.

Section 6.10

6.23 Fit a stochastic volatility model to the returns of one (or more) of the four financial time series available in the R datasets package as [EuStockMarkets](#).

6.24 Fit a stochastic volatility model to the residuals of the GNP ([gnp](#)) returns analyzed in [Example 3.38](#).

Section 6.11

6.25 Repeat [Example 6.23](#) on the log of the Johnson & Johnson data.

6.26 Fit an AR(1) to the returns of the US GNP ([gnp](#)) using a Bayesian approach via MCMC.

Chapter 7

Statistical Methods in the Frequency Domain

7.1 Introduction

In previous chapters, we saw many applied time series problems that involved relating series to each other or to evaluating the effects of treatments or design parameters that arise when time-varying phenomena are subjected to periodic stimuli. In many cases, the nature of the physical or biological phenomena under study are best described by their Fourier components rather than by the difference equations involved in ARIMA or state-space models. The fundamental tools we use in studying periodic phenomena are the discrete Fourier transforms (DFTs) of the processes and their statistical properties. Hence, in §7.2, we review the properties of the DFT of a multivariate time series and discuss various approximations to the likelihood function based on the large-sample properties and the properties of the complex multivariate normal distribution. This enables extension of the classical techniques discussed in the following paragraphs to the multivariate time series case.

An extremely important class of problems in classical statistics develops when we are interested in relating a collection of input series to some output series. For example, in Chapter 2, we have previously considered relating temperature and various pollutant levels to daily mortality, but have not investigated the frequencies that appear to be driving the relation and have not looked at the possibility of leading or lagging effects. In Chapter 4, we isolated a definite lag structure that could be used to relate sea surface temperature to the number of new recruits. In Problem 5.10, the possible driving processes that could be used to explain inflow to Lake Shasta were hypothesized in terms of the possible inputs precipitation, cloud cover, temperature, and other variables. Identifying the combination of input factors that produce the best prediction for inflow is an example of multiple regression in the frequency domain, with the models treated theoretically by considering the regression, conditional on the random input processes.

A situation somewhat different from that above would be one in which the input series are regarded as fixed and known. In this case, we have a model analogous to that occurring in analysis of variance, in which the analysis now can be performed on a frequency by frequency basis. This analysis works especially well when the inputs




Fig. 7.1. Mean response of subjects to various combinations of periodic stimuli measured at the cortex (primary somatosensory, contralateral). In the first column, the subjects are awake, in the second column the subjects are under mild anesthesia. In the first row, the stimulus is a brush on the hand, the second row involves the application of heat, and the third row involves a low level shock.

are dummy variables, depending on some configuration of treatment and other design effects and when effects are largely dependent on periodic stimuli. As an example, we will look at a designed experiment measuring the fMRI brain responses of a number of awake and mildly anesthetized subjects to several levels of periodic brushing, heat, and shock effects. Some limited data from this experiment have been discussed previously in [Example 1.6](#) of Chapter 1. [Figure 7.1](#) shows mean responses to various levels of periodic heat, brushing, and shock stimuli for subjects awake and subjects under mild anesthesia. The stimuli were periodic in nature, applied alternately for 32 seconds (16 points) and then stopped for 32 seconds. The periodic input signal comes through under all three design conditions when the subjects are awake, but is somewhat attenuated under anesthesia. The mean shock level response hardly shows on the input signal; shock levels were designed to simulate surgical incision without inflicting tissue damage. The means in [Figure 7.1](#) are from a single location. Actually, for each individual, some nine series were recorded at various locations in the brain. It is natural to consider testing the effects of brushing, heat, and shock under the two levels of consciousness, using a time series generalization of analysis of variance. The R code used to generate [Figure 7.1](#) is:

```
x = matrix(0, 128, 6)
for (i in 1:6) x[,i] = rowMeans(fmri[[i]])
```

```
colnames(x)=c("Brush", "Heat", "Shock", "Brush", "Heat", "Shock")
plot.ts(x, main="")
mtext("Awake", side=3, line=1.2, adj=.05, cex=1.2)
mtext("Sedated", side=3, line=1.2, adj=.85, cex=1.2)
```

A generalization to random coefficient regression is also considered, paralleling the univariate approach to signal extraction and detection presented in §???. This method enables a treatment of multivariate ridge-type regressions and inversion problems. Also, the usual random effects analysis of variance in the frequency domain becomes a special case of the random coefficient model.

The extension of frequency domain methodology to more classical approaches to multivariate discrimination and clustering is of interest in the frequency dependent case. Many time series differ in their means and in their autocovariance functions, making the use of both the mean function and the spectral density matrices relevant. As an example of such data, consider the bivariate series consisting of the P and S components derived from several earthquakes and explosions, such as those shown in [Figure 7.2](#), where the P and S components, representing different arrivals have been separated from the first and second halves, respectively, of waveforms like those shown originally in [Figure 1.7](#).

Two earthquakes and two explosions from a set of eight earthquakes and explosions are shown in [Figure 7.2](#) and some essential differences exist that might be used to characterize the two classes of events. Also, the frequency content of the two components of the earthquakes appears to be lower than those of the explosions, and relative amplitudes of the two classes appear to differ. For example, the ratio of the S to P amplitudes in the earthquake group is much higher for this restricted subset. Spectral differences were also noticed in Chapter 4, where the explosion processes had a stronger high-frequency component relative to the low-frequency contributions. Examples like these are typical of applications in which the essential differences between multivariate time series can be expressed by the behavior of either the frequency-dependent mean value functions or the spectral matrix. In discriminant analysis, these types of differences are exploited to develop combinations of linear and quadratic classification criteria. Such functions can then be used to classify events of unknown origin, such as the Novaya Zemlya event shown in [Figure 7.2](#), which tends to bear a visual resemblance to the explosion group.

The R code used to produce [Figure 7.2](#) is:

```
attach(eqexp)
P = 1:1024; S = P+1024
x = cbind(EQ5[P], EQ6[P], EX5[P], EX6[P], NZ[P], EQ5[S], EQ6[S], EX5[S],
          EX6[S], NZ[S])
x.name = c("EQ5", "EQ6", "EX5", "EX6", "NZ")
colnames(x) = c(x.name, x.name)
plot.ts(x, main="")
mtext("P waves", side=3, line=1.2, adj=.05, cex=1.2)
mtext("S waves", side=3, line=1.2, adj=.85, cex=1.2)
```

Finally, for multivariate processes, the structure of the spectral matrix is also of great interest. We might reduce the dimension of the underlying process to a smaller set of input processes that explain most of the variability in the cross-spectral matrix as a function of frequency. Principal component analysis can be used to decompose




Fig. 7.2. Various bivariate earthquakes (EQ) and explosions (EX) recorded at 40 pts/sec compared with an event NZ (Novaya Zemlya) of unknown origin. Compressional waves, also known as primary or P waves, travel fastest in the Earth's crust and are first to arrive. Shear waves propagate more slowly through the Earth and arrive second, hence they are called secondary or S waves.

the spectral matrix into a smaller subset of component factors that explain decreasing amounts of power. For example, the hydrological data might be explained in terms of a component process that weights heavily on precipitation and inflow and one that weights heavily on temperature and cloud cover. Perhaps these two components could explain most of the power in the spectral matrix at a given frequency. The ideas behind principal component analysis can also be generalized to include an optimal scaling methodology for categorical data called the spectral envelope (see Stoffer et al., 1993). In succeeding sections, we also give an introduction to dynamic Fourier analysis and to wavelet analysis.

7.2 Spectral Matrices and Likelihood Functions

We have previously argued for an approximation to the log likelihood based on the joint distribution of the DFTs in (4.81), where we used approximation as an aid in estimating parameters for certain parameterized spectra. In this chapter, we make heavy use of the fact that the sine and cosine transforms of the $p \times 1$ vector process $\mathbf{x}_t = (x_{t1}, x_{t2}, \dots, x_{tp})'$ with mean $E\mathbf{x}_t = \boldsymbol{\mu}_t$, say, with DFT¹

$$\mathbf{X}(\omega_k) = n^{-1/2} \sum_{t=1}^n \mathbf{x}_t e^{-2\pi i \omega_k t} = \mathbf{X}_c(\omega_k) - i\mathbf{X}_s(\omega_k) \quad (7.1)$$

and mean

$$\mathbf{M}(\omega_k) = n^{-1/2} \sum_{t=1}^n \boldsymbol{\mu}_t e^{-2\pi i \omega_k t} = \mathbf{M}_c(\omega_k) - i\mathbf{M}_s(\omega_k) \quad (7.2)$$

will be approximately uncorrelated, where we evaluate at the usual Fourier frequencies $\{\omega_k = k/n, 0 < |\omega_k| < 1/2\}$. By **Theorem C.6**, the approximate $2p \times 2p$ covariance matrix of the cosine and sine transforms, say, $\mathbf{X}(\omega_k) = (\mathbf{X}_c(\omega_k)', \mathbf{X}_s(\omega_k)')'$, is

$$\Sigma(\omega_k) = \frac{1}{2} \begin{pmatrix} C(\omega_k) & -Q(\omega_k) \\ Q(\omega_k) & C(\omega_k) \end{pmatrix}, \quad (7.3)$$

and the real and imaginary parts are jointly normal. This result implies, by the results stated in Appendix C, the density function of the vector DFT, say, $\mathbf{X}(\omega_k)$, can be approximated as

$$p(\omega_k) \approx |f(\omega_k)|^{-1} \exp\{-(\mathbf{X}(\omega_k) - \mathbf{M}(\omega_k))^* f^{-1}(\omega_k)(\mathbf{X}(\omega_k) - \mathbf{M}(\omega_k))\},$$

where the spectral matrix is the usual

$$f(\omega_k) = C(\omega_k) - iQ(\omega_k). \quad (7.4)$$

Certain computations that we do in the section on discriminant analysis will involve approximating the joint likelihood by the product of densities like the one given above over subsets of the frequency band $0 < \omega_k < 1/2$.

To use the likelihood function for estimating the spectral matrix, for example, we appeal to the limiting result implied by **Theorem C.7** and again choose L frequencies in the neighborhood of some target frequency ω , say, $\mathbf{X}(\omega_k \pm k/n)$, for $k = 1, \dots, m$ and $L = 2m + 1$. Then, let $\mathbf{X}_\ell$, for $\ell = 1, \dots, L$ denote the indexed values, and note the DFTs of the mean adjusted vector process are approximately jointly normal with mean zero and complex covariance matrix $f = f(\omega)$. Then, write the log likelihood over the L sub-frequencies as

¹ In previous chapters, the DFT of a process x_t was denoted by $d_x(\omega_k)$. In this chapter, we will consider the Fourier transforms of many different processes and so, to avoid the overuse of subscripts and to ease the notation, we use a capital letter, e.g., $X(\omega_k)$, to denote the DFT of x_t . This notation is standard in the digital signal processing (DSP) literature.




Fig. 7.3. Monthly values of weather and inflow at Lake Shasta.

$$\ln L(\mathbf{X}_1, \dots, \mathbf{X}_L; f) \approx -L \ln |f| - \sum_{\ell=1}^L (\mathbf{X}_\ell - \mathbf{M}_\ell)^* f^{-1} (\mathbf{X}_\ell - \mathbf{M}_\ell), \quad (7.5)$$

where we have suppressed the argument of $f = f(\omega)$ for ease of notation. The use of spectral approximations to the likelihood has been fairly standard, beginning with the work of Whittle (1961) and continuing in Brillinger (1981) and Hannan (1970). Assuming the mean adjusted series are available, i.e., $\mathbf{M}_\ell$ is known, we obtain the maximum likelihood estimator for f , namely,

$$\hat{f} = L^{-1} \sum_{\ell=1}^L (\mathbf{X}_\ell - \mathbf{M}_\ell)(\mathbf{X}_\ell - \mathbf{M}_\ell)^*; \quad (7.6)$$

see [Problem 7.2](#).

7.3 Regression for Jointly Stationary Series

In §4.8, we considered a model of the form

$$y_t = \sum_{r=-\infty}^{\infty} \beta_{1r} x_{t-r,1} + v_t, \quad (7.7)$$

where x_{t1} is a single observed input series and y_t is the observed output series, and we are interested in estimating the filter coefficients β_{1r} relating the adjacent lagged values of x_{t1} to the output series y_t . In the case of the SOI and Recruitment series, we identified the El Niño driving series as x_{t1} , the input and y_t , the Recruitment series, as the output. In general, more than a single plausible input series may exist. For example, the Lake Shasta inflow hydrological data ([climhyd](#)) shown in [Figure 7.3](#) suggests there may be at least five possible series driving the inflow; see [Example 7.1](#) for more details. Hence, we may envision a $q \times 1$ input vector of driving series, say, $\mathbf{x}_t = (x_{t1}, x_{t2}, \dots, x_{tq})'$, and a set of $q \times 1$ vector of regression functions $\boldsymbol{\beta}_r = (\beta_{1r}, \beta_{2r}, \dots, \beta_{qr})'$, which are related as

$$y_t = \sum_{r=-\infty}^{\infty} \boldsymbol{\beta}_r' \mathbf{x}_{t-r} + v_t = \sum_{j=1}^q \sum_{r=-\infty}^{\infty} \beta_{jr} x_{t-r,j} + v_t, \quad (7.8)$$

which shows that the output is a sum of linearly filtered versions of the input processes and a stationary noise process v_t , assumed to be uncorrelated with $\mathbf{x}_t$. Each filtered component in the sum over j gives the contribution of lagged values of the j -th input series to the output series. We assume the regression functions β_{jr} are fixed and unknown.

The model given by (7.8) is useful under several different scenarios, corresponding to a number of different assumptions that can be made about the components. Assuming the input and output processes are jointly stationary with zero means leads to the conventional regression analysis given in this section. The analysis depends on theory that assumes we observe the output process y_t conditional on fixed values of the input vector $\mathbf{x}_t$; this is the same as the assumptions made in conventional regression analysis. Assumptions considered later involve letting the coefficient vector $\boldsymbol{\beta}_t$ be a random unknown signal vector that can be estimated by Bayesian arguments, using the conditional expectation given the data. The answers to this approach, given in [§7.5](#), allow signal extraction and deconvolution problems to be handled. Assuming the inputs are fixed allows various experimental designs and analysis of variance to be done for both fixed and random effects models. Estimation of the frequency-dependent random effects variance components in the analysis of variance model is also considered in [§7.5](#).

For the approach in this section, assume the inputs and outputs have zero means and are jointly stationary with the $(q+1) \times 1$ vector process $(\mathbf{x}_t', y_t)'$ of inputs $\mathbf{x}_t$ and outputs y_t assumed to have a spectral matrix of the form

$$f(\omega) = \begin{pmatrix} f_{xx}(\omega) & f_{xy}(\omega) \\ f_{yx}(\omega) & f_{yy}(\omega) \end{pmatrix}, \quad (7.9)$$

where $f_{yx}(\omega) = (f_{yx_1}(\omega), f_{yx_2}(\omega), \dots, f_{yx_q}(\omega))$ is the $1 \times q$ vector of cross-spectra relating the q inputs to the output and $f_{xx}(\omega)$ is the $q \times q$ spectral matrix of the inputs. Generally, we observe the inputs and search for the vector of regression functions $\boldsymbol{\beta}_t$ relating the inputs to the outputs. We assume all autocovariance functions satisfy the absolute summability conditions of the form

$$\sum_{h=-\infty}^{\infty} |h| |\gamma_{jk}(h)| < \infty. \quad (7.10)$$

($j, k = 1, \dots, q+1$), where $\gamma_{jk}(h)$ is the autocovariance corresponding to the cross-spectrum $f_{jk}(\omega)$ in (7.9). We also need to assume a linear process of the form (C.35) as a condition for using Theorem C.7 on the joint distribution of the discrete Fourier transforms in the neighborhood of some fixed frequency.

ESTIMATION OF THE REGRESSION FUNCTION

In order to estimate the regression function β_r , the Projection Theorem (Appendix B) applied to minimizing

$$MSE = E \left[(y_t - \sum_{r=-\infty}^{\infty} \beta'_r \mathbf{x}_{t-r})^2 \right] \quad (7.11)$$

leads to the orthogonality conditions

$$E \left[(y_t - \sum_{r=-\infty}^{\infty} \beta'_r \mathbf{x}_{t-r}) \mathbf{x}'_{t-s} \right] = \mathbf{0}' \quad (7.12)$$

for all $s = 0, \pm 1, \pm 2, \dots$, where $\mathbf{0}'$ denotes the $1 \times q$ zero vector. Taking the expectations inside and substituting for the definitions of the autocovariance functions appearing and leads to the normal equations

$$\sum_{r=-\infty}^{\infty} \beta'_r \Gamma_{xx}(s-r) = \gamma'_{yx}(s), \quad (7.13)$$

for $s = 0, \pm 1, \pm 2, \dots$, where $\Gamma_{xx}(s)$ denotes the $q \times q$ autocovariance matrix of the vector series $\mathbf{x}_t$ at lag s and $\gamma_{yx}(s) = (\gamma_{yx_1}(s), \dots, \gamma_{yx_q}(s))$ is a $1 \times q$ vector containing the lagged covariances between y_t and $\mathbf{x}_t$. Again, a frequency domain approximate solution is easier in this case because the computations can be done frequency by frequency using cross-spectra that can be estimated from sample data using the DFT. In order to develop the frequency domain solution, substitute the representation into the normal equations, using the same approach as used in the simple case derived in §4.8. This approach yields

$$\int_{-1/2}^{1/2} \sum_{r=-\infty}^{\infty} \beta'_r e^{2\pi i \omega(s-r)} f_{xx}(\omega) d\omega = \gamma'_{yx}(s).$$

Now, because $\gamma'_{yx}(s)$ is the Fourier transform of the cross-spectral vector $f_{yx}(\omega) = f_{xy}^*(\omega)$, we might write the system of equations in the frequency domain, using the uniqueness of the Fourier transform, as

$$\mathbf{B}'(\omega) f_{xx}(\omega) = f_{xy}^*(\omega), \quad (7.14)$$

where $f_{xx}(\omega)$ is the $q \times q$ spectral matrix of the inputs and $\mathbf{B}(\omega)$ is the $q \times 1$ vector Fourier transform of $\boldsymbol{\beta}_t$. Multiplying (7.14) on the right by $f_{xx}^{-1}(\omega)$, assuming $f_{xx}(\omega)$ is nonsingular at ω , leads to the frequency domain estimator

$$\mathbf{B}'(\omega) = f_{xy}^*(\omega) f_{xx}^{-1}(\omega). \quad (7.15)$$

Note, (7.15) implies the regression function would take the form

$$\boldsymbol{\beta}_t = \int_{-1/2}^{1/2} \mathbf{B}(\omega) e^{2\pi i \omega t} d\omega. \quad (7.16)$$

As before, it is conventional to introduce the DFT as the approximate estimator for the integral (7.16) and write

$$\boldsymbol{\beta}_t^M = M^{-1} \sum_{k=0}^{M-1} \mathbf{B}(\omega_k) e^{2\pi i \omega_k t}, \quad (7.17)$$

where $\omega_k = k/M$, $M \ll n$. The approximation was shown in Problem 4.32 to hold exactly as long as $\boldsymbol{\beta}_t = \mathbf{0}$ for $|t| \geq M/2$ and to have a mean-squared-error bounded by a function of the zero-lag autocovariance and the absolute sum of the neglected coefficients.

The mean-squared-error (7.11) can be written using the orthogonality principle, giving

$$MSE = \int_{-1/2}^{1/2} f_{y \cdot x}(\omega) d\omega, \quad (7.18)$$

where

$$f_{y \cdot x}(\omega) = f_{yy}(\omega) - f_{xy}^*(\omega) f_{xx}^{-1}(\omega) f_{xy}(\omega) \quad (7.19)$$

denotes the residual or error spectrum. The resemblance of (7.19) to the usual equations in regression analysis is striking. It is useful to pursue the multiple regression analogy further by noting a squared multiple coherence can be defined as

$$\rho_{y \cdot x}^2(\omega) = \frac{f_{xy}^*(\omega) f_{xx}^{-1}(\omega) f_{xy}(\omega)}{f_{yy}(\omega)}. \quad (7.20)$$

This expression leads to the mean squared error in the form

$$MSE = \int_{-1/2}^{1/2} f_{yy}(\omega) [1 - \rho_{y \cdot x}^2(\omega)] d\omega, \quad (7.21)$$

and we have an interpretation of $\rho_{y \cdot x}^2(\omega)$ as the proportion of power accounted for by the lagged regression on $\mathbf{x}_t$ at frequency ω . If $\rho_{y \cdot x}^2(\omega) = 0$ for all ω , we have

$$MSE = \int_{-1/2}^{1/2} f_{yy}(\omega) d\omega = E[y_t^2],$$

which is the mean squared error when no predictive power exists. As long as $f_{xx}(\omega)$ is positive definite at all frequencies, $MSE \geq 0$, and we will have

$$0 \leq \rho_{y \cdot x}^2(\omega) \leq 1 \quad (7.22)$$

for all ω . If the multiple coherence is unity for all frequencies, the mean squared error in (7.21) is zero and the output series is perfectly predicted by a linearly filtered combination of the inputs. Problem 7.3 shows the ordinary squared coherence between the series y_t and the linearly filtered combinations of the inputs appearing in (7.11) is exactly (7.20).

ESTIMATION USING SAMPLED DATA

Clearly, the matrices of spectra and cross-spectra will not ordinarily be known, so the regression computations need to be based on sampled data. We assume, therefore, the inputs $x_{t1}, x_{t2}, \dots, x_{tq}$ and output y_t series are available at the time points $t = 1, 2, \dots, n$, as in Chapter 4. In order to develop reasonable estimates for the spectral quantities, some replication must be assumed. Often, only one replication of each of the inputs and the output will exist, so it is necessary to assume a band exists over which the spectra and cross-spectra are approximately equal to $f_{xx}(\omega)$ and $f_{xy}(\omega)$, respectively. Then, let $Y(\omega_k + \ell/n)$ and $\mathbf{X}(\omega_k + \ell/n)$ be the DFTs of y_t and $\mathbf{x}_t$ over the band, say, at frequencies of the form

$$\omega_k \pm \ell/n, \quad \ell = 1, \dots, m,$$

where $L = 2m + 1$ as before. Then, simply substitute the sample spectral matrix

$$\widehat{f}_{xx}(\omega) = L^{-1} \sum_{\ell=-m}^m \mathbf{X}(\omega_k + \ell/n) \mathbf{X}^*(\omega_k + \ell/n) \quad (7.23)$$

and the vector of sample cross-spectra

$$\widehat{f}_{xy}(\omega) = L^{-1} \sum_{\ell=-m}^m \mathbf{X}(\omega_k + \ell/n) \overline{Y(\omega_k + \ell/n)} \quad (7.24)$$

for the respective terms in (7.15) to get the regression estimator $\widehat{\mathbf{B}}(\omega)$. For the regression estimator (7.17), we may use

$$\widehat{\beta}_t^M = \frac{1}{M} \sum_{k=0}^{M-1} \widehat{f}_{xy}^*(\omega_k) \widehat{f}_{xx}^{-1}(\omega_k) e^{2\pi i \omega_k t} \quad (7.25)$$

for $t = 0, \pm 1, \pm 2, \dots, \pm(M/2 - 1)$, as the estimated regression function.

Table 7.1. ANOPOW for the Partitioned Regression Model

Source	Power	Degrees of Freedom
$x_{t,q_1+1}, \dots, x_{t,q_1+q_2}$	$SSR(\omega)$ (7.34)	$2q_2$
Error	$SSE(\omega)$ (7.35)	$2(L - q_1 - q_2)$
Total	$L\widehat{f}_{y \cdot 1}(\omega)$	$2(L - q_1)$

TESTS OF HYPOTHESES

The estimated squared multiple coherence, corresponding to the theoretical coherence (7.20), becomes

$$\widehat{\rho}_{y \cdot x}^2(\omega) = \frac{\widehat{f}_{xy}^*(\omega) \widehat{f}_{xx}^{-1}(\omega) \widehat{f}_{xy}(\omega)}{\widehat{f}_{yy}(\omega)}. \quad (7.26)$$

We may obtain a distributional result for the multiple coherence function analogous to that obtained in the univariate case by writing the multiple regression model in the frequency domain, as was done in §4.6. We obtain the statistic

$$F_{2q, 2(L-q)} = \frac{(L-q)}{q} \frac{\widehat{\rho}_{y \cdot x}^2(\omega)}{[1 - \widehat{\rho}_{y \cdot x}^2(\omega)]}, \quad (7.27)$$

which has an F -distribution with $2q$ and $2(L-q)$ degrees of freedom under the null hypothesis that $\rho_{y \cdot x}^2(\omega) = 0$, or equivalently, that $\mathbf{B}(\omega) = 0$, in the model

$$Y(\omega_k + \ell/n) = \mathbf{B}'(\omega)X(\omega_k + \ell/n) + V(\omega_k + \ell/n), \quad (7.28)$$

where the spectral density of the error $V(\omega_k + \ell/n)$ is $f_{y \cdot x}(\omega)$. **Problem 7.4** sketches a derivation of this result.

A second kind of hypothesis of interest is one that might be used to test whether a full model with q inputs is significantly better than some submodel with $q_1 < q$ components. In the time domain, this hypothesis implies, for a partition of the vector of inputs into q_1 and q_2 components ($q_1 + q_2 = q$), say, $\mathbf{x}_t = (\mathbf{x}'_{t1}, \mathbf{x}'_{t2})'$, and the similarly partitioned vector of regression functions $\boldsymbol{\beta}_t = (\boldsymbol{\beta}'_{1t}, \boldsymbol{\beta}'_{2t})'$, we would be interested in testing whether $\boldsymbol{\beta}_{2t} = \mathbf{0}$ in the partitioned regression model

$$y_t = \sum_{r=-\infty}^{\infty} \boldsymbol{\beta}'_{1r} \mathbf{x}_{t-r,1} + \sum_{r=-\infty}^{\infty} \boldsymbol{\beta}'_{2r} \mathbf{x}_{t-r,2} + v_t. \quad (7.29)$$

Rewriting the regression model (7.29) in the frequency domain in a form that is similar to (7.28) establishes that, under the partitions of the spectral matrix into its $q_i \times q_j$ ($i, j = 1, 2$) submatrices, say,

$$\widehat{f}_{xx}(\omega) = \begin{pmatrix} \widehat{f}_{11}(\omega) & \widehat{f}_{12}(\omega) \\ \widehat{f}_{21}(\omega) & \widehat{f}_{22}(\omega) \end{pmatrix}, \quad (7.30)$$

and the cross-spectral vector into its $q_i \times 1$ ($i = 1, 2$) subvectors,

$$\widehat{f}_{xy}(\omega) = \begin{pmatrix} \widehat{f}_{1y}(\omega) \\ \widehat{f}_{2y}(\omega) \end{pmatrix}, \quad (7.31)$$

we may test the hypothesis $\beta_{2t} = \mathbf{0}$ at frequency ω by comparing the estimated residual power

$$\widehat{f}_{y \cdot x}(\omega) = \widehat{f}_{yy}(\omega) - \widehat{f}_{xy}^*(\omega) \widehat{f}_{xx}^{-1}(\omega) \widehat{f}_{xy}(\omega) \quad (7.32)$$

under the full model with that under the reduced model, given by

$$\widehat{f}_{y \cdot 1}(\omega) = \widehat{f}_{yy}(\omega) - \widehat{f}_{1y}^*(\omega) \widehat{f}_{11}^{-1}(\omega) \widehat{f}_{1y}(\omega). \quad (7.33)$$

The power due to regression can be written as

$$\text{SSR}(\omega) = L[\widehat{f}_{y \cdot 1}(\omega) - \widehat{f}_{y \cdot x}(\omega)], \quad (7.34)$$

with the usual error power given by

$$\text{SSE}(\omega) = L\widehat{f}_{y \cdot x}(\omega). \quad (7.35)$$

The test of no regression proceeds using the F -statistic

$$F_{2q_2, 2(L-q)} = \frac{(L-q) \text{SSR}(\omega)}{q_2 \text{SSE}(\omega)}. \quad (7.36)$$

The distribution of this F -statistic with $2q_2$ numerator degrees of freedom and $2(L-q)$ denominator degrees of freedom follows from an argument paralleling that given in Chapter 4 for the case of a single input. The test results can be summarized in an Analysis of Power (ANOPOW) table that parallels the usual analysis of variance (ANOVA) table. **Table 7.1** shows the components of power for testing $\beta_{2t} = \mathbf{0}$ at a particular frequency ω . The ratio of the two components divided by their respective degrees of freedom just yields the F -statistic (7.36) used for testing whether the q_2 add significantly to the predictive power of the regression on the q_1 series.

Example 7.1 Predicting Lake Shasta Inflow

We illustrate some of the preceding ideas by considering the problem of predicting the transformed (logged) inflow series shown in **Figure 7.3** from some combination of the inputs. First, look for the best single input predictor using the squared coherence function (7.26). The results, exhibited in **Figure 7.4(a)-(e)**, show transformed (square root) precipitation produces the most consistently high squared coherence values at all frequencies ($L = 25$), with the seasonal period contributing most significantly. Other inputs, with the exception of wind speed, also appear to be plausible contributors. **Figure 7.4(a)-(e)** shows a .001 threshold corresponding to the F -statistic, separately, for each possible predictor of inflow.

Next, we focus on the analysis with two predictor series, temperature and transformed precipitation. The additional contribution of temperature to the model




Fig. 7.4. Squared coherency between Lake Shasta inflow and (a) temperature; (b) dew point; (c) cloud cover; (d) wind speed; (e) precipitation. The multiple coherency between inflow and temperature – precipitation jointly is displayed in (f). In each case, the .001 threshold is exhibited as a horizontal line.

seems somewhat marginal because the multiple coherence (7.26), shown in the top panel of Figure 7.4(f) seems only slightly better than the univariate coherence with precipitation shown in Figure 7.4(e). It is, however, instructive to produce the multiple regression functions, using (7.25) to see if a simple model for inflow exists that would involve some regression combination of inputs temperature and precipitation that would be useful for predicting inflow to Shasta Lake. The top of Figure 7.5 shows the partial F -statistic, (7.36), for testing if temperature is predictive of inflow when precipitation is in the model. In addition, threshold values corresponding to a false discovery rate (FDR) of .001 (see Benjamini & Hochberg, 1995) and the corresponding null F quantile are displayed in that figure.

Although the contribution of temperature is marginal, it is instructive to produce the multiple regression functions, using (7.25), to see if a simple model for inflow exists that would involve some regression combination of inputs temperature and precipitation that would be useful for predicting inflow to Lake Shasta. With this in mind, denoting the possible inputs P_t for transformed precipitation and T_t for transformed temperature, the regression function has been plotted in the lower two




Fig. 7.5. Partial F -statistics [top] for testing whether temperature adds to the ability to predict Lake Shasta inflow when precipitation is included in the model. The dashed line indicates the .001 FDR level and the solid line represents the corresponding quantile of the null F distribution. Multiple impulse response functions for the regression relations of temperature [middle] and precipitation [bottom].

panels of **Figure 7.5** using a value of $M = 100$ for each of the two inputs. In that figure, the time index runs over both positive and negative values and are centered at time $t = 0$. Hence, the relation with temperature seems to be instantaneous and positive and an exponentially decaying relation to precipitation exists that has been noticed previously in the analysis in **Problem 4.34**. The plots suggest a transfer function model of the general form fitted to the Recruitment and SOI series in **Example 5.7**. We might propose fitting the inflow output, say, I_t , using the model

$$I_t = \alpha_0 + \frac{\delta_0}{(1 - \omega_1 B)} P_t + \alpha_2 T_t + \eta_t,$$

which is the transfer function model, without the temperature component, considered in that section.

The R code for this example is as follows.

```
attach(climhyd)
plot.ts(climhyd) # figure 7.3
Y = climhyd # Y holds the transformed series
```

```

Y[,6] = log(Y[,6]) # log inflow
Y[,5] = sqrt(Y[,5]) # sqrt precipitation
L = 25; M = 100; alpha = .001; fdr = .001
nq = 2 # number of inputs (Temp and Precip)
# Spectral Matrix
Yspec = mvspec(Y, spans=L, kernel="daniell", detrend=TRUE, demean=FALSE,
  taper=.1)
n = Yspec$n.used # effective sample size
Fr = Yspec$freq # fundamental freqs
n.freq = length(Fr) # number of frequencies
Yspec$bandwidth*sqrt(12) # = 0.050 - the bandwidth
# Coherencies (see sec 4.7 also)
Fq = qf(1-alpha, 2, L-2); cn = Fq/(L-1+Fq)
plt.name = c("(a)","(b)","(c)","(d)","(e)","(f)")
dev.new(); par(mfrow=c(2,3), cex.lab=1.2)
# The coherencies are listed as 1,2,...,15=choose(6,2)
for (i in 1:15){
  plot(Fr, Yspec$coh[,i], type="l", ylab="Sq Coherence", xlab="Frequency",
    ylim=c(0,1), main=c("Inflow with", names(climhyd[i-10])))
  abline(h = cn); text(.45,.98, plt.name[i-10], cex=1.2) }
# Multiple Coherency
coh.15 = stoch.reg(Y, cols.full = c(1,5), cols.red = NULL, alpha, L, M,
  plot.which = "coh")
text(.45 ,.98, plt.name[6], cex=1.2)
title(main = c("Inflow with", "Temp and Precip"))
# Partial F (called eF; avoid use of F alone)
numer.df = 2*nq; denom.df = Yspec$df-2*nq
dev.new()
par(mfrow=c(3,1), mar=c(3,3,2,1)+.5, mgp = c(1.5,0.4,0), cex.lab=1.2)
out.15 = stoch.reg(Y, cols.full = c(1,5), cols.red = 5, alpha, L, M,
  plot.which = "F.stat")
eF = out.15$eF
pvals = pf(eF, numer.df, denom.df, lower.tail = FALSE)
pID = FDR(pvals, fdr); abline(h=c(eF[pID]), lty=2)
title(main = "Partial F Statistic")
# Regression Coefficients
S = seq(from = -M/2+1, to = M/2 - 1, length = M-1)
plot(S, coh.15$Betahat[,1], type = "h", xlab = "", ylab = names(climhyd[1]),
  ylim = c(-.025, .055), lwd=2)
abline(h=0); title(main = "Impulse Response Functions")
plot(S, coh.15$Betahat[,2], type = "h", xlab = "Index", ylab =
  names(climhyd[5]), ylim = c(-.015, .055), lwd=2)
abline(h=0)

```

7.4 Regression with Deterministic Inputs

The previous section considered the case in which the input and output series were jointly stationary, but there are many circumstances in which we might want to assume that the input functions are fixed and have a known functional form. This happens in the analysis of data from designed experiments. For example, we may want to take a collection of earthquakes and explosions such as are shown in [Figure 7.2](#) and test whether the mean functions are the same for either the P or S components or, perhaps, for them jointly. In certain other signal detection problems using arrays, the inputs

are used as dummy variables to express lags corresponding to the arrival times of the signal at various elements, under a model corresponding to that of a plane wave from a fixed source propagating across the array. In [Figure 7.1](#), we plotted the mean responses of the cortex as a function of various underlying design configurations corresponding to various stimuli applied to awake and mildly anesthetized subjects.

It is necessary to introduce a replicated version of the underlying model to handle even the univariate situation, and we replace (7.8) by

$$y_{jt} = \sum_{r=-\infty}^{\infty} \beta'_r z_{j,t-r} + v_{jt} \quad (7.37)$$

for $j = 1, 2, \dots, N$ series, where we assume the vector of known deterministic inputs, $\mathbf{z}_{jt} = (z_{jt1}, \dots, z_{jtq})'$, satisfies

$$\sum_{t=-\infty}^{\infty} |t| |z_{jtk}| < \infty$$

for $j = 1, \dots, N$ replicates of an underlying process involving $k = 1, \dots, q$ regression functions. The model can also be treated under the assumption that the deterministic function satisfy Grenander's conditions, as in Hannan (1970), but we do not need those conditions here and simply follow the approach in Shumway (1983, 1988).

It will sometimes be convenient in what follows to represent the model in matrix notation, writing (7.37) as

$$\mathbf{y}_t = \sum_{r=-\infty}^{\infty} \mathbf{z}_{t-r} \boldsymbol{\beta}_r + \mathbf{v}_t, \quad (7.38)$$

where $\mathbf{z}_t = (z_{1t}, \dots, z_{Nt})'$ are the $N \times q$ matrices of independent inputs and $\mathbf{y}_t$ and $\mathbf{v}_t$ are the $N \times 1$ output and error vectors. The error vector $\mathbf{v}_t = (v_{1t}, \dots, v_{Nt})'$ is assumed to be a multivariate, zero-mean, stationary, normal process with spectral matrix $f_v(\omega)I_N$ that is proportional to the $N \times N$ identity matrix. That is, we assume the error series v_{jt} are independently and identically distributed with spectral densities $f_v(\omega)$.

Example 7.2 An Infrasonic Signal from a Nuclear Explosion

Often, we will observe a common signal, say, β_t on an array of sensors, with the response at the j -th sensor denoted by y_{jt} , $j = 1, \dots, N$. For example, [Figure 7.6](#) shows an infrasonic or low-frequency acoustic signal from a nuclear explosion, as observed on a small triangular array of $N = 3$ acoustic sensors. These signals appear at slightly different times. Because of the way signals propagate, a plane wave signal of this kind, from a given source, traveling at a given velocity, will arrive at elements in the array at predictable time delays. In the case of the infrasonic signal in [Figure 7.6](#), the delays were approximated by computing the cross-correlation between elements and simply reading off the time delay corresponding to the maximum. For a detailed discussion of the statistical analysis of array signals, see Shumway et al. (1999).




Fig. 7.6. Three series for a nuclear explosion detonated 25 km south of Christmas Island and the delayed average or beam. The time scale is 10 points per second.

A simple additive signal-plus-noise model of the form

$$y_{jt} = \beta_{t-\tau_j} + v_{jt} \quad (7.39)$$

can be assumed, where $\tau_j, j = 1, 2, \dots, N$ are the time delays that determine the start point of the signal at each element of the array. The model (7.39) is written in the form (7.37) by letting $z_{jt} = \delta_{t-\tau_j}$, where $\delta_t = 1$ when $t = 0$ and is zero otherwise. In this case, we are interested in both the problem of detecting the presence of the signal and in estimating its waveform β_t . In this case, a plausible estimator of the waveform would be the unbiased beam, say,

$$\hat{\beta}_t = \frac{\sum_{j=1}^N y_{j,t+\tau_j}}{N}, \quad (7.40)$$

where time delays in this case were measured as $\tau_1 = 17$, $\tau_2 = 0$, and $\tau_3 = -22$ from the cross-correlation function. The bottom panel of Figure 7.6 shows the computed beam in this case, and the noise in the individual channels has been reduced and

the essential characteristics of the common signal are retained in the average. The R code for this example is

```
attach(beamd)
tau = rep(0,3)
u = ccf(sensor1, sensor2, plot=FALSE)
tau[1] = u$lag[which.max(u$acf)] # 17
u = ccf(sensor3, sensor2, plot=FALSE)
tau[3] = u$lag[which.max(u$acf)] # -22
Y = ts.union(lag(sensor1,tau[1]), lag(sensor2, tau[2]), lag(sensor3, tau[3]))
beam = rowMeans(Y)
par(mfrow=c(4,1), mar=c(0,5.1,0,5.1), oma=c(6,0,5,0))
plot.ts(sensor1, xaxt="no")
title(main="Infrasonic Signals and Beam", outer=TRUE)
plot.ts(sensor2, xaxt="no"); plot.ts(sensor3, xaxt="no")
plot.ts(beam); title(xlab="Time", outer=TRUE)
```

The above discussion and example serve to motivate a more detailed look at the estimation and detection problems in the case in which the input series $\mathbf{z}_{jt}$ are fixed and known. We consider the modifications needed for this case in the following sections.

ESTIMATION OF THE REGRESSION RELATION

Because the regression model (7.37) involves fixed functions, we may parallel the usual approach using the Gauss–Markov theorem to search for linear-filtered estimators of the form

$$\hat{\boldsymbol{\beta}}_t = \sum_{j=1}^N \sum_{r=-\infty}^{\infty} \mathbf{h}_{jr} y_{j,t-r}, \quad (7.41)$$

where $\mathbf{h}_{jt} = (h_{jt1} \dots, h_{jtq})'$ is a vector of filter coefficients, determined so the estimators are unbiased and have minimum variance. The equivalent matrix form is

$$\hat{\boldsymbol{\beta}}_t = \sum_{r=-\infty}^{\infty} \mathbf{h}_r \mathbf{y}_{t-r}, \quad (7.42)$$

where $\mathbf{h}_t = (\mathbf{h}_{1t}, \dots, \mathbf{h}_{Nt})$ is a $q \times N$ matrix of filter functions. The matrix form resembles the usual classical regression case and is more convenient for extending the Gauss–Markov Theorem to lagged regression. The unbiased condition is considered in **Problem 7.6**. It can be shown (see Shumway and Dean, 1968) that $\mathbf{h}_{js}$ can be taken as the Fourier transform of

$$\mathbf{H}_j(\omega) = S_z^{-1}(\omega) \overline{\mathbf{Z}_j(\omega)}, \quad (7.43)$$

where

$$\mathbf{Z}_j(\omega) = \sum_{t=-\infty}^{\infty} \mathbf{z}_{jt} e^{-2\pi i \omega t} \quad (7.44)$$

is the infinite Fourier transform of $\mathbf{z}_{jt}$. The matrix

$$S_z(\omega) = \sum_{j=1}^N \overline{\mathbf{Z}_j(\omega)} \mathbf{Z}'_j(\omega) \quad (7.45)$$

can be written in the form

$$S_z(\omega) = \mathbf{Z}^*(\omega) \mathbf{Z}(\omega), \quad (7.46)$$

where the $N \times q$ matrix $\mathbf{Z}(\omega)$ is defined by $\mathbf{Z}(\omega) = (\mathbf{Z}_1(\omega), \dots, \mathbf{Z}_N(\omega))'$. In matrix notation, the Fourier transform of the optimal filter becomes

$$\mathbf{H}(\omega) = S_z^{-1}(\omega) \mathbf{Z}^*(\omega), \quad (7.47)$$

where $\mathbf{H}(\omega) = (\mathbf{H}_1(\omega), \dots, \mathbf{H}_N(\omega))$ is the $q \times N$ matrix of frequency response functions. The optimal filter then becomes the Fourier transform

$$h_t = \int_{-1/2}^{1/2} \mathbf{H}(\omega) e^{2\pi i \omega t} d\omega. \quad (7.48)$$

If the transform is not tractable to compute, an approximation analogous to (7.25) may be used.

Example 7.3 Estimation of the Infrasonic Signal in Example 7.2

We consider the problem of producing a best linearly filtered unbiased estimator for the infrasonic signal in Example 7.2. In this case, $q = 1$ and (7.44) becomes

$$\mathbf{Z}_j(\omega) = \sum_{t=-\infty}^{\infty} \delta_{t-\tau_j} e^{-2\pi i \omega t} = e^{-2\pi i \omega \tau_j}$$

and $S_z(\omega) = N$. Hence, we have

$$H_j(\omega) = \frac{1}{N} e^{2\pi i \omega \tau_j}.$$

Using (7.48), we obtain $h_{jt} = \frac{1}{N} \delta(t + \tau_j)$. Substituting in (7.41), we obtain the best linear unbiased estimator as the beam, computed as in (7.40).

TESTS OF HYPOTHESES

We consider first testing the hypothesis that the complete vector $\boldsymbol{\beta}_t$ is zero, i.e., that the vector signal is absent. We develop a test at each frequency ω by taking single adjacent frequencies of the form $\omega_k = k/n$, as in the initial section. We may approximate the DFT of the observed vector in the model (7.37) using a representation of the form

$$Y_j(\omega_k) = \mathbf{B}'(\omega_k) \mathbf{Z}_j(\omega_k) + V_j(\omega_k) \quad (7.49)$$

for $j = 1, \dots, N$, where the error terms will be uncorrelated with common variance $f(\omega_k)$, the spectral density of the error term. The independent variables $\mathbf{Z}_j(\omega_k)$ can either be the infinite Fourier transform, or they can be approximated by the DFT.

Hence, we can obtain the matrix version of a complex regression model, written in the form

$$\mathbf{Y}(\omega_k) = Z(\omega_k)\mathbf{B}(\omega_k) + \mathbf{V}(\omega_k), \quad (7.50)$$

where the $N \times q$ matrix $Z(\omega_k)$ has been defined previously below (7.46) and $\mathbf{Y}(\omega_k)$ and $\mathbf{V}(\omega_k)$ are $N \times 1$ vectors with the error vector $\mathbf{V}(\omega_k)$ having mean zero, with covariance matrix $f(\omega_k)I_N$. The usual regression arguments show that the maximum likelihood estimator for the regression coefficient will be

$$\widehat{\mathbf{B}}(\omega_k) = S_z^{-1}(\omega_k)\mathbf{s}_{zy}(\omega_k), \quad (7.51)$$

where $S_z(\omega_k)$ is given by (7.46) and

$$\mathbf{s}_{zy}(\omega_k) = Z^*(\omega_k)\mathbf{Y}(\omega_k) = \sum_{j=1}^N \overline{Z_j(\omega_k)} Y_j(\omega_k). \quad (7.52)$$

Also, the maximum likelihood estimator for the error spectral matrix is proportional to

$$\begin{aligned} s_{y \cdot z}^2(\omega_k) &= \sum_{j=1}^N |Y_j(\omega_k) - \widehat{\mathbf{B}}(\omega_k)' \mathbf{Z}_j(\omega_k)|^2 \\ &= \mathbf{Y}^*(\omega_k)\mathbf{Y}(\omega_k) - \mathbf{Y}^*(\omega_k)Z(\omega_k)[Z^*(\omega_k)Z(\omega_k)]^{-1}Z^*(\omega_k)\mathbf{Y}(\omega_k) \\ &= s_y^2(\omega_k) - \mathbf{s}_{zy}^*(\omega_k)S_z^{-1}(\omega_k)\mathbf{s}_{zy}(\omega_k), \end{aligned} \quad (7.53)$$

where

$$s_y^2(\omega_k) = \sum_{j=1}^N |Y_j(\omega_k)|^2. \quad (7.54)$$

Under the null hypothesis that the regression coefficient $\mathbf{B}(\omega_k) = \mathbf{0}$, the estimator for the error power is just $s_y^2(\omega_k)$. If smoothing is needed, we may replace the (7.53) and (7.54) by smoothed components over the frequencies $\omega_k + \ell/n$, for $\ell = -m, \dots, m$ and $L = 2m + 1$, close to ω . In that case, we obtain the regression and error spectral components as

$$\text{SSR}(\omega) = \sum_{\ell=-m}^m \mathbf{s}_{zy}^*(\omega_k + \ell/n) S_z^{-1}(\omega_k + \ell/n) \mathbf{s}_{zy}(\omega_k + \ell/n) \quad (7.55)$$

and

$$\text{SSE}(\omega) = \sum_{\ell=-m}^m s_{y \cdot z}^2(\omega_k + \ell/n). \quad (7.56)$$

The F -statistic for testing no regression relation is

$$F_{2Lq, 2L(N-q)} = \frac{N-q}{q} \frac{\text{SSR}(\omega)}{\text{SSE}(\omega)}. \quad (7.57)$$

The analysis of power pertaining to this situation appears in Table 7.2.

Table 7.2. Analysis of Power (ANOPOW) for Testing No Contribution from the Independent Series at Frequency ω in the Fixed Input Case

Source	Power	Degrees of Freedom
Regression	$SSR(\omega)$ (7.55)	$2Lq$
Error	$SSE(\omega)$ (7.56)	$2L(N - q)$
Total	$SST(\omega)$	$2LN$

In the fixed regression case, the partitioned hypothesis that is the analog of $\beta_{2t} = 0$ in (7.27) with $\mathbf{x}_{t1}, \mathbf{x}_{t2}$ replaced by $\mathbf{z}_{t1}, \mathbf{z}_{t2}$. Here, we partition $S_z(\omega)$ into $q_i \times q_j$ ($i, j = 1, 2$) submatrices, say,

$$S_z(\omega_k) = \begin{pmatrix} S_{11}(\omega_k) & S_{12}(\omega_k) \\ S_{21}(\omega_k) & S_{22}(\omega_k) \end{pmatrix}, \quad (7.58)$$

and the cross-spectral vector into its $q_i \times 1$, for $i = 1, 2$, subvectors

$$\mathbf{s}_{zy}(\omega_k) = \begin{pmatrix} \mathbf{s}_{1y}(\omega_k) \\ \mathbf{s}_{2y}(\omega_k) \end{pmatrix}. \quad (7.59)$$

Here, we test the hypothesis $\beta_{2t} = \mathbf{0}$ at frequency ω by comparing the residual power (7.53) under the full model with the residual power under the reduced model, given by

$$s_{y \cdot 1}^2(\omega_k) = s_y^2(\omega_k) - \mathbf{s}_{1y}^*(\omega_k) S_{11}^{-1}(\omega_k) \mathbf{s}_{1y}(\omega_k). \quad (7.60)$$

Again, it is desirable to add over adjacent frequencies with roughly comparable spectra so the regression and error power components can be taken as

$$SSR(\omega) = \sum_{\ell=-m}^m \left[s_{y \cdot 1}^2(\omega_k + \ell/n) - s_{y \cdot z}^2(\omega_k + \ell/n) \right] \quad (7.61)$$

and

$$SSE(\omega) = \sum_{\ell=-m}^m s_{y \cdot z}^2(\omega_k + \ell/n). \quad (7.62)$$

The information can again be summarized as in Table 7.3, where the ratio of mean power regression and error components leads to the F -statistic

$$F_{2Lq_2, 2L(N-q)} = \frac{(N - q) SSR(\omega)}{q_2 SSE(\omega)}. \quad (7.63)$$

We illustrate the analysis of power procedure using the infrasonic signal detection procedure of Example 7.2.

Table 7.3. Analysis of Power (ANOPOW) for Testing No Contribution from the Last q_2 Inputs in the Fixed Input Case

Source	Power	Degrees of Freedom
Regression	$SSR(\omega)$ (7.61)	$2Lq_2$
Error	$SSE(\omega)$ (7.62)	$2L(N - q)$
Total	$SST(\omega)$	$2L(N - q_1)$

Example 7.4 Detecting the Infrasonic Signal Using ANOPOW

We consider the problem of detecting the common signal for the three infrasonic series observing the common signal, as shown in Figure 7.4. The presence of the signal is obvious in the waveforms shown, so the test here mainly confirms the statistical significance and isolates the frequencies containing the strongest signal components. Each series contained $n = 2048$ points, sampled at 10 points per second. We use the model in (7.39) so $Z_j(\omega) = e^{-2\pi i \omega \tau_j}$ and $S_z(\omega) = N$ as in Example 7.3, with $s_{zy}(\omega_k)$ given as

$$s_{zy}(\omega_k) = \sum_{j=1}^N e^{2\pi i \omega \tau_j} Y_j(\omega_k),$$

using (7.45) and (7.52). The above expression can be interpreted as being proportional to the weighted mean or beam, computed in frequency, and we introduce the notation

$$B_w(\omega_k) = \frac{1}{N} \sum_{j=1}^N e^{2\pi i \omega \tau_j} Y_j(\omega_k) \quad (7.64)$$

for that term. Substituting for the power components in Table 7.3 yields

$$\mathbf{s}_{zy}^*(\omega_k) S_z^{-1}(\omega_k) \mathbf{s}_{zy}(\omega_k) = N |B_w(\omega_k)|^2$$

and

$$s_{y \cdot z}^2(\omega_k) = \sum_{j=1}^N |Y_j(\omega_k) - B_w(\omega_k)|^2 = \sum_{j=1}^N |Y_j(\omega_k)|^2 - N |B_w(\omega_k)|^2$$

for the regression signal and error components, respectively. Because only three elements in the array and a reasonable number of points in time exist, it seems advisable to employ some smoothing over frequency to obtain additional degrees of freedom. In this case, $L = 9$, yielding $2(9) = 18$ and $2(9)(3 - 1) = 36$ degrees of freedom for the numerator and denominator of the F -statistic (7.57). The top of Figure 7.7 shows the analysis of power components due to error and the total power. The power is maximum at about .002 cycles per point or about .02 cycles per second. The F -statistic is compared with the .001 FDR and the corresponding null F significance in the bottom panel and has the strongest detection at about .02 cycles




Fig. 7.7. Analysis of power for infrasound array on a log scale (top panel) with $SST(\omega)$ shown as a solid line and $SSE(\omega)$ as a dashed line. The F -statistics (bottom panel) showing detections with the dashed line based on an FDR level of .001 and the solid line corresponding null F quantile.

per second. Little power of consequence appears to exist elsewhere, however, there is some marginally significant signal power near the .5 cycles per second frequency band.

The R code for this example is as follows.

```
attach(beamd)
L = 9; fdr = .001; N = 3
Y = cbind(beamd, beam=rowMeans(beamd))
n = nextn(nrow(Y))
Y.fft = mvfft(as.ts(Y))/sqrt(n)
Df = Y.fft[,1:3] # fft of the data
Bf = Y.fft[,4]   # beam fft
ssr = N*Re(Bf*Conj(Bf)) # raw signal spectrum
sse = Re(rowSums(Df*Conj(Df))) - ssr # raw error spectrum
# Smooth
SSE = filter(sse, sides=2, filter=rep(1/L,L), circular=TRUE)
SSR = filter(ssr, sides=2, filter=rep(1/L,L), circular=TRUE)
SST = SSE + SSR
par(mfrow=c(2,1), mar=c(4,4,2,1)+.1)
Fr = 0:(n-1)/n; nFr = 1:200 # number of freqs to plot
plot(Fr[nFr], SST[nFr], type="l", ylab="log Power", xlab="", main="Sum of
  Squares", log="y")
lines(Fr[nFr], SSE[nFr], type="l", lty=2)
eF = (N-1)*SSR/SSE; df1 = 2*L; df2 = 2*L*(N-1)
pvals = pf(eF, df1, df2, lower=FALSE) # p values for FDR
pID = FDR(pvals, fdr); Fq = qf(1-fdr, df1, df2)
```

```
plot(Fr[nFr], eF[nFr], type="l", ylab="F-statistic", xlab="Frequency",
     main="F Statistic")
abline(h=c(Fq, eF[pID]), lty=1:2)
```

Although there are examples of detecting multiple regression functions of the general type considered above (see, for example, Shumway, 1983), we do not consider additional examples of partitioning in the fixed input case here. The reason is that several examples exist in the section on designed experiments that illustrate the partitioned approach.

7.5 Random Coefficient Regression

The lagged regression models considered so far have assumed the input process is either stochastic or fixed and the components of the vector of regression function β_t are fixed and unknown parameters to be estimated. There are many cases in time series analysis in which it is more natural to regard the regression vector as an unknown stochastic signal. For example, we have studied the state-space model in Chapter 6, where the state equation can be considered as involving a random parameter vector that is essentially a multivariate autoregressive process. In §4.10, we considered estimating the univariate regression function β_t as a signal extraction problem.

In this section, we consider a random coefficient regression model of (7.38) in the equivalent form

$$\mathbf{y}_t = \sum_{r=-\infty}^{\infty} z_{t-r} \beta_r + \mathbf{v}_t, \quad (7.65)$$

where $\mathbf{y}_t = (y_{1t}, \dots, y_{Nt})'$ is the $N \times 1$ response vector and $\mathbf{z}_t = (z_{1t}, \dots, z_{Nt})'$ are the $N \times q$ matrices containing the fixed input processes. Here, the components of the $q \times 1$ regression vector β_t are zero-mean, uncorrelated, stationary series with common spectral matrix $f_\beta(\omega)I_q$ and the error series $\mathbf{v}_t$ have zero-means and spectral matrix $f_v(\omega)I_N$, where I_N is the $N \times N$ identity matrix. Then, defining the $N \times q$ matrix $Z(\omega) = (\mathbf{Z}_1(\omega), \mathbf{Z}_2(\omega), \dots, \mathbf{Z}_N(\omega))'$ of Fourier transforms of $\mathbf{z}_t$, as in (7.44), it is easy to show the spectral matrix of the response vector $\mathbf{y}_t$ is given by

$$f_y(\omega) = f_\beta(\omega)Z(\omega)Z^*(\omega) + f_v(\omega)I_N. \quad (7.66)$$

The regression model with a stochastic stationary signal component is a general version of the simple additive noise model

$$y_t = \beta_t + v_t,$$

considered by Wiener (1949) and Kolmogorov (1941), who derived the minimum mean squared error estimators for β_t , as in §4.10. The more general multivariate version (7.65) represents the series as a convolution of the signal vector β_t and a known set of vector input series contained in the matrix $\mathbf{z}_t$. Restricting the the covariance matrices of signal and noise to diagonal form is consistent with what is done in

statistics using random effects models, which we consider here in a later section. The problem of estimating the regression function β_t is often called *deconvolution* in the engineering and geophysical literature.

ESTIMATION OF THE REGRESSION RELATION

The regression function β_t can be estimated by a general filter of the form (7.42), where we write that estimator in matrix form

$$\hat{\beta}_t = \sum_{r=-\infty}^{\infty} h_r y_{t-r}, \quad (7.67)$$

where $h_t = (h_{1t}, \dots, h_{Nt})$, and apply the orthogonality principle, as in §3.9. A generalization of the argument in that section (see Problem 7.7) leads to the estimator

$$H(\omega) = [S_z(\omega) + \theta(\omega)I_q]^{-1}Z^*(\omega) \quad (7.68)$$

for the Fourier transform of the minimum mean-squared error filter, where the parameter

$$\theta(\omega) = \frac{f_v(\omega)}{f_\beta(\omega)} \quad (7.69)$$

is the inverse of the signal-to-noise ratio. It is clear from the frequency domain version of the linear model (7.50), the comparable version of the estimator (7.51) can be written as

$$\hat{\mathbf{B}}(\omega) = [S_z(\omega) + \theta(\omega)I_q]^{-1}\mathbf{s}_{zy}(\omega). \quad (7.70)$$

This version exhibits the estimator in the stochastic regressor case as the usual estimator, with a ridge correction, $\theta(\omega)$, that is proportional to the inverse of the signal-to-noise ratio.

The mean-squared covariance of the estimator is shown to be

$$E[(\hat{\mathbf{B}} - \mathbf{B})(\hat{\mathbf{B}} - \mathbf{B})^*] = f_v(\omega)[S_z(\omega) + \theta(\omega)I_q]^{-1}, \quad (7.71)$$

which again exhibits the close connection between this case and the variance of the estimator (7.51), which can be shown to be $f_v(\omega)S_z^{-1}(\omega)$.

Example 7.5 Estimating the Random Infrasonic Signal

In Example 7.4, we have already determined the components needed in (7.68) and (7.69) to obtain the estimators for the random signal. The Fourier transform of the optimum filter at series j has the form

$$H_j(\omega) = \frac{e^{2\pi i \omega \tau_j}}{N + \theta(\omega)} \quad (7.72)$$

with the mean-squared error given by $f_v(\omega)/[N + \theta(\omega)]$ from (7.71). The net effect of applying the filters will be the same as filtering the beam with the frequency response function

$$H_0(\omega) = \frac{N}{N + \theta(\omega)} = \frac{N f_\beta(\omega)}{f_v(\omega) + N f_\beta(\omega)}, \quad (7.73)$$

where the last form is more convenient in cases in which portions of the signal spectrum are essentially zero.

The optimal filters h_t have frequency response functions that depend on the signal spectrum $f_\beta(\omega)$ and noise spectrum $f_v(\omega)$, so we will need estimators for these parameters to apply the optimal filters. Sometimes, there will be values, suggested from experience, for the signal-to-noise ratio $1/\theta(\omega)$ as a function of frequency. The analogy between the model here and the usual variance components model in statistics, however, suggests we try an approach along those lines as in the next section.

DETECTION AND PARAMETER ESTIMATION

The analogy to the usual variance components situation suggests looking at the regression and error components of Table 7.2 under the stochastic signal assumptions. We consider the components of (7.55) and (7.56) at a single frequency ω_k . In order to estimate the spectral components $f_\beta(\omega)$ and $f_v(\omega)$, we reconsider the linear model (7.50) under the assumption that $\mathbf{B}(\omega_k)$ is a random process with spectral matrix $f_\beta(\omega_k)I_q$. Then, the spectral matrix of the observed process is (7.66), evaluated at frequency ω_k .

Consider first the component of the regression power, defined as

$$\begin{aligned} SSR(\omega_k) &= \mathbf{s}_{zy}^*(\omega_k) S_z^{-1}(\omega_k) \mathbf{s}_{zy}(\omega_k) \\ &= \mathbf{Y}^*(\omega_k) \mathbf{Z}(\omega_k) S_z^{-1}(\omega_k) \mathbf{Z}^*(\omega_k) \mathbf{Y}(\omega_k). \end{aligned}$$

A computation shows

$$E[SSR(\omega_k)] = f_\beta(\omega_k) \text{tr}\{S_z(\omega_k)\} + q f_v(\omega_k),$$

where tr denotes the trace of a matrix. If we can find a set of frequencies of the form $\omega_k + \ell/n$, where the spectra and the Fourier transforms $S_z(\omega_k + \ell/n) \approx S_z(\omega)$ are relatively constant, the expectation of the averaged values in (7.55) yields

$$E[SSR(\omega)] = L f_\beta(\omega) \text{tr}[S_z(\omega)] + L q f_v(\omega). \quad (7.74)$$

A similar computation establishes

$$E[SSE(\omega)] = L(N - q) f_v(\omega). \quad (7.75)$$

We may obtain an approximately unbiased estimator for the spectra $f_v(\omega)$ and $f_\beta(\omega)$ by replacing the expected power components by their values and solving (7.74) and (7.75).

Example 7.6 Estimating the Power Components and the Random Infrasonic Signal

In order to provide an optimum estimator for the infrasonic signal, we need to have estimators for the signal and noise spectra $f_\beta(\omega)$ and $f_v(\omega)$ for the case considered in [Example 7.5](#). The form of the filter is $H_0(\omega)$, given in (7.73), and with $q = 1$ and the matrix $S_z(\omega) = N$ at all frequencies in this example simplifies the computations considerably. We may estimate

$$\hat{f}_v(\omega) = \frac{SSE(\omega)}{L(N-1)} \quad (7.76)$$

and

$$\hat{f}_\beta(\omega) = (LN)^{-1} \left(SSR(\omega) - \frac{SSE(\omega)}{(N-1)} \right), \quad (7.77)$$

using (7.74) and (7.75) for this special case. Cases will exist in which (7.77) is negative and the estimated signal spectrum can be set to zero for those frequencies. The estimators can be substituted into the optimal filters to apply to the beam, say, $H_0(\omega)$ in (7.73), or to use in the filter applied to each level (7.72).

The analysis of variance estimators can be computed using the analysis of power given in [Figure 7.7](#), and the results of that computation and applying (7.76) and (7.77) are shown in [Figure 7.8\(a\)](#) for a bandwidth of $B = 9/2048$ cycles per point or about .04 cycles per second (Hz). Neither spectrum contains any significant power for frequencies greater than .1 Hz. As expected, the signal spectral estimator is substantial over a narrow band, and this leads to an estimated filter, with estimated frequency response function $\hat{H}_0(\omega)$, shown in [Figure 7.8\(b\)](#). The estimated optimal filter essentially deletes frequencies above .1 Hz and, subject to slight modification, differs little from a standard low-pass filter with that cutoff. Computing the time version with a cutoff at $M = 101$ points and using a taper leads to the estimated impulse response function $\hat{h}_0(t)$, as shown in [Figure 7.8\(c\)](#). Finally, we apply the optimal filter to the beam and get the filtered beam $\hat{\beta}_t$ shown in [Figure 7.4\(d\)](#). The analysis shows the primary signal as basically a low-frequency signal with primary power at about .05 Hz or, essentially, a wave with a 20-second period.

The R code for this example is as follows.

```
attach(beamd)
L = 9; M = 100; M2 = M/2; N = 3
Y = cbind(beamd, beam <- rowMeans(beamd))
n = nextn(nrow(Y)); n.freq=n/2
Y[,1:3] = Y[,1:3] - Y[,4]
Y.fft = mvfft(as.ts(Y))/sqrt(n)
Ef = Y.fft[,1:3] # fft of the error
Bf = Y.fft[,4]   # beam fft
ssr = N*Re(Bf*Conj(Bf)) # Raw Signal Spectrum
sse = Re(rowSums(Ef*Conj(Ef))) # Raw Error Spectrum
# Smooth
SSE = filter(sse, sides=2, filter=rep(1/L,L), circular=TRUE)
SSR = filter(ssr, sides=2, filter=rep(1/L,L), circular=TRUE)
# Estimate Signal and Noise Spectra
fv = SSE/(L*(N-1)) # Equation (7.77)
fb = (SSR-SSE/(N-1))/(L*N) # Equation (7.78)
```




Fig. 7.8. (a) Estimated signal (solid line) and noise (dashed line) spectra; (b) Estimated frequency and (c) corresponding impulse filter responses; (d) Raw beam (dashed line) with superimposed filtered beam (solid line).

```
fb[fb<0] = 0; H0 = N*fb/(fv+N*fb)
H0[ceiling(.04*n):n] = 0 # zero out H0 beyond frequency .04
# Extend components to make it a valid transform
H0 = c(H0[1:n.freq], rev(H0[2:(n.freq+1)]))
h0 = Re(fft(H0, inverse = TRUE)) # Impulse Response
h0 = c(rev(h0[2:(M2+1)]), h0[1:(M2+1)]) # center it
h1 = spec.taper(h0, p = .5); k1 = h1/sum(h1) # taper/normalize it
f.beam = filter(Y$beam, filter=k1, sides=2) # filter it
# Graphics
nFr = 1:50 # number of freqs displayed
Fr = (nFr-1)/n # frequencies
layout(matrix(c(1, 2, 4, 1, 3, 4), nc=2)); par(mar=c(4,4,2,1)+.1)
plot(10*Fr, fb[nFr], type="l", ylab="Power", xlab="Frequency (Hz)")
lines(10*Fr, fv[nFr], lty=2); text(.24, 5, "(a)", cex=1.2)
plot(10*Fr, H0[nFr], type="l", ylab="Frequency Response", xlab="Frequency
(Hz)")
text(.23, .84, "(b)", cex=1.2)
plot(-M2:M2, k1, type="l", ylab="Impulse Response", xlab="Index", lwd=1.5)
text(45, .022, "(c)", cex=1.2)
```

```
ts.plot(cbind(f.beam,beam), lty=1:2, ylab="beam")
text(2040, 2, "(d)", cex=1.2)
```

7.6 Analysis of Designed Experiments

An important special case (see Brillinger, 1973, 1980) of the regression model (7.49) occurs when the regression (7.38) is of the form

$$\mathbf{y}_t = \mathbf{z}\boldsymbol{\beta}_t + \mathbf{v}_t, \quad (7.78)$$

where $\mathbf{z} = (\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_N)'$ is a matrix that determines what is observed by the j -th series; i.e.,

$$y_{jt} = \mathbf{z}'_j \boldsymbol{\beta}_t + v_{jt}. \quad (7.79)$$

In this case, the the matrix $\mathbf{z}$ of independent variables is constant and we will have the frequency domain model.

$$\mathbf{Y}(\omega_k) = \mathbf{Z}\mathbf{B}(\omega_k) + \mathbf{V}(\omega_k) \quad (7.80)$$

corresponding to (7.50), where the matrix $\mathbf{Z}(\omega_k)$ was a function of frequency ω_k . The matrix is purely real, in this case, but the equations (7.51)-(7.57) can be applied with $\mathbf{Z}(\omega_k)$ replaced by the constant matrix $\mathbf{Z}$.

EQUALITY OF MEANS

A typical general problem that we encounter in analyzing real data is a simple equality of means test in which there might be a collection of time series y_{ijt} , $i = 1, \dots, I$, $j = 1, \dots, N_i$, belonging to I possible groups, with N_i series in group i . To test equality of means, we may write the regression model in the form

$$y_{ijt} = \mu_t + \alpha_{it} + v_{ijt}, \quad (7.81)$$

where μ_t denotes the overall mean and α_{it} denotes the effect of the i -th group at time t and we require that $\sum_i \alpha_{it} = 0$ for all t . In this case, the full model can be written in the general regression notation as

$$y_{ijt} = \mathbf{z}'_{ij} \boldsymbol{\beta}_t + v_{ijt}$$

where

$$\boldsymbol{\beta}_t = (\mu_t, \alpha_{1t}, \alpha_{2t}, \dots, \alpha_{I-1,t})'$$

denotes the regression vector, subject to the constraint. The reduced model becomes

$$y_{ijt} = \mu_t + v_{ijt} \quad (7.82)$$

under the assumption that the group means are equal. In the full model, there are I possible values for the $I \times 1$ design vectors $\mathbf{z}_{ij}$; the first component is always one for the mean, and the rest have a one in the i -th position for $i = 1, \dots, I - 1$ and zeros

elsewhere. The vectors for the last group take the value -1 for $i = 2, 3, \dots, I - 1$. Under the reduced model, each $\mathbf{z}_{ij}$ is a single column of ones. The rest of the analysis follows the approach summarized in (7.51)-(7.57). In this particular case, the power components in Table 7.3 (before smoothing) simplify to

$$SSR(\omega_k) = \sum_{i=1}^I \sum_{j=1}^{N_i} |Y_{i.}(\omega_k) - Y_{..}(\omega_k)|^2 \quad (7.83)$$

and

$$SSE(\omega_k) = \sum_{i=1}^I \sum_{j=1}^{N_i} |Y_{ij}(\omega_k) - Y_{i.}(\omega_k)|^2, \quad (7.84)$$

which are analogous to the usual sums of squares in analysis of variance. Note that a dot ($\cdot$) stands for a mean, taken over the appropriate subscript, so the regression power component $SSR(\omega_k)$ is basically the power in the residuals of the group means from the overall mean and the error power component $SSE(\omega_k)$ reflects the departures of the group means from the original data values. Smoothing each component over L frequencies leads to the usual F -statistic (7.63) with $2L(I - 1)$ and $2L(\sum_i N_i - I)$ degrees of freedom at each frequency ω of interest.

Example 7.7 Means Test for the fMRI Data

Figure 7.1 showed the mean responses of subjects to various levels of periodic stimulation while awake and while under anesthesia, as collected in a pain perception experiment of Antognini et al. (1997). Three types of periodic stimuli were presented to awake and anesthetized subjects, namely, brushing, heat, and shock. The periodicity was introduced by applying the stimuli, brushing, heat, and shocks in on-off sequences lasting 32 seconds each and the sampling rate was one point every two seconds. The blood oxygenation level (BOLD) signal intensity (Ogawa et al., 1990) was measured at nine locations in the brain. Areas of activation were determined using a technique first described by Bandettini et al. (1993). The specific locations of the brain where the signal was measured were Cortex 1: Primary Somatosensory, Contralateral, Cortex 2: Primary Somatosensory, Ipsilateral, Cortex 3: Secondary Somatosensory, Contralateral, Cortex 4: Secondary Somatosensory, Ipsilateral, Caudate, Thalamus 1: Contralateral, Thalamus 2: Ipsilateral, Cerebellum 1: Contralateral and Cerebellum 2: Ipsilateral. Figure 7.1 shows the mean response of subjects at Cortex 1 for each of the six treatment combinations, 1: Awake-Brush (5 subjects), 2: Awake-Heat (4 subjects), 3: Awake-Shock (5 subjects), 4: Low-Brush (3 subjects), 5: Low-Heat (5 subjects), and 6: Low-Shock (4 subjects). The objective of this first analysis is to test equality of these six group means, paying special attention to the 64-second period band (1/64 cycles per second) expected from the periodic driving stimuli. Because a test of equality is needed at each of the nine brain locations, we took $\alpha = .001$ to control for the overall error rate. Figure 7.9 shows F -statistics, computed from (7.63), with $L = 3$, and we see substantial signals for the four cortex locations and for the second cerebellum trace, but the effects are nonsignificant in the caudate and thalamus regions. Hence, we

will retain the four cortex locations and the second cerebellum location for further analysis.

The R code for this example is as follows.

```
n = 128          # length of series
n.freq = 1 + n/2 # number of frequencies
Fr = (0:(n.freq-1))/n # the frequencies
N = c(5,4,5,3,5,4) # number of series for each cell
n.subject = sum(N) # number of subjects (26)
n.trt = 6 # number of treatments
L = 3      # for smoothing
num.df = 2*L*(n.trt-1) # df for F test
den.df = 2*L*(n.subject-n.trt)
# Design Matrix (Z):
Z1 = outer(rep(1,N[1]), c(1,1,0,0,0,0))
Z2 = outer(rep(1,N[2]), c(1,0,1,0,0,0))
Z3 = outer(rep(1,N[3]), c(1,0,0,1,0,0))
Z4 = outer(rep(1,N[4]), c(1,0,0,0,1,0))
Z5 = outer(rep(1,N[5]), c(1,0,0,0,0,1))
Z6 = outer(rep(1,N[6]), c(1,-1,-1,-1,-1,-1))
Z = rbind(Z1, Z2, Z3, Z4, Z5, Z6)
ZZ = t(Z)%*%Z
SSEF <- rep(NA, n) -> SSER
HatF = Z%*%solve(ZZ, t(Z))
HatR = Z[,1]%*%t(Z[,1])/ZZ[1,1]
par(mfrow=c(3,3), mar=c(3.5,4,0,0), oma=c(0,0,2,2), mgp = c(1.6,.6,0))
loc.name = c("Cortex 1","Cortex 2","Cortex 3","Cortex 4","Caudate","Thalamus
1","Thalamus 2","Cerebellum 1","Cerebellum 2")
for(Loc in 1:9) {
  i = n.trt*(Loc-1)
  Y = cbind(fmri[[i+1]], fmri[[i+2]], fmri[[i+3]], fmri[[i+4]], fmri[[i+5]],
    fmri[[i+6]])
  Y = mvfft(spec.taper(Y, p=.5))/sqrt(n)
  Y = t(Y) # Y is now 26 x 128 FFTs
# Calculation of Error Spectra
for (k in 1:n) {
  SSY = Re(Conj(t(Y[,k])))%*%Y[,k])
  SSReg = Re(Conj(t(Y[,k])))%*%HatF%*%Y[,k])
  SSEF[k] = SSY - SSReg
  SSReg = Re(Conj(t(Y[,k])))%*%HatR%*%Y[,k])
  SSER[k] = SSY - SSReg }
# Smooth
sSSEF = filter(SSEF, rep(1/L, L), circular = TRUE)
sSSER = filter(SSER, rep(1/L, L), circular = TRUE)
eF = (den.df/num.df)*(sSSER-sSSEF)/sSSEF
plot(Fr, eF[1:n.freq], type="l", xlab="Frequency", ylab="F Statistic",
  ylim=c(0,7))
abline(h=qf(.999, num.df, den.df),lty=2)
text(.25, 6.5, loc.name[Loc], cex=1.2) }
```

AN ANALYSIS OF VARIANCE MODEL

The arrangement of treatments for the fMRI data in [Figure 7.1](#) suggests more information might be available than was obtained from the simple equality of means test. Separate effects caused by state of consciousness as well as the separate treatments




Fig. 7.9. Frequency-dependent equality of means tests for fMRI data at 9 brain locations. $L = 3$ and critical value $F_{.001}(30, 120) = 2.26$.

brush, heat, and shock might exist. The reduced signal present in the low shock mean suggests a possible interaction between the treatments and level of consciousness. The arrangement in the classical two-way table suggests looking at the analog of the two factor analysis of variance as a function of frequency. In this case, we would obtain a different version of the regression model (7.81) of the form

$$y_{ijkt} = \mu_t + \alpha_{it} + \beta_{jt} + \gamma_{ijt} + v_{ijkt} \quad (7.85)$$

for the k -th individual undergoing the i -th level of some factor A and the j -th level of some other factor B, $i = 1, \dots, I$, $j = 1, \dots, J$, $k = 1, \dots, n_{ij}$. The number of individuals in each cell can be different, as for the fMRI data in the next example. In the above model, we assume the response can be modeled as the sum of a mean, μ_t , a row effect (type of stimulus), α_{it} , a column effect (level of consciousness), β_{jt} and an interaction, γ_{ijt} , with the usual restrictions

Table 7.4. Rows of the Design Matrix $\mathbf{z}'_j$ for fMRI Data.

	Awake							Low Anesthesia						
Brush	1	1	0	1	1	0	(5)	1	1	0	-1	-1	0	(3)
Heat	1	0	1	1	0	1	(4)	1	0	1	-1	0	-1	(5)
Shock	1	-1	-1	1	-1	-1	(5)	1	-1	-1	-1	1	1	(4)

Number of Observations per Cell in Parentheses

$$\sum_i \alpha_{it} = \sum_j \beta_{jt} = \sum_i \gamma_{ijt} = \sum_j \gamma_{ijt} = 0$$

required for a full rank design matrix $\mathbf{Z}$ in the overall regression model (7.80). If the number of observations in each cell were the same, the usual simple analogous version of the power components (7.83) and (7.84) would exist for testing various hypotheses. In the case of (7.85), we are interested in testing hypotheses obtained by dropping one set of terms at a time out of (7.85), so an A factor (testing $\alpha_{it} = 0$), a B factor ($\beta_{jt} = 0$), and an interaction term ($\gamma_{ijt} = 0$) will appear as components in the analysis of power. Because of the unequal numbers of observations in each cell, we often put the model in the form of the regression model (7.78)–(7.80).

Example 7.8 Analysis of Power Tests for the fMRI Series

For the fMRI data given as the means in Figure 7.1, a model of the form (7.85) is plausible and will yield more detailed information than the simple equality of means test described earlier. The results of that test, shown in Figure 7.9, were that the means were different for the four cortex locations and for the second cerebellum location. We may examine these differences further by testing whether the mean differences are because of the nature of the stimulus or the consciousness level, or perhaps due to an interaction between the two factors. Unequal numbers of observations exist in the cells that contributed the means in Figure 7.1. For the regression vector,

$$(\mu_t, \alpha_{1t}, \alpha_{2t}, \beta_{1t}, \gamma_{11t}, \gamma_{21t})',$$

the rows of the design matrix are as specified in Table 7.4. Note the restrictions given above for the parameters.

The results of testing the three hypotheses are shown in Figure 7.10 for the four cortex locations and the cerebellum, the components that showed some significant differences in the means in Figure 7.9. Again, the regression power components were smoothed over $L = 3$ frequencies. Appealing to the ANOPOW results summarized in Table 7.3 for each of the subhypotheses, $q_2 = 1$ when the stimulus effect is dropped, and $q_2 = 2$ when either the consciousness effect or the interaction terms are dropped. Hence, $2Lq_2 = 6, 12$ for the two cases, with $N = \sum_{ij} n_{ij} = 26$ total observations. Here, the state of consciousness (Awake, Sedated) has the major effect at the signal frequency. The level of stimulus was less significant at the signal frequency. A significant interaction occurred, however, at the ipsilateral component of the primary somatosensory cortex location.

The R code for this example is similar to Example 7.7.




Fig. 7.10. Analysis of power for fMRI data at five locations, $L = 3$ and critical values $F_{.001}(6, 120) = 4.04$ for stimulus and $F_{.001}(12, 120) = 3.02$ for consciousness and interaction.

```

n = 128; n.freq = 1 + n/2
Fr = (0:(n.freq-1))/n; nFr = 1:(n.freq/2)
N = c(5,4,5,3,5,4); n.subject = sum(N)
n.para = 6      # number of parameters
L = 3
df.stm=2*L*(3-1)  # stimulus (3 levels: Brush,Heat,Shock)
df.con=2*L*(2-1)  # conscious (2 levels: Awake,Sedated)
df.int=2*L*(3-1)*(2-1) # interaction
den.df= 2*L*(n.subject-n.para) # df for full model
# Design Matrix:      mu  a1  a2   b  g1  g2
Z1 = outer(rep(1,N[1]), c(1, 1, 0, 1, 1, 0))
Z2 = outer(rep(1,N[2]), c(1, 0, 1, 1, 0, 1))
Z3 = outer(rep(1,N[3]), c(1, -1, -1, 1, -1, -1))
Z4 = outer(rep(1,N[4]), c(1, 1, 0, -1, -1, 0))
Z5 = outer(rep(1,N[5]), c(1, 0, 1, -1, -1, -1))
Z6 = outer(rep(1,N[6]), c(1, -1, -1, -1, 1, 1))
Z = rbind(Z1, Z2, Z3, Z4, Z5, Z6); ZZ = t(Z)%*%Z
rep(NA, n)-> SSEF-> SSE.stm-> SSE.con-> SSE.int

```

```

HatF = Z%%solve(ZZ,t(Z))
Hat.stm = Z[,-(2:3)]%%solve(ZZ[-(2:3),-(2:3)], t(Z[,-(2:3)]))
Hat.con = Z[, -4]%%solve(ZZ[-4,-4], t(Z[, -4]))
Hat.int = Z[,-(5:6)]%%solve(ZZ[-(5:6),-(5:6)], t(Z[,-(5:6)]))
par(mfrow=c(5,3), mar=c(3.5,4,0,0), oma=c(0,0,2,2), mgp = c(1.6,.6,0))
loc.name = c("Cortex 1","Cortex 2","Cortex 3","Cortex 4","Caudate",
             "Thalamus 1","Thalamus 2","Cerebellum 1","Cerebellum 2")
for(Loc in c(1:4,9)) { # only Loc 1 to 4 and 9 used
  i = 6*(Loc-1)
  Y = cbind(fmri[[i+1]], fmri[[i+2]], fmri[[i+3]], fmri[[i+4]], fmri[[i+5]],
            fmri[[i+6]])
  Y = mvfft(spec.taper(Y, p=.5))/sqrt(n); Y = t(Y)
  for (k in 1:n) {
    SSY=Re(Conj(t(Y[,k])))*%Y[,k])
    SSReg= Re(Conj(t(Y[,k])))*%HatF*%Y[,k])
    SSEF[k]=SSY-SSReg
    SSReg=Re(Conj(t(Y[,k])))*%Hat.stm*%Y[,k])
    SSE.stm[k] = SSY-SSReg
    SSReg=Re(Conj(t(Y[,k])))*%Hat.con*%Y[,k])
    SSE.con[k]=SSY-SSReg
    SSReg=Re(Conj(t(Y[,k])))*%Hat.int*%Y[,k])
    SSE.int[k]=SSY-SSReg
  }
  # Smooth
  sSEF = filter(SSEF, rep(1/L, L), circular = TRUE)
  sSSE.stm = filter(SSE.stm, rep(1/L, L), circular = TRUE)
  sSSE.con = filter(SSE.con, rep(1/L, L), circular = TRUE)
  sSSE.int = filter(SSE.int, rep(1/L, L), circular = TRUE)
  eF.stm = (den.df/df.stm)*(sSSE.stm-sSEF)/sSEF
  eF.con = (den.df/df.con)*(sSSE.con-sSEF)/sSEF
  eF.int = (den.df/df.int)*(sSSE.int-sSEF)/sSEF
  plot(Fr[nFr], eF.stm[nFr], type="l", xlab="Frequency", ylab="F Statistic",
       ylim=c(0,12))
  abline(h=qf(.999, df.stm, den.df),lty=2)
  if(Loc==1) mtext("Stimulus", side=3, line=.3, cex=1)
  mtext(loc.name[Loc], side=2, line=3, cex=.9)
  plot(Fr[nFr], eF.con[nFr], type="l", xlab="Frequency", ylab="F Statistic",
       ylim=c(0,12))
  abline(h=qf(.999, df.con, den.df),lty=2)
  if(Loc==1) mtext("Consciousness", side=3, line=.3, cex=1)
  plot(Fr[nFr], eF.int[nFr], type="l", xlab="Frequency", ylab="F Statistic",
       ylim=c(0,12))
  abline(h=qf(.999, df.int, den.df),lty=2)
  if(Loc==1) mtext("Interaction", side=3, line= .3, cex=1)
}

```

SIMULTANEOUS INFERENCE

In the previous examples involving the fMRI data, it would be helpful to focus on the components that contributed most to the rejection of the equal means hypothesis. One way to accomplish this is to develop a test for the significance of an arbitrary linear compound of the form

$$\Psi(\omega_k) = \mathbf{A}^*(\omega_k)\mathbf{B}(\omega_k), \quad (7.86)$$

where the components of the vector $\mathbf{A}(\omega_k) = (A_1(\omega_k), A_2(\omega_k), \dots, A_q(\omega_k))'$ are chosen in such a way as to isolate particular linear functions of parameters in the regression vector $\mathbf{B}(\omega_k)$ in the regression model (7.80). This argument suggests developing a test of the hypothesis $\Psi(\omega_k) = 0$ for all possible values of the linear coefficients in the compound (7.86) as is done in the conventional analysis of variance approach (see, for example, Scheffé, 1959).

Recalling the material involving the regression models of the form (7.50), the linear compound (7.86) can be estimated by

$$\widehat{\Psi}(\omega_k) = \mathbf{A}^*(\omega_k) \widehat{\mathbf{B}}(\omega_k), \quad (7.87)$$

where $\widehat{\mathbf{B}}(\omega_k)$ is the estimated vector of regression coefficients given by (7.51) and independent of the error spectrum $s_{y \cdot z}^2(\omega_k)$ in (7.53). It is possible to show the maximum of the ratio

$$F(\mathbf{A}) = \frac{N - q}{q} \frac{|\widehat{\Psi}(\omega_k) - \Psi(\omega_k)|^2}{s_{y \cdot z}^2(\omega_k) Q(\mathbf{A})}, \quad (7.88)$$

where

$$Q(\mathbf{A}) = \mathbf{A}^*(\omega_k) S_z^{-1}(\omega_k) \mathbf{A}(\omega_k) \quad (7.89)$$

is bounded by a statistic that has an F -distribution with $2q$ and $2(N - q)$ degrees of freedom. Testing the hypothesis that the compound has a particular value, usually $\Psi(\omega_k) = 0$, then proceeds naturally, by comparing the statistic (7.88) evaluated at the hypothesized value with the α level point on an $F_{2q, 2(N-q)}$ distribution. We can choose an infinite number of compounds of the form (7.86) and the test will still be valid at level α . As before, arguing the error spectrum is relatively constant over a band enables us to smooth the numerator and denominator of (7.88) separately over L frequencies so distribution involving the smooth components is $F_{2Lq, 2L(N-q)}$.

Example 7.9 Simultaneous Inference for the fMRI Series

As an example, consider the previous tests for significance of the fMRI factors, in which we have indicated the primary effects are among the stimuli but have not investigated which of the stimuli, heat, brushing, or shock, had the most effect. To analyze this further, consider the means model (7.81) and a 6×1 contrast vector of the form

$$\widehat{\Psi} = \mathbf{A}^*(\omega_k) \widehat{\mathbf{B}}(\omega_k) = \sum_{i=1}^6 A_i^*(\omega_k) \mathbf{Y}_{i \cdot}(\omega_k), \quad (7.90)$$

where the means are easily shown to be the regression coefficients in this particular case. In this case, the means are ordered by columns; the first three means are the three levels of stimuli for the awake state, and the last three means are the levels for the anesthetized state. In this special case, the denominator terms are

$$Q = \sum_{i=1}^6 \frac{|A_i(\omega_k)|^2}{N_i}, \quad (7.91)$$

with $SSE(\omega_k)$ available in (7.84). In order to evaluate the effect of a particular stimulus, like brushing over the two levels of consciousness, we may take $A_1(\omega_k) = A_4(\omega_k) = 1$ for the two brush levels and $A(\omega_k) = 0$ zero otherwise. From Figure 7.11, we see that, at the first and third cortex locations, brush and heat are both significant, whereas the fourth cortex shows only brush and the second cerebellum shows only heat. Shock appears to be transmitted relatively weakly, when averaged over the awake and mildly anesthetized states.

The R code for this example is as follows.

```
n = 128; n.freq = 1 + n/2
Fr = (0:(n.freq-1))/n; nFr = 1:(n.freq/2)
N = c(5,4,5,3,5,4); n.subject = sum(N); L = 3
# Design Matrix
Z1 = outer(rep(1,N[1]), c(1,0,0,0,0,0))
Z2 = outer(rep(1,N[2]), c(0,1,0,0,0,0))
Z3 = outer(rep(1,N[3]), c(0,0,1,0,0,0))
Z4 = outer(rep(1,N[4]), c(0,0,0,1,0,0))
Z5 = outer(rep(1,N[5]), c(0,0,0,0,1,0))
Z6 = outer(rep(1,N[6]), c(0,0,0,0,0,1))
Z = rbind(Z1, Z2, Z3, Z4, Z5, Z6); ZZ = t(Z)%*%Z
# Contrasts: 6 by 3
A = rbind(diag(1,3), diag(1,3))
nq = nrow(A); num.df = 2*L*nq; den.df = 2*L*(n.subject-nq)
HatF = Z%*%solve(ZZ, t(Z)) # full model
rep(NA, n)-> SSEF -> SSER; eF = matrix(0,n,3)
par(mfrow=c(5,3), mar=c(3.5,4,0,0), oma=c(0,0,2,2), mgp = c(1.6,.6,0))
loc.name = c("Cortex 1", "Cortex 2", "Cortex 3", "Cortex 4", "Caudate",
"Thalamus 1", "Thalamus 2", "Cerebellum 1", "Cerebellum 2")
cond.name = c("Brush", "Heat", "Shock")
for(Loc in c(1:4,9)) {
  i = 6*(Loc-1)
  Y = cbind(fmri[[i+1]], fmri[[i+2]], fmri[[i+3]], fmri[[i+4]], fmri[[i+5]],
fmri[[i+6]])
  Y = mvfft(spec.taper(Y, p=.5))/sqrt(n); Y = t(Y)
  for (cond in 1:3){
    Q = t(A[,cond])%*%solve(ZZ, A[,cond])
    HR = A[,cond]%*%solve(ZZ, t(Z))
    for (k in 1:n){
      SSY = Re(Conj(t(Y[,k]))%*%Y[,k])
      SSReg = Re(Conj(t(Y[,k]))%*%HatF%*%Y[,k])
      SSEF[k] = (SSY-SSReg)*Q
      SSReg = HR%*%Y[,k]
      SSER[k] = Re(SSReg*Conj(SSReg))
    }
  }
# Smooth
sSSEF = filter(SSEF, rep(1/L, L), circular = TRUE)
sSSER = filter(SSER, rep(1/L, L), circular = TRUE)
eF[,cond] = (den.df/num.df)*(sSSER/sSSEF) }
plot(Fr[nFr], eF[nFr,1], type="l", xlab="Frequency", ylab="F Statistic",
ylim=c(0,5))
abline(h=qf(.999, num.df, den.df),lty=2)
if(Loc==1) mtext("Brush", side=3, line=.3, cex=1)
mtext(loc.name[Loc], side=2, line=3, cex=.9)
plot(Fr[nFr], eF[nFr,2], type="l", xlab="Frequency", ylab="F Statistic",
ylim=c(0,5))
```




Fig. 7.11. Power in simultaneous linear compounds at five locations, enhancing brush, heat, and shock effects, $L = 3$, $F_{.001}(36, 120) = 2.16$.

```

abline(h=qf(.999, num.df, den.df),lty=2)
if(Loc==1) mtext("Heat", side=3, line=.3, cex=1)
plot(Fr[nFr], eF[nFr,3], type="l", xlab="Frequency", ylab="F Statistic",
      ylim=c(0,5))
abline(h= qf(.999, num.df, den.df) ,lty=2)
if(Loc==1) mtext("Shock", side=3, line=.3, cex=1)
}

```

MULTIVARIATE TESTS

Although it is possible to develop multivariate regression along lines analogous to the usual real valued case, we will only look at tests involving equality of group means and spectral matrices, because these tests appear to be used most often in applications. For these results, consider the p -variate time series $\mathbf{y}_{ijt} = (y_{ijt1}, \dots, y_{ijt p})'$ to have arisen from observations on $j = 1, \dots, N_i$ individuals in group i , all having mean $\boldsymbol{\mu}_{it}$ and stationary autocovariance matrix $\Gamma_i(h)$. Denote the DFTs of the group mean

vectors as $\mathbf{Y}_{i.}(\omega_k)$ and the $p \times p$ spectral matrices as $\widehat{f}_i(\omega_k)$ for the $i = 1, 2, \dots, I$ groups. Assume the same general properties as for the vector series considered in §7.3.

In the multivariate case, we obtain the analogous versions of (7.83) and (7.84) as the between cross-power and within cross-power matrices

$$SPR(\omega_k) = \sum_{i=1}^I \sum_{j=1}^{N_i} (\mathbf{Y}_{i.}(\omega_k) - \mathbf{Y}_{..}(\omega_k)) (\mathbf{Y}_{i.}(\omega_k) - \mathbf{Y}_{..}(\omega_k))^* \quad (7.92)$$

and

$$SPE(\omega_k) = \sum_{i=1}^I \sum_{j=1}^{N_i} (\mathbf{Y}_{ij}(\omega_k) - \mathbf{Y}_{i.}(\omega_k)) (\mathbf{Y}_{ij}(\omega_k) - \mathbf{Y}_{i.}(\omega_k))^*. \quad (7.93)$$

The equality of means test is rejected using the fact that the likelihood ratio test yields a monotone function of

$$\Lambda(\omega_k) = \frac{|SPE(\omega_k)|}{|SPE(\omega_k) + SPR(\omega_k)|}. \quad (7.94)$$

Khatri (1965) and Hannan (1970) give the approximate distribution of the statistic

$$\chi_{2(I-1)p}^2 = -2 \left(\sum N_i - I - p - 1 \right) \log \Lambda(\omega_k) \quad (7.95)$$

as chi-squared with $2(I-1)p$ degrees of freedom when the group means are equal.

The case of $I = 2$ groups reduces to Hotelling's T^2 , as has been shown by Giri (1965), where

$$T^2 = \frac{N_1 N_2}{(N_1 + N_2)} [\mathbf{Y}_{1.}(\omega_k) - \mathbf{Y}_{2.}(\omega_k)]^* \widehat{f}_v^{-1}(\omega_k) [\mathbf{Y}_{1.}(\omega_k) - \mathbf{Y}_{2.}(\omega_k)], \quad (7.96)$$

where

$$\widehat{f}_v(\omega_k) = \frac{SPE(\omega_k)}{\sum_i N_i - I} \quad (7.97)$$

is the pooled error spectrum given in (7.93), with $I = 2$. The test statistic, in this case, is

$$F_{2p, 2(N_1+N_2-p-1)} = \frac{(N_1 + N_2 - 2)p}{(N_1 + N_2 - p - 1)} T^2, \quad (7.98)$$

which was shown by Giri (1965) to have the indicated limiting F -distribution with $2p$ and $2(N_1 + N_2 - p - 1)$ degrees of freedom when the means are the same. The classical t -test for inequality of two univariate means will be just (7.97) and (7.98) with $p = 1$.

Testing equality of the spectral matrices is also of interest, not only for discrimination and pattern recognition, as considered in the next section, but also as a test indicating whether the equality of means test, which assumes equal spectral matrices,

is valid. The test evolves from the likelihood ration criterion, which compares the single group spectral matrices

$$\widehat{f}_i(\omega_k) = \frac{1}{N_i - 1} \sum_{j=1}^{N_i} (\mathbf{Y}_{ij}(\omega_k) - \mathbf{Y}_{i\cdot}(\omega_k)) (\mathbf{Y}_{ij}(\omega_k) - \mathbf{Y}_{i\cdot}(\omega_k))^* \quad (7.99)$$

with the pooled spectral matrix (7.97). A modification of the likelihood ratio test, which incorporates the degrees of freedom $M_i = N_i - 1$ and $M = \sum M_i$ rather than the sample sizes into the likelihood ratio statistic, uses

$$L'(\omega_k) = \frac{M^{Mp}}{\prod_{i=1}^I M_i^{M_i p}} \frac{\prod |M_i \widehat{f}_i(\omega_k)|^{M_i}}{|M \widehat{f}_v(\omega_k)|^M}. \quad (7.100)$$

Krishnaiah et al. (1976) have given the moments of $L'(\omega_k)$ and calculated 95% critical points for $p = 3, 4$ using a Pearson Type I approximation. For reasonably large samples involving smoothed spectral estimators, the approximation involving the first term of the usual chi-squared series will suffice and Shumway (1982) has given

$$\chi_{(I-1)p^2}^2 = -2r \log L'(\omega_k), \quad (7.101)$$

where

$$1 - r = \frac{(p+1)(p-1)}{6p(I-1)} \left(\sum_i M_i^{-1} - M^{-1} \right), \quad (7.102)$$

with an approximate chi-squared distribution with $(I-1)p^2$ degrees of freedom when the spectral matrices are equal. Introduction of smoothing over L frequencies leads to replacing M_j and M by LM_j and LM in the equations above.

Of course, it is often of great interest to use the above result for testing equality of two univariate spectra, and it is obvious from the material in Chapter 4

$$F_{2LM_1, 2LM_2} = \frac{\widehat{f}_1(\omega)}{\widehat{f}_2(\omega)} \quad (7.103)$$

will have the requisite F -distribution with $2LM_1$ and $2LM_2$ degrees of freedom when spectra are smoothed over L frequencies.

Example 7.10 Equality of Means and Spectral Matrices for Earthquakes and Explosions

An interesting problem arises when attempting to develop a methodology for discriminating between waveforms originating from explosions and those that came from the more commonly occurring earthquakes. Figure 7.2 shows a small subset of a larger population of bivariate series consisting of two phases from each of eight earthquakes and eight explosions. If the large-sample approximations to normality hold for the DFTs of these series, it is of interest to know whether the differences between the two classes are better represented by the mean functions or by the spectral matrices. The tests described above can be applied to look at these two questions. The upper left panel of Figure 7.12 shows the test statistic (7.98) with the straight line denoting the critical level for $\alpha = .001$, i.e., $F_{.001}(4, 26) = 7.36$, for equal means using $L = 1$, and the test statistic remains well below its critical value at all frequencies, implying that the means of the two classes of series are not significantly different. Checking Figure 7.2 shows little reason exists to suspect that either the earthquakes or explosions have a nonzero mean signal. Checking the equality of the spectra and the spectral matrices, however, leads to a different conclusion. Some smoothing ($L = 21$) is useful here, and univariate tests on both the P and S components using (7.103) and $N_1 = N_2 = 8$ lead to strong rejections of the equal spectra hypotheses. The rejection seems stronger for the S component and we might tentatively identify that component as being dominant. Testing equality of the spectral matrices using (7.101) and $\chi^2_{.001}(4) = 18.47$ shows a similar strong rejection of the equality of spectral matrices. We use these results to suggest optimal discriminant functions based on spectral differences in the next section.

The R code for this example is as follows. We make use of the recycling feature of R and the fact that the data are bivariate to produce simple code specific to this problem in order to avoid having to use multiple arrays.

```
P=1:1024; S=P+1024; N=8; n=1024; p.dim=2; m=10; L=2*m+1
eq.P = as.ts(eqexp[P,1:8]); eq.S = as.ts(eqexp[S,1:8])
eq.m = cbind(rowMeans(eq.P), rowMeans(eq.S))
ex.P = as.ts(eqexp[P,9:16]); ex.S = as.ts(eqexp[S,9:16])
ex.m = cbind(rowMeans(ex.P), rowMeans(ex.S))
m.diff = mvfft(eq.m - ex.m)/sqrt(n)
eq.Pf = mvfft(eq.P-eq.m[,1])/sqrt(n)
eq.Sf = mvfft(eq.S-eq.m[,2])/sqrt(n)
ex.Pf = mvfft(ex.P-ex.m[,1])/sqrt(n)
ex.Sf = mvfft(ex.S-ex.m[,2])/sqrt(n)
fv11=rowSums(eq.Pf*Conj(eq.Pf))+rowSums(ex.Pf*Conj(ex.Pf))/(2*(N-1))
fv12=rowSums(eq.Pf*Conj(eq.Sf))+rowSums(ex.Pf*Conj(ex.Sf))/(2*(N-1))
fv22=rowSums(eq.Sf*Conj(eq.Sf))+rowSums(ex.Sf*Conj(ex.Sf))/(2*(N-1))
fv21 = Conj(fv12)
# Equal Means
T2 = rep(NA, 512)
for (k in 1:512){
  fvk = matrix(c(fv11[k], fv21[k], fv12[k], fv22[k]), 2, 2)
  dk = as.matrix(m.diff[k,])
  T2[k] = Re((N/2)*Conj(t(dk))%*%solve(fvk,dk)) }
eF = T2*(2*p.dim*(N-1))/(2*N-p.dim-1)
par(mfrow=c(2,2), mar=c(3,3,2,1), mgp = c(1.6,.6,0), cex.main=1.1)
```





Fig. 7.12. Tests for equality of means, spectra, and spectral matrices for the earthquake and explosion data $p = 2$, $L = 21$, $n = 1024$ points at 40 points per second.

```

freq = 40*(0:511)/n # Hz
plot(freq, eF, type="l", xlab="Frequency (Hz)", ylab="F Statistic",
      main="Equal Means")
abline(h=qf(.999, 2*p.dim, 2*(2*N-p.dim-1)))
# Equal P
kd = kernel("daniell",m);
u = Re(rowSums(eq.Pf*Conj(eq.Pf))/(N-1))
feq.P = kernapply(u, kd, circular=TRUE)
u = Re(rowSums(ex.Pf*Conj(ex.Pf))/(N-1))
fex.P = kernapply(u, kd, circular=TRUE)
plot(freq, feq.P[1:512]/fex.P[1:512], type="l", xlab="Frequency (Hz)",
      ylab="F Statistic", main="Equal P-Spectra")
abline(h=qf(.999, 2*L*(N-1), 2*L*(N-1)))
# Equal S
u = Re(rowSums(eq.Sf*Conj(eq.Sf))/(N-1))
feq.S = kernapply(u, kd, circular=TRUE)
u = Re(rowSums(ex.Sf*Conj(ex.Sf))/(N-1))
fex.S = kernapply(u, kd, circular=TRUE)
plot(freq, feq.S[1:512]/fex.S[1:512], type="l", xlab="Frequency (Hz)",
      ylab="F Statistic", main="Equal S-Spectra")
abline(h=qf(.999, 2*L*(N-1), 2*L*(N-1)))
# Equal Spectra
u = rowSums(eq.Pf*Conj(eq.Sf))/(N-1)
feq.PS = kernapply(u, kd, circular=TRUE)
u = rowSums(ex.Pf*Conj(ex.Sf))/(N-1)
fex.PS = kernapply(u, kd, circular=TRUE)
fv11 = kernapply(fv11, kd, circular=TRUE)
fv22 = kernapply(fv22, kd, circular=TRUE)
fv12 = kernapply(fv12, kd, circular=TRUE)
Mi = L*(N-1); M = 2*Mi

```

```

TS = rep(NA,512)
for (k in 1:512){
  det.feq.k= Re(feq.P[k]*feq.S[k] - feq.PS[k]*Conj(feq.PS[k]))
  det.fex.k= Re(fex.P[k]*fex.S[k] - fex.PS[k]*Conj(fex.PS[k]))
  det.fv.k = Re(fv11[k]*fv22[k] - fv12[k]*Conj(fv12[k]))
  log.n1 = log(M)*(M*p.dim); log.d1 = log(Mi)*(2*Mi*p.dim)
  log.n2 = log(Mi)*2 +log(det.feq.k)*Mi + log(det.fex.k)*Mi
  log.d2 = (log(M)+log(det.fv.k))*M
  r = 1 - ((p.dim+1)*(p.dim-1)/6*p.dim*(2-1))*(2/Mi - 1/M)
  TS[k] = -2*r*(log.n1+log.n2-log.d1-log.d2) }
plot(freq, TS, type="l", xlab="Frequency (Hz)", ylab="Chi-Sq Statistic",
      main="Equal Spectral Matrices")
abline(h = qchisq(.9999, p.dim^2))

```

7.7 Discriminant and Cluster Analysis

The extension of classical pattern-recognition techniques to experimental time series is a problem of great practical interest. A series of observations indexed in time often produces a pattern that may form a basis for discriminating between different classes of events. As an example, consider [Figure 7.2](#), which shows regional (100-2000 km) recordings of several typical Scandinavian earthquakes and mining explosions measured by stations in Scandinavia. A listing of the events is given in Kakizawa et al. (1998). The problem of discriminating between mining explosions and earthquakes is a reasonable proxy for the problem of discriminating between nuclear explosions and earthquakes. This latter problem is one of critical importance for monitoring a comprehensive test-ban treaty. Time series classification problems are not restricted to geophysical applications, but occur under many and varied circumstances in other fields. Traditionally, the detection of a signal embedded in a noise series has been analyzed in the engineering literature by statistical pattern recognition techniques (see [Problem 7.10](#) and [Problem 7.11](#)).

The historical approaches to the problem of discriminating among different classes of time series can be divided into two distinct categories. The optimality approach, as found in the engineering and statistics literature, makes specific Gaussian assumptions about the probability density functions of the separate groups and then develops solutions that satisfy well-defined minimum error criteria. Typically, in the time series case, we might assume the difference between classes is expressed through differences in the theoretical mean and covariance functions and use likelihood methods to develop an optimal classification function. A second class of techniques, which might be described as a feature extraction approach, proceeds more heuristically by looking at quantities that tend to be good visual discriminators for well-separated populations and have some basis in physical theory or intuition. Less attention is paid to finding functions that are approximations to some well-defined optimality criterion.

As in the case of regression, both time domain and frequency domain approaches to discrimination will exist. For relatively short univariate series, a time domain approach that follows conventional multivariate discriminant analysis as described in

conventional multivariate texts, such as Anderson (1984) or Johnson and Wichern (1992) may be preferable. We might even characterize differences by the autocovariance functions generated by different ARMA or state-space models. For longer multivariate time series that can be regarded as stationary after the common mean has been subtracted, the frequency domain approach will be easier computationally because the np dimensional vector in the time domain, represented here as $\mathbf{x} = (\mathbf{x}'_1, \mathbf{x}'_t, \dots, \mathbf{x}'_n)'$, with $\mathbf{x}_t = (x_{t1}, \dots, x_{tp})'$, will be reduced to separate computations made on the p -dimensional DFTs. This happens because of the approximate independence of the DFTs, $\mathbf{X}(\omega_k)$, $0 \leq \omega_k \leq 1$, a property that we have often used in preceding chapters.

Finally, the grouping properties of measures like the discrimination information and likelihood-based statistics can be used to develop measures of disparity for clustering multivariate time series. In this section, we define a measure of disparity between two multivariate time series by the spectral matrices of the two processes and then apply hierarchical clustering and partitioning techniques to identify natural groupings within the bivariate earthquake and explosion populations.

THE GENERAL DISCRIMINATION PROBLEM

The general problem of classifying a vector time series $\mathbf{x}$ occurs in the following way. We observe a time series $\mathbf{x}$ known to belong to one of g populations, denoted by $\Pi_1, \Pi_2, \dots, \Pi_g$. The general problem is to assign or classify this observation into one of the g groups in some optimal fashion. An example might be the $g = 2$ populations of earthquakes and explosions shown in Figure 7.2. We would like to classify the unknown event, shown as NZ in the bottom two panels, as belonging to either the earthquake (Π_1) or explosion (Π_2) populations. To solve this problem, we need an optimality criterion that leads to a statistic $T(\mathbf{x})$ that can be used to assign the NZ event to either the earthquake or explosion populations. To measure the success of the classification, we need to evaluate errors that can be expected in the future relating to the number of earthquakes classified as explosions (false alarms) and the number of explosions classified as earthquakes (missed signals).

The problem can be formulated by assuming the observed series $\mathbf{x}$ has a probability density $p_i(\mathbf{x})$ when the observed series is from population Π_i for $i = 1, \dots, g$. Then, partition the space spanned by the np -dimensional process $\mathbf{x}$ into g mutually exclusive regions $R_1, R_2, \dots, R_g$ such that, if $\mathbf{x}$ falls in R_i , we assign $\mathbf{x}$ to population Π_i . The misclassification probability is defined as the probability of classifying the observation into population Π_j when it belongs to Π_i , for $j \neq i$ and would be given by the expression

$$P(j|i) = \int_{R_j} p_i(\mathbf{x}) d\mathbf{x}. \quad (7.104)$$

The overall total error probability depends also on the prior probabilities, say, $\pi_1, \pi_2, \dots, \pi_g$, of belonging to one of the g groups. For example, the probability that an observation $\mathbf{x}$ originates from Π_i and is then classified into Π_j is obviously $\pi_i P(j|i)$, and the total error probability becomes

$$P_e = \sum_{i=1}^g \pi_i \sum_{j \neq i} P(j|i). \quad (7.105)$$

Although costs have not been incorporated into (7.105), it is easy to do so by multiplying $P(j|i)$ by $C(j|i)$, the cost of assigning a series from population Π_i to Π_j .

The overall error P_e is minimized by classifying $\mathbf{x}$ into Π_i if

$$\frac{p_i(\mathbf{x})}{p_j(\mathbf{x})} > \frac{\pi_j}{\pi_i} \quad (7.106)$$

for all $j \neq i$ (see, for example, Anderson, 1984). A quantity of interest, from the Bayesian perspective, is the posterior probability an observation belongs to population Π_i , conditional on observing $\mathbf{x}$, say,

$$P(\Pi_i|\mathbf{x}) = \frac{\pi_i p_i(\mathbf{x})}{\sum_j \pi_j p_j(\mathbf{x})}. \quad (7.107)$$

The procedure that classifies $\mathbf{x}$ into the population Π_i for which the posterior probability is largest is equivalent to that implied by using the criterion (7.106). The posterior probabilities give an intuitive idea of the relative odds of belonging to each of the plausible populations.

Many situations occur, such as in the classification of earthquakes and explosions, in which there are only $g = 2$ populations of interest. For two populations, the Neyman–Pearson lemma implies, in the absence of prior probabilities, classifying an observation into Π_1 when

$$\frac{p_1(\mathbf{x})}{p_2(\mathbf{x})} > K \quad (7.108)$$

minimizes each of the error probabilities for a fixed value of the other. The rule is identical to the Bayes rule (7.106) when $K = \pi_2/\pi_1$.

The theory given above takes a simple form when the vector $\mathbf{x}$ has a p -variate normal distribution with mean vectors $\boldsymbol{\mu}_j$ and covariance matrices Σ_j under Π_j for $j = 1, 2, \dots, g$. In this case, simply use

$$p_j(\mathbf{x}) = (2\pi)^{-p/2} |\Sigma_j|^{-1/2} \exp\left\{-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_j)' \Sigma_j^{-1}(\mathbf{x} - \boldsymbol{\mu}_j)\right\}. \quad (7.109)$$

The classification functions are conveniently expressed by quantities that are proportional to the logarithms of the densities, say,

$$g_j(\mathbf{x}) = -\frac{1}{2} \ln |\Sigma_j| - \frac{1}{2} \mathbf{x}' \Sigma_j^{-1} \mathbf{x} + \boldsymbol{\mu}_j' \Sigma_j^{-1} \mathbf{x} - \frac{1}{2} \boldsymbol{\mu}_j' \Sigma_j^{-1} \boldsymbol{\mu}_j + \ln \pi_j. \quad (7.110)$$

In expressions involving the log likelihood, we will generally ignore terms involving the constant $-\ln 2\pi$. For this case, we may assign an observation $\mathbf{x}$ to population Π_i whenever

$$g_i(\mathbf{x}) > g_j(\mathbf{x}) \quad (7.111)$$

for $j \neq i, j = 1, \dots, g$ and the posterior probability (7.107) has the form

$$P(\Pi_i|\mathbf{x}) = \frac{\exp\{g_i(\mathbf{x})\}}{\sum_j \exp\{g_j(\mathbf{x})\}}.$$

A common situation occurring in applications involves classification for $g = 2$ groups under the assumption of multivariate normality and equal covariance matrices; i.e., $\Sigma_1 = \Sigma_2 = \Sigma$. Then, the criterion (7.111) can be expressed in terms of the linear discriminant function

$$\begin{aligned} d_l(\mathbf{x}) &= g_1(\mathbf{x}) - g_2(\mathbf{x}) \\ &= (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)' \Sigma^{-1} \mathbf{x} - \frac{1}{2} (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)' \Sigma^{-1} (\boldsymbol{\mu}_1 + \boldsymbol{\mu}_2) + \ln \frac{\pi_1}{\pi_2}, \end{aligned} \quad (7.112)$$

where we classify into Π_1 or Π_2 according to whether $d_l(\mathbf{x}) \geq 0$ or $d_l(\mathbf{x}) < 0$. The linear discriminant function is clearly a combination of normal variables and, for the case $\pi_1 = \pi_2 = .5$, will have mean $D^2/2$ under Π_1 and mean $-D^2/2$ under Π_2 , with variances given by D^2 under both hypotheses, where

$$D^2 = (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)' \Sigma^{-1} (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2) \quad (7.113)$$

is the Mahalanobis distance between the mean vectors $\boldsymbol{\mu}_1$ and $\boldsymbol{\mu}_2$. In this case, the two misclassification probabilities (7.1) are

$$P(1|2) = P(2|1) = \Phi\left(-\frac{D}{2}\right), \quad (7.114)$$

and the performance is directly related to the Mahalanobis distance (7.113).

For the case in which the covariance matrices cannot be assumed to be the same, the discriminant function takes a different form, with the difference $g_1(\mathbf{x}) - g_2(\mathbf{x})$ taking the form

$$\begin{aligned} d_q(\mathbf{x}) &= -\frac{1}{2} \ln \frac{|\Sigma_1|}{|\Sigma_2|} - \frac{1}{2} \mathbf{x}' (\Sigma_1^{-1} - \Sigma_2^{-1}) \mathbf{x} \\ &\quad + (\boldsymbol{\mu}_1' \Sigma_1^{-1} - \boldsymbol{\mu}_2' \Sigma_2^{-1}) \mathbf{x} + \ln \frac{\pi_1}{\pi_2} \end{aligned} \quad (7.115)$$

for $g = 2$ groups. This discriminant function differs from the equal covariance case in the linear term and in a nonlinear quadratic term involving the differing covariance matrices. The distribution theory is not tractable for the quadratic case so no convenient expression like (7.114) is available for the error probabilities for the quadratic discriminant function.

A difficulty in applying the above theory to real data is that the group mean vectors $\boldsymbol{\mu}_j$ and covariance matrices Σ_j are seldom known. Some engineering problems, such as the detection of a signal in white noise, assume the means and covariance parameters are known exactly, and this can lead to an optimal solution (see Problems 7.14 and 7.15). In the classical multivariate situation, it is possible to collect a sample of N_i training vectors from group Π_i , say, $\mathbf{x}_{ij}$, for $j = 1, \dots, N_i$, and use them to estimate the mean vectors and covariance matrices for each of the groups $i = 1, 2, \dots, g$; i.e., simply choose $\mathbf{x}_i$ and

$$S_i = (N_i - 1)^{-1} \sum_{j=1}^{N_i} (\mathbf{x}_{ij} - \mathbf{x}_{i\cdot})(\mathbf{x}_{ij} - \mathbf{x}_{i\cdot})' \quad (7.116)$$

as the estimators for $\boldsymbol{\mu}_i$ and Σ_i , respectively. In the case in which the covariance matrices are assumed to be equal, simply use the pooled estimator

$$S = \left(\sum_i N_i - g \right)^{-1} \sum_i (N_i - 1) S_i. \quad (7.117)$$

For the case of a linear discriminant function, we may use

$$\widehat{g_i(\mathbf{x})} = \mathbf{x}'_i \cdot S^{-1} \mathbf{x} - \frac{1}{2} \mathbf{x}'_i \cdot S^{-1} \mathbf{x}_i + \log \pi_i \quad (7.118)$$

as a simple estimator for $g_i(\mathbf{x})$. For large samples, $\mathbf{x}_{i\cdot}$ and S converge to $\boldsymbol{\mu}_i$ and Σ in probability so $\widehat{g_i(\mathbf{x})}$ converges in distribution to $g_i(\mathbf{x})$ in that case. The procedure works reasonably well for the case in which $N_i, i = 1, \dots, g$ are large, relative to the length of the series n , a case that is relatively rare in time series analysis. For this reason, we will resort to using spectral approximations for the case in which data are given as long time series.

The performance of sample discriminant functions can be evaluated in several different ways. If the population parameters are known, (7.113) and (7.114) can be evaluated directly. If the parameters are estimated, the estimated Mahalanobis distance $\widehat{D^2}$ can be substituted for the theoretical value in very large samples. Another approach is to calculate the apparent error rates using the result of applying the classification procedure to the training samples. If n_{ij} denotes the number of observations from population Π_j classified into Π_i , the sample error rates can be estimated by the ratio

$$\widehat{P(i|j)} = \frac{n_{ij}}{\sum_i n_{ij}} \quad (7.119)$$

for $i \neq j$. If the training samples are not large, this procedure may be biased and a resampling option like cross-validation or the bootstrap can be employed. A simple version of cross-validation is the jackknife procedure proposed by Lachenbruch and Mickey (1968), which holds out the observation to be classified, deriving the classification function from the remaining observations. Repeating this procedure for each of the members of the training sample and computing (7.119) for the holdout samples leads to better estimators of the error rates.

Example 7.11 Discriminant Analysis Using Amplitudes from Earthquakes and Explosions

We can give a simple example of applying the above procedures to the logarithms of the amplitudes of the separate P and S components of the original earthquake and explosion traces. The logarithms (base 10) of the maximum peak-to-peak amplitudes of the P and S components, denoted by $\log_{10} P$ and $\log_{10} S$, can be considered as two-dimensional feature vectors, say, $\mathbf{x} = (x_1, x_2)' = (\log_{10} P, \log_{10} S)'$, from a




Fig. 7.13. Classification of earthquakes and explosions based on linear discriminant analysis using the magnitude features.

bivariate normal population with differing means and covariances. The original data, from Kakizawa et al. (1998), are shown in Figure 7.13. The figure includes the Novaya Zemlya (NZ) event of unknown origin. The tendency of the earthquakes to have higher values for $\log_{10} S$, relative to $\log_{10} P$ has been noted by many and the use of the logarithm of the ratio, i.e., $\log_{10} P - \log_{10} S$ in some references (see Lay, 1997, pp. 40–41) is a tacit indicator that a linear function of the two parameters will be a useful discriminant.

The sample means $\mathbf{x}_1 = (.346, 1.024)'$ and $\mathbf{x}_2 = (.922, .993)'$, and covariance matrices

$$S_1 = \begin{pmatrix} .026 & -.007 \\ -.007 & .010 \end{pmatrix} \quad \text{and} \quad S_2 = \begin{pmatrix} .025 & -.001 \\ -.001 & .010 \end{pmatrix}$$

are immediate from (7.116), with the pooled covariance matrix given by

$$S = \begin{pmatrix} .026 & -.004 \\ -.004 & .010 \end{pmatrix}$$

from (7.117). Although the covariance matrices are not equal, we try the linear discriminant function anyway, which yields (with equal prior probabilities $\pi_1 =$

$\pi_2 = .5$) the sample discriminant functions

$$\widehat{g_1}(\mathbf{x}) = 30.668x_1 + 111.411x_2 - 62.401$$

and

$$\widehat{g_2}(\mathbf{x}) = 54.048x_1 + 117.255x_2 - 83.142$$

from (7.118), with the estimated linear discriminant function (7.112) as

$$\widehat{d_l}(\mathbf{x}) = -23.380x_1 - 5.843x_2 + 20.740.$$

The jackknifed posterior probabilities of being an earthquake for the earthquake group ranged from .621 to 1.000, whereas the explosion probabilities for the explosion group ranged from .717 to 1.000. The unknown event, NZ, was classified as an explosion, with posterior probability .960.

The R code for this example is as follows.

```
P = 1:1024; S = P+1024
mag.P = log10(apply(eqexp[P,], 2, max) - apply(eqexp[P,], 2, min))
mag.S = log10(apply(eqexp[S,], 2, max) - apply(eqexp[S,], 2, min))
eq.P = mag.P[1:8]; eq.S = mag.S[1:8]
ex.P = mag.P[9:16]; ex.S = mag.S[9:16]
NZ.P = mag.P[17]; NZ.S = mag.S[17]
# Compute linear discriminant function
cov.eq = var(cbind(eq.P, eq.S))
cov.ex = var(cbind(ex.P, ex.S))
cov.pooled = (cov.ex + cov.eq)/2
means.eq = colMeans(cbind(eq.P, eq.S))
means.ex = colMeans(cbind(ex.P, ex.S))
slopes.eq = solve(cov.pooled, means.eq)
inter.eq = -sum(slopes.eq*means.eq)/2
slopes.ex = solve(cov.pooled, means.ex)
inter.ex = -sum(slopes.ex*means.ex)/2
d.slopes = slopes.eq - slopes.ex
d.inter = inter.eq - inter.ex
# Classify new observation
new.data = cbind(NZ.P, NZ.S)
d = sum(d.slopes*new.data) + d.inter
post.eq = exp(d)/(1+exp(d))
# Print (disc function, posteriors) and plot results
cat(d.slopes[1], "mag.P +", d.slopes[2], "mag.S +", d.inter, "\n")
cat("P(EQ|data) =", post.eq, " P(EX|data) =", 1-post.eq, "\n")
plot(eq.P, eq.S, xlim=c(0,1.5), ylim=c(.75,1.25), xlab="log mag(P)", ylab=
  "log mag(S)", pch = 8, cex=1.1, lwd=2, main="Classification Based on
  Magnitude Features")
points(ex.P, ex.S, pch = 6, cex=1.1, lwd=2)
points(new.data, pch = 3, cex=1.1, lwd=2)
abline(a = -d.inter/d.slopes[2], b = -d.slopes[1]/d.slopes[2])
text(eq.P+.07, eq.S+.005, label=names(eqexp[1:8]), cex=.8)
text(ex.P+.07, ex.S+.003, label=names(eqexp[9:16]), cex=.8)
text(NZ.P+.05, NZ.S+.003, label=names(eqexp[17]), cex=.8)
legend("topright", c("EQ", "EX", "NZ"), pch=c(8,6,3), pt.lwd=2, cex=1.1)
# Cross-validation
all.data = rbind(cbind(eq.P, eq.S), cbind(ex.P, ex.S))
post.eq <- rep(NA, 8) -> post.ex
```



```

for(j in 1:16) {
  if (j <= 8){samp.eq = all.data[-c(j, 9:16),]
  samp.ex = all.data[9:16,]}
  if (j > 8){samp.eq = all.data[1:8,]
  samp.ex = all.data[-c(j, 1:8),] }
  df.eq = nrow(samp.eq)-1; df.ex = nrow(samp.ex)-1
  mean.eq = colMeans(samp.eq); mean.ex = colMeans(samp.ex)
  cov.eq = var(samp.eq); cov.ex = var(samp.ex)
  cov.pooled = (df.eq*cov.eq + df.ex*cov.ex)/(df.eq + df.ex)
  slopes.eq = solve(cov.pooled, mean.eq)
  inter.eq = -sum(slopes.eq*mean.eq)/2
  slopes.ex = solve(cov.pooled, mean.ex)
  inter.ex = -sum(slopes.ex*mean.ex)/2
  d.slopes = slopes.eq - slopes.ex
  d.inter = inter.eq - inter.ex
  d = sum(d.slopes*all.data[j,]) + d.inter
  if (j <= 8) post.eq[j] = exp(d)/(1+exp(d))
  if (j > 8) post.ex[j-8] = 1/(1+exp(d)) }
Posterior = cbind(1:8, post.eq, 1:8, post.ex)
colnames(Posterior) = c("EQ", "P(EQ|data)", "EX", "P(EX|data)")
round(Posterior,3) # Results from Cross-validation (not shown)

```

FREQUENCY DOMAIN DISCRIMINATION

The feature extraction approach often works well for discriminating between classes of univariate or multivariate series when there is a simple low-dimensional vector that seems to capture the essence of the differences between the classes. It still seems sensible, however, to develop optimal methods for classification that exploit the differences between the multivariate means and covariance matrices in the time series case. Such methods can be based on the Whittle approximation to the log likelihood given in §7.2. In this case, the vector DFTs, say, $\mathbf{X}(\omega_k)$, are assumed to be approximately normal, with means $\mathbf{M}_j(\omega_k)$ and spectral matrices $f_j(\omega_k)$ for population $\mathcal{I}_j$ at frequencies $\omega_k = k/n$, for $k = 0, 1, \dots, [n/2]$, and are approximately uncorrelated at different frequencies, say, ω_k and ω_ℓ for $k \neq \ell$. Then, writing the complex normal densities as in §7.2 leads to a criterion similar to (7.110); namely,

$$\begin{aligned}
 g_j(\mathbf{X}) = \ln \pi_j - \sum_{0 < \omega_k < 1/2} & \left[\ln |f_j(\omega_k)| + \mathbf{X}^*(\omega_k) f_j^{-1}(\omega_k) \mathbf{X}(\omega_k) \right. \\
 & \left. - 2 \mathbf{M}_j^*(\omega_k) f_j^{-1}(\omega_k) \mathbf{X}(\omega_k) + \mathbf{M}_j^*(k) f_j^{-1}(\omega_k) \mathbf{M}_j(\omega_k) \right],
 \end{aligned} \tag{7.120}$$

where the sum goes over frequencies for which $|f_j(\omega_k)| \neq 0$. The periodicity of the spectral density matrix and DFT allows adding over $0 < k < 1/2$. The classification rule is as in (7.111).

In the time series case, it is more likely the discriminant analysis involves assuming the covariance matrices are different and the means are equal. For example, the tests, shown in Figure 7.12, imply, for the earthquakes and explosions, the primary differences are in the bivariate spectral matrices and the means are essentially the same. For this case, it will be convenient to write the Whittle approximation to the log likelihood in the form

$$\ln p_j(\mathbf{X}) = \sum_{0 < \omega_k < 1/2} \left[-\ln |f_j(\omega_k)| - \mathbf{X}^*(\omega_k) f_j^{-1}(\omega_k) \mathbf{X}(\omega_k) \right], \quad (7.121)$$

where we have omitted the prior probabilities from the equation. The quadratic detector in this case can be written in the form

$$\ln p_j(\mathbf{X}) = \sum_{0 < \omega_k < 1/2} \left[-\ln |f_j(\omega_k)| - \text{tr}\{I(\omega_k) f_j^{-1}(\omega_k)\} \right], \quad (7.122)$$

where

$$I(\omega_k) = \mathbf{X}(\omega_k) \mathbf{X}^*(\omega_k) \quad (7.123)$$

denotes the periodogram matrix. For equal prior probabilities, we may assign an observation $\mathbf{x}$ into population Π_i whenever

$$\ln p_i(\mathbf{X}) > \ln p_j(\mathbf{X}) \quad (7.124)$$

for $j \neq i, j = 1, 2, \dots, g$.

Numerous authors have considered various versions of discriminant analysis in the frequency domain. Shumway and Unger (1974) considered (7.120) for $p = 1$ and equal covariance matrices, so the criterion reduces to a simple linear one. They apply the criterion to discriminating between earthquakes and explosions using teleseismic P wave data in which the means over the two groups might be considered as fixed. Alagón (1989) and Dargahi-Noubary and Laycock (1981) considered discriminant functions of the form (7.120) in the univariate case when the means are zero and the spectra for the two groups are different. Taniguchi et al. (1994) adopted (7.121) as a criterion and discussed its non-Gaussian robustness. Shumway (1982) reviews general discriminant functions in both the univariate and multivariate time series cases.

MEASURES OF DISPARITY

Before proceeding to examples of discriminant and cluster analysis, it is useful to consider the relation to the Kullback–Leibler (K-L) discrimination information, as defined in Problem 2.4. Using the spectral approximation and noting the periodogram matrix has the approximate expectation

$$E_j I(\omega_k) = f_j(\omega_k)$$

under the assumption that the data come from population Π_j , and approximating the ratio of the densities by

$$\ln \frac{p_1(\mathbf{X})}{p_2(\mathbf{X})} = \sum_{0 < \omega_k < 1/2} \left[-\ln \frac{|f_1(\omega_k)|}{|f_2(\omega_k)|} - \text{tr}\left\{(f_2^{-1}(\omega_k) - f_1^{-1}(\omega_k))I(\omega_k)\right\} \right],$$

we may write the approximate discrimination information as

$$\begin{aligned}
I(f_1; f_2) &= \frac{1}{n} E_1 \ln \frac{p_1(\mathbf{X})}{p_2(\mathbf{X})} \\
&= \frac{1}{n} \sum_{0 < \omega_k < 1/2} \left[\text{tr} \{f_1(\omega_k) f_2^{-1}(\omega_k)\} - \ln \frac{|f_1(\omega_k)|}{|f_2(\omega_k)|} - p \right]. \quad (7.125)
\end{aligned}$$

The approximation may be carefully justified by noting the multivariate normal time series $\mathbf{x} = (\mathbf{x}'_1, \mathbf{x}'_2, \dots, \mathbf{x}'_n)$ with zero means and $np \times np$ stationary covariance matrices Γ_1 and Γ_2 will have $p, n \times n$ blocks, with elements of the form $\gamma_{ij}^{(\ell)}(s-t)$, $s, t = 1, \dots, n, i, j = 1, \dots, p$ for population Π_ℓ , $\ell = 1, 2$. The discrimination information, under these conditions, becomes

$$I(1; 2 : \mathbf{x}) = \frac{1}{n} E_1 \ln \frac{p_1(\mathbf{x})}{p_2(\mathbf{x})} = \frac{1}{n} \left[\text{tr} \{ \Gamma_1 \Gamma_2^{-1} \} - \ln \frac{|\Gamma_1|}{|\Gamma_2|} - np \right]. \quad (7.126)$$

The limiting result

$$\lim_{n \rightarrow \infty} I(1; 2 : \mathbf{x}) = \frac{1}{2} \int_{-1/2}^{1/2} \left[\text{tr} \{f_1(\omega) f_2^{-1}(\omega)\} - \ln \frac{|f_1(\omega)|}{|f_2(\omega)|} - p \right] d\omega$$

has been shown, in various forms, by Pinsker (1964), Hannan (1970), and Kazakos and Papantoni-Kazakos (1980). The discrete version of (7.125) is just the approximation to the integral of the limiting form. The K-L measure of disparity is not a true distance, but it can be shown that $I(1; 2) \geq 0$, with equality if and only if $f_1(\omega) = f_2(\omega)$ almost everywhere. This result makes it potentially suitable as a measure of disparity between the two densities.

A connection exists, of course, between the discrimination information number, which is just the expectation of the likelihood criterion and the likelihood itself. For example, we may measure the disparity between the sample and the process defined by the theoretical spectrum $f_j(\omega_k)$ corresponding to population Π_j in the sense of Kullback (1958), as $I(\hat{f}; f_j)$, where

$$\hat{f}(\omega_k) = \sum_{\ell=-m}^m h_\ell I(\omega_k + \ell/n) \quad (7.127)$$

denotes the smoothed spectral matrix with weights $\{h_\ell\}$. The likelihood ratio criterion can be thought of as measuring the disparity between the periodogram and the theoretical spectrum for each of the populations. To make the discrimination information finite, we replace the periodogram implied by the log likelihood by the sample spectrum. In this case, the classification procedure can be regarded as finding the population closest, in the sense of minimizing disparity between the sample and theoretical spectral matrices. The classification in this case proceeds by simply choosing the population Π_j that minimizes $I(\hat{f}; f_j)$, i.e., assigning $\mathbf{x}$ to population Π_i whenever

$$I(\hat{f}; f_i) < I(\hat{f}; f_j) \quad (7.128)$$

for $j \neq i, j = 1, 2, \dots, g$.

Kakizawa et al. (1998) proposed using the Chernoff (CH) information measure (Chernoff, 1952, Renyi, 1961), defined as

$$B_\alpha(1; 2) = -\ln E_2 \left\{ \left(\frac{p_2(\mathbf{x})}{p_1(\mathbf{x})} \right)^\alpha \right\}, \quad (7.129)$$

where the measure is indexed by a regularizing parameter α , for $0 < \alpha < 1$. When $\alpha = .5$, the Chernoff measure is the symmetric divergence proposed by Bhattacharya (1943). For the multivariate normal case,

$$B_\alpha(1; 2 : \mathbf{x}) = \frac{1}{n} \left[\ln \frac{|\alpha \Gamma_1 + (1 - \alpha) \Gamma_2|}{|\Gamma_2|} - \alpha \ln \frac{|\Gamma_1|}{|\Gamma_2|} \right]. \quad (7.130)$$

The large sample spectral approximation to the Chernoff information measure is analogous to that for the discrimination information, namely,

$$B_\alpha(f_1; f_2) = \frac{1}{2n} \sum_{0 < \omega_k < 1/2} \left[\ln \frac{|\alpha f_1(\omega_k) + (1 - \alpha) f_2(\omega_k)|}{|f_2(\omega_k)|} - \alpha \ln \frac{|f_1(\omega_k)|}{|f_2(\omega_k)|} \right]. \quad (7.131)$$

The Chernoff measure, when divided by $\alpha(1 - \alpha)$, behaves like the discrimination information in the limit in the sense that it converges to $I(1; 2 : \mathbf{x})$ for $\alpha \rightarrow 0$ and to $I(2; 1 : \mathbf{x})$ for $\alpha \rightarrow 1$. Hence, near the boundaries of the parameter α , it tends to behave like discrimination information and for other values represents a compromise between the two information measures. The classification rule for the Chernoff measure reduces to assigning $\mathbf{x}$ to population Π_i whenever

$$B_\alpha(\hat{f}; f_i) < B_\alpha(\hat{f}; f_j) \quad (7.132)$$

for $j \neq i, j = 1, 2, \dots, g$.

Although the classification rules above are well defined if the group spectral matrices are known, this will not be the case in general. If there are g training samples, $\mathbf{x}_{ij}, j = 1, \dots, N_i, i = 1, \dots, g$, with N_i vector observations available in each group, the natural estimator for the spectral matrix of the group i is just the average spectral matrix (7.99), namely, with $\hat{f}_{ij}(\omega_k)$ denoting the estimated spectral matrix of series j from the i -th population,

$$\hat{f}_i(\omega_k) = \frac{1}{N_i} \sum_{j=1}^{N_i} \hat{f}_{ij}(\omega_k). \quad (7.133)$$

A second consideration is the choice of the regularization parameter α for the Chernoff criterion, (7.131). For the case of $g = 2$ groups, it should be chosen to maximize the disparity between the two group spectra, as defined in (7.131). Kakizawa et al. (1998) simply plot (7.131) as a function of α , using the estimated group spectra in (7.133), choosing the value that gives the maximum disparity between the two groups.

Example 7.12 Discriminant Analysis on Seismic Data

The simplest approaches to discriminating between the earthquake and explosion groups have been based on either the relative amplitudes of the P and S phases, as in [Figure 7.5](#) or on relative power components in various frequency bands. Considerable effort has been expended on using various spectral ratios involving the bivariate P and S phases as discrimination features. Kakizawa et al. (1998) mention a number of measures that have been used in the seismological literature as features. These features include ratios of power for the two phases and ratios of power components in high- and low-frequency bands. The use of such features of the spectrum suggests an optimal procedure based on discriminating between the spectral matrices of two stationary processes would be reasonable. The fact that the hypothesis that the spectral matrices were equal, tested in [Example 7.10](#), was also soundly rejected suggests the use of a discriminant function based on spectral differences. Recall the sampling rate is 40 points per second, leading to a folding frequency of 20 Hz.

[Figure 7.14](#) displays the diagonal elements of the average spectral matrices for each group. The maximum value of the estimated Chernoff disparity $B_\alpha(\hat{f}_1; \hat{f}_2)$ occurs for $\alpha = .4$, and we use that value in the discriminant criterion (7.131). [Figure 7.15](#) shows the results of using the Chernoff differences along with the Kullback-Leibler differences for classification. The differences are the measures for earthquakes minus explosions, so negative values of the differences indicate earthquake and positive values indicate explosion. Hence, points in the first quadrant of [Figure 7.15](#) are classified an explosion and points in the third quadrant are classified as earthquakes. We note that Explosion 6 is misclassified as an earthquake. Also, Earthquake 1, which falls in the fourth quadrant has an uncertain classification, the Chernoff distance classifies it as an earthquake, however, the Kullback-Leibler difference classifies it as an explosion.

The NZ event of unknown origin was also classified using these distance measures, and, as in [Example 7.11](#), it is classified as an explosion. The Russians have asserted no mine blasting or nuclear testing occurred in the area in question, so the event remains as somewhat of a mystery. The fact that it was relatively removed geographically from the test set may also have introduced some uncertainties into the procedure.

The R code for this example is as follows.

```
P = 1:1024; S = P+1024; p.dim = 2; n = 1024
eq = as.ts(eqexp[, 1:8])
ex = as.ts(eqexp[, 9:16])
nz = as.ts(eqexp[, 17])
f.eq <- array(dim=c(8, 2, 2, 512)) -> f.ex
f.NZ = array(dim=c(2, 2, 2, 512))
# below calculates determinant for 2x2 Hermitian matrix
det.c <- function(mat){return(Re(mat[1,1]*mat[2,2]-mat[1,2]*mat[2,1]))}
L = c(15,13,5) # for smoothing
for (i in 1:8){ # compute spectral matrices
  f.eq[i,,] = mvspec(cbind(eq[P,i], eq[S,i]), spans=L, taper=.5)$fxx
  f.ex[i,,] = mvspec(cbind(ex[P,i], ex[S,i]), spans=L, taper=.5)$fxx
  u = mvspec(cbind(nz[P], nz[S]), spans=L, taper=.5)
```




Fig. 7.14. Average P-spectra and S-spectra of the earthquake and explosion series. The insert on the upper right shows the smoothing kernel used; the resulting bandwidth is about .75 Hz.

```
f.NZ = u$fx
bndwidth = u$bandwidth*sqrt(12)*40 # about .75 Hz
fhat.eq = apply(f.eq, 2:4, mean) # average spectra
fhat.ex = apply(f.ex, 2:4, mean)
# plot the average spectra
par(mfrow=c(2,2), mar=c(3,3,2,1), mgp = c(1.6,.6,0))
Fr = 40*(1:512)/n
plot(Fr,Re(fhat.eq[1,1,]),type="l",xlab="Frequency (Hz)",ylab="")
plot(Fr,Re(fhat.eq[2,2,]),type="l",xlab="Frequency (Hz)",ylab="")
plot(Fr,Re(fhat.ex[1,1,]),type="l",xlab="Frequency (Hz)",ylab="")
plot(Fr,Re(fhat.ex[2,2,]),type="l",xlab="Frequency (Hz)",ylab="")
mtext("Average P-spectra", side=3, line=-1.5, adj=.2, outer=TRUE)
mtext("Earthquakes", side=2, line=-1, adj=.8, outer=TRUE)
mtext("Average S-spectra", side=3, line=-1.5, adj=.82, outer=TRUE)
mtext("Explosions", side=2, line=-1, adj=.2, outer=TRUE)
par(fig = c(.75, 1, .75, 1), new = TRUE)
ker = kernel("modified.daniell", L)$coef; ker = c(rev(ker),ker[-1])
plot((-33:33)/40, ker, type="l", ylab="", xlab="", cex.axis=.7,
      yaxp=c(0,.04,2))
```




Fig. 7.15. Classification (by quadrant) of earthquakes and explosions using the Chernoff and Kullback-Leibler differences.

```
# Choose alpha
Balpaha = rep(0,19)
for (i in 1:19){ alf=i/20
  for (k in 1:256) {
    Balpaha[i]= Balpaha[i] + Re(log(det.c(alf*fhat.ex[, ,k] +
      (1-alf)*fhat.eq[, ,k])/det.c(fhat.eq[, ,k])) -
      alf*log(det.c(fhat.ex[, ,k])/det.c(fhat.eq[, ,k]))) }
  alf = which.max(Balpaha)/20 # alpha = .4
# Calculate Information Criteria
rep(0,17) -> KLDiff -> BDiff -> KLeq -> KLex -> Beq -> Bex
for (i in 1:17){
  if (i <= 8) f0 = f.eq[i, ,]
  if (i > 8 & i <= 16) f0 = f.ex[i-8, ,]
  if (i == 17) f0 = f.NZ
  for (k in 1:256) { # only use freqs out to .25
    tr = Re(sum(diag(solve(fhat.eq[, ,k], f0[, ,k])))
    KLeq[i] = KLeq[i] + tr + log(det.c(fhat.eq[, ,k])) - log(det.c(f0[, ,k]))
    Beq[i] = Beq[i] +
      Re(log(det.c(alf*f0[, ,k]+(1-alf)*fhat.eq[, ,k])/det.c(fhat.eq[, ,k])) -
      alf*log(det.c(f0[, ,k])/det.c(fhat.eq[, ,k])))
    tr = Re(sum(diag(solve(fhat.ex[, ,k], f0[, ,k])))
    KLex[i] = KLex[i] + tr + log(det.c(fhat.ex[, ,k])) - log(det.c(f0[, ,k]))
```

```

Bex[i] = Bex[i] +
  Re(log(det.c(alf*f0[,k]+(1-alf)*fhat.ex[,k])/det.c(fhat.ex[,k])) -
    alf*log(det.c(f0[,k])/det.c(fhat.ex[,k]))) }
KLDiff[i] = (KLeq[i] - KLex[i])/n
BDiff[i] = (Beq[i] - Bex[i])/(2*n) }
x.b = max(KLDiff)+.1; x.a = min(KLDiff)-.1
y.b = max(BDiff)+.01; y.a = min(BDiff)-.01
dev.new()
plot(KLDiff[9:16], BDiff[9:16], type="p", xlim=c(x.a,x.b), ylim=c(y.a,y.b),
  cex=1.1,lwd=2, xlab="Kullback-Leibler Difference",ylab="Chernoff
  Difference", main="Classification Based on Chernoff and K-L Distances",
  pch=6)
points(KLDiff[1:8], BDiff[1:8], pch=8, cex=1.1, lwd=2)
points(KLDiff[17], BDiff[17], pch=3, cex=1.1, lwd=2)
legend("topleft", legend=c("EQ", "EX", "NZ"), pch=c(8,6,3), pt.lwd=2)
abline(h=0, v=0, lty=2, col="gray")
text(KLDiff[-c(1,2,3,7,14)]-.075, BDiff[-c(1,2,3,7,14)],
  label=names(eqexp[-c(1,2,3,7,14)]), cex=.7)
text(KLDiff[c(1,2,3,7,14)]+.075, BDiff[c(1,2,3,7,14)],
  label=names(eqexp[c(1,2,3,7,14)]), cex=.7)

```

CLUSTER ANALYSIS

For the purpose of clustering, it may be more useful to consider a symmetric disparity measures and we introduce the J-Divergence measure

$$J(f_1; f_2) = I(f_1; f_2) + I(f_2; f_1) \quad (7.134)$$

and the symmetric Chernoff number

$$JB_\alpha(f_1; f_2) = B_\alpha(f_1; f_2) + B_\alpha(f_2; f_1) \quad (7.135)$$

for that purpose. In this case, we define the disparity between the sample spectral matrix of a single vector, $\mathbf{x}$, and the population Π_j as

$$J(\hat{f}; f_j) = I(\hat{f}; f_j) + I(f_j; \hat{f}) \quad (7.136)$$

and

$$JB_\alpha(\hat{f}; f_j) = B_\alpha(\hat{f}; f_j) + B_\alpha(f_j; \hat{f}), \quad (7.137)$$

respectively and use these as quasi-distances between the vector and population Π_j .

The measures of disparity can be used to cluster multivariate time series. The symmetric measures of disparity, as defined above ensure that the disparity between f_i and f_j is the same as the disparity between f_j and f_i . Hence, we will consider the symmetric forms (7.136) and (7.137) as quasi-distances for the purpose of defining a distance matrix for input into one of the standard clustering procedures (see Johnson and Wichern, 1992). In general, we may consider either hierarchical or partitioned clustering methods using the quasi-distance matrix as an input.

For purposes of illustration, we may use the symmetric divergence (7.136), which implies the quasi-distance between sample series with estimated spectral matrices $\hat{f}_i$ and $\hat{f}_j$ would be (7.136); i.e.,

$$J(\widehat{f}_i; \widehat{f}_j) = \frac{1}{n} \sum_{0 < \omega_k < 1/2} \left[\text{tr}\{\widehat{f}_i(\omega_k) \widehat{f}_j^{-1}(\omega_k)\} + \text{tr}\{\widehat{f}_j(\omega_k) \widehat{f}_i^{-1}(\omega_k)\} - 2p \right], \quad (7.138)$$

for $i \neq j$. We can also use the comparable form for the Chernoff divergence, but we may not want to make an assumption for the regularization parameter α .

For hierarchical clustering, we begin by clustering the two members of the population that minimize the disparity measure (7.138). Then, these two items form a cluster, and we can compute distances between unclustered items as before. The distance between unclustered items and a current cluster is defined here as the average of the distances to elements in the cluster. Again, we combine objects that are closest together. We may also compute the distance between the unclustered items and clustered items as the closest distance, rather than the average. Once a series is in a cluster, it stays there. At each stage, we have a fixed number of clusters, depending on the merging stage.

Alternatively, we may think of clustering as a partitioning of the sample into a prespecified number of groups. MacQueen (1967) has proposed this using k-means clustering, using the Mahalanobis distance between an observation and the group mean vectors. At each stage, a reassignment of an observation into its closest affinity group is possible. To see how this procedure applies in the current context, consider a preliminary partition into a fixed number of groups and define the disparity between the spectral matrix of the observation, say, $\widehat{f}$, and the average spectral matrix of the group, say, $\widehat{f}_i$, as $J(\widehat{f}; \widehat{f}_i)$, where the group spectral matrix can be estimated by (7.133). At any pass, a single series is reassigned to the group for which its disparity is minimized. The reassignment procedure is repeated until all observations stay in their current groups. Of course, the number of groups must be specified for each repetition of the partitioning algorithm and a starting partition must be chosen. This assignment can either be random or chosen from a preliminary hierarchical clustering, as described above.

Example 7.13 Cluster Analysis for Earthquakes and Explosions

It is instructive to try a clustering procedure on the population of known earthquakes and explosions. Figure 7.16 shows the results of applying the Partitioning Around Medoids (PAM) clustering algorithm, which is essentially a robustification of the k-means procedure (see Ch. 2 of Kaufman & Rousseeuw, 1990), under the assumption that two groups are appropriate. The two-group partition tends to produce a final partition that agrees closely with the known configuration with earthquake 1 (EQ1) and explosion 8 (EX8) being misclassified; as in previous examples, the NZ event is classified as an explosion.

The R code for this example uses the `cluster` package and our `mvspec` script for estimating spectral matrices.

```
library(cluster)
P = 1:1024; S = P+1024; p.dim = 2; n = 1024
eq = as.ts(eqexp[, 1:8])
ex = as.ts(eqexp[, 9:16])
nz = as.ts(eqexp[, 17])
f = array(dim=c(17, 2, 2, 512))
```




Fig. 7.16. Clustering results for the earthquake and explosion series based on symmetric divergence using a robust version of k-means clustering with two groups. Circles indicate Group I classification, triangles indicate Group II classification.

```

L = c(15, 15)      # for smoothing
for (i in 1:8){    # compute spectral matrices
  f[i,,] = mvspec(cbind(eq[P,i], eq[S,i]), spans=L, taper=.5)$fxx
  f[i+8,,] = mvspec(cbind(ex[P,i], ex[S,i]), spans=L, taper=.5)$fxx }
f[17,,] = mvspec(cbind(nz[P], nz[S]), spans=L, taper=.5)$fxx
JD = matrix(0, 17, 17)
# Calculate Symmetric Information Criteria
for (i in 1:16){
  for (j in (i+1):17){
    for (k in 1:256) { # only use freqs out to .25
      tr1 = Re(sum(diag(solve(f[i,,k], f[j,,k]))))
      tr2 = Re(sum(diag(solve(f[j,,k], f[i,,k]))))
      JD[i,j] = JD[i,j] + (tr1 + tr2 - 2*p.dim)}}
  JD = (JD + t(JD))/n
colnames(JD) = c(colnames(eq), colnames(ex), "NZ")
rownames(JD) = colnames(JD)
cluster.2 = pam(JD, k = 2, diss = TRUE)
summary(cluster.2) # print results
par(mgp = c(1.6,.6,0), cex=3/4, cex.lab=4/3, cex.main=4/3)

```

```
clusplot(JD, cluster.2$cluster, col.clus=1, labels=3, lines=0, col.p=1,
  main="Clustering Results for Explosions and Earthquakes")
text(-7,-.5, "Group I", cex=1.1, font=2)
text(1, 5, "Group II", cex=1.1, font=2)
```

7.8 Principal Components and Factor Analysis

In this section, we introduce the related topics of spectral domain principal components analysis and factor analysis for time series. The topics of principal components and canonical analysis in the frequency domain are rigorously presented in Brillinger (1981, Chapters 9 and 10) and many of the details concerning these concepts can be found there.

The techniques presented here are related to each other in that they focus on extracting pertinent information from spectral matrices. This information is important because dealing directly with a high-dimensional spectral matrix $f(\omega)$ itself is somewhat cumbersome because it is a function into the set of complex, nonnegative-definite, Hermitian matrices. We can view these techniques as easily understood, parsimonious tools for exploring the behavior of vector-valued time series in the frequency domain with minimal loss of information. Because our focus is on spectral matrices, we assume for convenience that the time series of interest have zero means; the techniques are easily adjusted in the case of nonzero means.

In this and subsequent sections, it will be convenient to work occasionally with complex-valued time series. A $p \times 1$ complex-valued time series can be represented as $\mathbf{x}_t = \mathbf{x}_{1t} - i\mathbf{x}_{2t}$, where $\mathbf{x}_{1t}$ is the real part and $\mathbf{x}_{2t}$ is the imaginary part of $\mathbf{x}_t$. The process is said to be stationary if $E(\mathbf{x}_t)$ and $E(\mathbf{x}_{t+h}\mathbf{x}_t^*)$ exist and are independent of time t . The $p \times p$ autocovariance function,

$$\Gamma_{xx}(h) = E(\mathbf{x}_{t+h}\mathbf{x}_t^*) - E(\mathbf{x}_{t+h})E(\mathbf{x}_t^*),$$

of $\mathbf{x}_t$ satisfies conditions similar to those of the real-valued case. Writing $\Gamma_{xx}(h) = \{\gamma_{ij}(h)\}$, for $i, j = 1, \dots, p$, we have (i) $\gamma_{ii}(0) \geq 0$ is real, (ii) $|\gamma_{ij}(h)|^2 \leq \gamma_{ii}(0)\gamma_{jj}(0)$ for all integers h , and (iii) $\Gamma_{xx}(h)$ is a non-negative definite function. The spectral theory of complex-valued vector time series is analogous to the real-valued case. For example, if $\sum_h \|\Gamma_{xx}(h)\| < \infty$, the spectral density matrix of the complex series $\mathbf{x}_t$ is given by

$$f_{xx}(\omega) = \sum_{h=-\infty}^{\infty} \Gamma_{xx}(h) \exp(-2\pi i h \omega).$$

PRINCIPAL COMPONENTS

Classical principal component analysis (PCA) is concerned with explaining the variance–covariance structure among p variables, $\mathbf{x} = (x_1, \dots, x_p)'$, through a few linear combinations of the components of $\mathbf{x}$. Suppose we wish to find a linear combination

$$y = \mathbf{c}'\mathbf{x} = c_1x_1 + \dots + c_px_p \quad (7.139)$$

of the components of $\mathbf{x}$ such that $\text{var}(y)$ is as large as possible. Because $\text{var}(y)$ can be increased by simply multiplying $\mathbf{c}$ by a constant, it is common to restrict $\mathbf{c}$ to be of unit length; that is, $\mathbf{c}'\mathbf{c} = 1$. Noting that $\text{var}(y) = \mathbf{c}'\Sigma_{xx}\mathbf{c}$, where Σ_{xx} is the $p \times p$ variance–covariance matrix of $\mathbf{x}$, another way of stating the problem is to find $\mathbf{c}$ such that

$$\max_{\mathbf{c} \neq \mathbf{0}} \frac{\mathbf{c}'\Sigma_{xx}\mathbf{c}}{\mathbf{c}'\mathbf{c}}. \quad (7.140)$$

Denote the eigenvalue–eigenvector pairs of Σ_{xx} by $\{(\lambda_1, \mathbf{e}_1), \dots, (\lambda_p, \mathbf{e}_p)\}$, where $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_p \geq 0$, and the eigenvectors are of unit length. The solution to (7.140) is to choose $\mathbf{c} = \mathbf{e}_1$, in which case the linear combination $y_1 = \mathbf{e}_1'\mathbf{x}$ has maximum variance, $\text{var}(y_1) = \lambda_1$. In other words,

$$\max_{\mathbf{c} \neq \mathbf{0}} \frac{\mathbf{c}'\Sigma_{xx}\mathbf{c}}{\mathbf{c}'\mathbf{c}} = \frac{\mathbf{e}_1'\Sigma_{xx}\mathbf{e}_1}{\mathbf{e}_1'\mathbf{e}_1} = \lambda_1. \quad (7.141)$$

The linear combination, $y_1 = \mathbf{e}_1'\mathbf{x}$, is called the first principal component. Because the eigenvalues of Σ_{xx} are not necessarily unique, the first principal component is not necessarily unique.

The second principal component is defined to be the linear combination $y_2 = \mathbf{c}'\mathbf{x}$ that maximizes $\text{var}(y_2)$ subject to $\mathbf{c}'\mathbf{c} = 1$ and such that $\text{cov}(y_1, y_2) = 0$. The solution is to choose $\mathbf{c} = \mathbf{e}_2$, in which case, $\text{var}(y_2) = \lambda_2$. In general, the k -th principal component, for $k = 1, 2, \dots, p$, is the linear combination $y_k = \mathbf{c}'\mathbf{x}$ that maximizes $\text{var}(y_k)$ subject to $\mathbf{c}'\mathbf{c} = 1$ and such that $\text{cov}(y_k, y_j) = 0$, for $j = 1, 2, \dots, k-1$. The solution is to choose $\mathbf{c} = \mathbf{e}_k$, in which case $\text{var}(y_k) = \lambda_k$.

One measure of the importance of a principal component is to assess the proportion of the total variance attributed to that principal component. The total variance of $\mathbf{x}$ is defined to be the sum of the variances of the individual components; that is, $\text{var}(x_1) + \dots + \text{var}(x_p) = \sigma_{11} + \dots + \sigma_{pp}$, where σ_{jj} is the j -th diagonal element of Σ_{xx} . This sum is also denoted as $\text{tr}(\Sigma_{xx})$, or the trace of Σ_{xx} . Because $\text{tr}(\Sigma_{xx}) = \lambda_1 + \dots + \lambda_p$, the proportion of the total variance attributed to the k -th principal component is given simply by $\text{var}(y_k) / \text{tr}(\Sigma_{xx}) = \lambda_k / \sum_{j=1}^p \lambda_j$.

Given a random sample $\mathbf{x}_1, \dots, \mathbf{x}_n$, the sample principal components are defined as above, but with Σ_{xx} replaced by the sample variance–covariance matrix, $S_{xx} = (n-1)^{-1} \sum_{i=1}^n (\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i - \bar{\mathbf{x}})'$. Further details can be found in the introduction to classical principal component analysis in Johnson and Wichern (1992, Chapter 9).

For the case of time series, suppose we have a zero mean, $p \times 1$, stationary vector process $\mathbf{x}_t$ that has a $p \times p$ spectral density matrix given by $f_{xx}(\omega)$. Recall $f_{xx}(\omega)$ is a complex-valued, nonnegative-definite, Hermitian matrix. Using the analogy of classical principal components, and in particular (7.139) and (7.140), suppose, for a fixed value of ω , we want to find a complex-valued univariate process $y_t(\omega) = \mathbf{c}(\omega)^* \mathbf{x}_t$, where $\mathbf{c}(\omega)$ is complex, such that the spectral density of $y_t(\omega)$ is maximized at frequency ω , and $\mathbf{c}(\omega)$ is of unit length, $\mathbf{c}(\omega)^* \mathbf{c}(\omega) = 1$. Because, at frequency ω , the spectral density of $y_t(\omega)$ is $f_y(\omega) = \mathbf{c}(\omega)^* f_{xx}(\omega) \mathbf{c}(\omega)$, the problem can be restated as: Find complex vector $\mathbf{c}(\omega)$ such that

$$\max_{\mathbf{c}(\omega) \neq \mathbf{0}} \frac{\mathbf{c}(\omega)^* f_{xx}(\omega) \mathbf{c}(\omega)}{\mathbf{c}(\omega)^* \mathbf{c}(\omega)}. \quad (7.142)$$

Let $\{(\lambda_1(\omega), \mathbf{e}_1(\omega)), \dots, (\lambda_p(\omega), \mathbf{e}_p(\omega))\}$ denote the eigenvalue–eigenvector pairs of $f_{xx}(\omega)$, where $\lambda_1(\omega) \geq \lambda_2(\omega) \geq \dots \geq \lambda_p(\omega) \geq 0$, and the eigenvectors are of unit length. We note that the eigenvalues of a Hermitian matrix are real. The solution to (7.142) is to choose $\mathbf{c}(\omega) = \mathbf{e}_1(\omega)$; in which case the desired linear combination is $y_t(\omega) = \mathbf{e}_1(\omega)^* \mathbf{x}_t$. For this choice,

$$\max_{\mathbf{c}(\omega) \neq \mathbf{0}} \frac{\mathbf{c}(\omega)^* f_{xx}(\omega) \mathbf{c}(\omega)}{\mathbf{c}(\omega)^* \mathbf{c}(\omega)} = \frac{\mathbf{e}_1(\omega)^* f_{xx}(\omega) \mathbf{e}_1(\omega)}{\mathbf{e}_1(\omega)^* \mathbf{e}_1(\omega)} = \lambda_1(\omega). \quad (7.143)$$

This process may be repeated for any frequency ω , and the complex-valued process, $y_{t1}(\omega) = \mathbf{e}_1(\omega)^* \mathbf{x}_t$, is called the first principal component at frequency ω . The k -th principal component at frequency ω , for $k = 1, 2, \dots, p$, is the complex-valued time series $y_{tk}(\omega) = \mathbf{e}_k(\omega)^* \mathbf{x}_t$, in analogy to the classical case. In this case, the spectral density of $y_{tk}(\omega)$ at frequency ω is $f_{y_k}(\omega) = \mathbf{e}_k(\omega)^* f_{xx}(\omega) \mathbf{e}_k(\omega) = \lambda_k(\omega)$.

The previous development of spectral domain principal components is related to the spectral envelope methodology first discussed in Stoffer et al. (1993). We will present the spectral envelope in the next section, where we motivate the use of principal components as it is presented above. Another way to motivate the use of principal components in the frequency domain was given in Brillinger (1981, Chapter 9). Although this technique leads to the same analysis, the motivation may be more satisfactory to the reader at this point. In this case, we suppose we have a stationary, p -dimensional, vector-valued process $\mathbf{x}_t$ and we are only able to keep a univariate process y_t such that, when needed, we may reconstruct the vector-valued process, $\mathbf{x}_t$, according to an optimality criterion.

Specifically, we suppose we want to approximate a mean-zero, stationary, vector-valued time series, $\mathbf{x}_t$, with spectral matrix $f_{xx}(\omega)$, by a univariate process y_t defined by

$$y_t = \sum_{j=-\infty}^{\infty} \mathbf{c}_{t-j}^* \mathbf{x}_j, \quad (7.144)$$

where $\{\mathbf{c}_j\}$ is a $p \times 1$ vector-valued filter, such that $\{\mathbf{c}_j\}$ is absolutely summable; that is, $\sum_{j=-\infty}^{\infty} |\mathbf{c}_j| < \infty$. The approximation is accomplished so the reconstruction of $\mathbf{x}_t$ from y_t , say,

$$\hat{\mathbf{x}}_t = \sum_{j=-\infty}^{\infty} \mathbf{b}_{t-j} y_j, \quad (7.145)$$

where $\{\mathbf{b}_j\}$ is an absolutely summable $p \times 1$ filter, is such that the mean square approximation error

$$E\{(\mathbf{x}_t - \hat{\mathbf{x}}_t)^* (\mathbf{x}_t - \hat{\mathbf{x}}_t)\} \quad (7.146)$$

is minimized.

Let $\mathbf{b}(\omega)$ and $\mathbf{c}(\omega)$ be the transforms of $\{\mathbf{b}_j\}$ and $\{\mathbf{c}_j\}$, respectively. For example,

$$\mathbf{c}(\omega) = \sum_{j=-\infty}^{\infty} \mathbf{c}_j \exp(-2\pi i j \omega), \quad (7.147)$$

and, consequently,

$$\mathbf{c}_j = \int_{-1/2}^{1/2} \mathbf{c}(\omega) \exp(2\pi i j \omega) d\omega. \quad (7.148)$$

Brillinger (1981, Theorem 9.3.1) shows the solution to the problem is to choose $\mathbf{c}(\omega)$ to satisfy (7.142) and to set $\mathbf{b}(\omega) = \overline{\mathbf{c}(\omega)}$. This is precisely the previous problem, with the solution given by (7.143). That is, we choose $\mathbf{c}(\omega) = \mathbf{e}_1(\omega)$ and $\mathbf{b}(\omega) = \overline{\mathbf{e}_1(\omega)}$; the filter values can be obtained via the inversion formula given by (7.148). Using these results, in view of (7.144), we may form the first principal component series, say y_{t1} .

This technique may be extended by requesting another series, say, y_{t2} , for approximating $\mathbf{x}_t$ with respect to minimum mean square error, but where the coherency between y_{t2} and y_{t1} is zero. In this case, we choose $\mathbf{c}(\omega) = \mathbf{e}_2(\omega)$. Continuing this way, we can obtain the first $q \leq p$ principal components series, say, $\mathbf{y}_t = (y_{t1}, \dots, y_{tq})'$, having spectral density $f_q(\omega) = \text{diag}\{\lambda_1(\omega), \dots, \lambda_q(\omega)\}$. The series y_{tk} is the k -th principal component series.

As in the classical case, given observations, $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n$, from the process $\mathbf{x}_t$, we can form an estimate $\widehat{f}_{xx}(\omega)$ of $f_{xx}(\omega)$ and define the sample principal component series by replacing $f_{xx}(\omega)$ with $\widehat{f}_{xx}(\omega)$ in the previous discussion. Precise details pertaining to the asymptotic ($n \rightarrow \infty$) behavior of the principal component series and their spectra can be found in Brillinger (1981, Chapter 9). To give a basic idea of what we can expect, we focus on the first principal component series and on the spectral estimator obtained by smoothing the periodogram matrix, $I_n(\omega_j)$; that is

$$\widehat{f}_{xx}(\omega_j) = \sum_{\ell=-m}^m h_{\ell} I_n(\omega_j + \ell/n), \quad (7.149)$$

where $L = 2m + 1$ is odd and the weights are chosen so $h_{\ell} = h_{-\ell}$ are positive and $\sum_{\ell} h_{\ell} = 1$. Under the conditions for which $\widehat{f}_{xx}(\omega_j)$ is a well-behaved estimator of $f_{xx}(\omega_j)$, and for which the largest eigenvalue of $f_{xx}(\omega_j)$ is unique,

$$\left\{ \eta_n \frac{\widehat{\lambda}_1(\omega_j) - \lambda_1(\omega_j)}{\lambda_1(\omega_j)}; \eta_n [\widehat{\mathbf{e}}_1(\omega_j) - \mathbf{e}_1(\omega_j)]; j = 1, \dots, J \right\} \quad (7.150)$$

converges ($n \rightarrow \infty$) jointly in distribution to independent, zero-mean normal distributions, the first of which is standard normal. In (7.150), $\eta_n^{-2} = \sum_{\ell=-m}^m h_{\ell}^2$, noting we must have $L \rightarrow \infty$ and $\eta_n \rightarrow \infty$, but $L/n \rightarrow 0$ as $n \rightarrow \infty$. The asymptotic variance-covariance matrix of $\widehat{\mathbf{e}}_1(\omega)$, say, $\Sigma_{e_1}(\omega)$, is given by

$$\Sigma_{e_1}(\omega) = \eta_n^{-2} \lambda_1(\omega) \sum_{\ell=2}^p \lambda_{\ell}(\omega) \{\lambda_1(\omega) - \lambda_{\ell}(\omega)\}^{-2} \mathbf{e}_{\ell}(\omega) \mathbf{e}_{\ell}^*(\omega). \quad (7.151)$$




Fig. 7.17. The individual periodograms of x_{tk} , for $k = 1, \dots, 8$, in [Example 7.14](#).

The distribution of $\widehat{\mathbf{e}}_1(\omega)$ depends on the other latent roots and vectors of $f_x(\omega)$. Writing $\widehat{\mathbf{e}}_1(\omega) = (\widehat{e}_{11}(\omega), \widehat{e}_{12}(\omega), \dots, \widehat{e}_{1p}(\omega))'$, we may use this result to form confidence regions for the components of $\widehat{\mathbf{e}}_1$ by approximating the distribution of

$$\frac{2|\widehat{\mathbf{e}}_{1,j}(\omega) - \mathbf{e}_{1,j}(\omega)|^2}{s_j^2(\omega)}, \quad (7.152)$$

for $j = 1, \dots, p$, by a χ^2 distribution with two degrees of freedom. In (7.152), $s_j^2(\omega)$ is the j -th diagonal element of $\widehat{\Sigma}_{e_1}(\omega)$, the estimate of $\Sigma_{e_1}(\omega)$. We can use (7.152) to check whether the value of zero is in the confidence region by comparing $2|\widehat{\mathbf{e}}_{1,j}(\omega)|^2/s_j^2(\omega)$ with $\chi_2^2(1 - \alpha)$, the $1 - \alpha$ upper tail cutoff of the χ_2^2 distribution.

Example 7.14 Principal Component Analysis of the fMRI Data

Recall [Example 1.6](#) of Chapter 1, where the vector time series $\mathbf{x}_t = (x_{t1}, \dots, x_{t8})'$, $t = 1, \dots, 128$, represents consecutive measures of average blood oxygenation level dependent (BOLD) signal intensity, which measures areas of activation in the brain. Recall subjects were given a non-painful brush on the hand and the stimulus was applied for 32 seconds and then stopped for 32 seconds; thus, the signal period is 64 seconds (the sampling rate was one observation every two seconds for 256 seconds). The series x_{tk} for $k = 1, 2, 3, 4$ represent locations in cortex, series x_{t5} and x_{t6} represent locations in the thalamus, and x_{t7} and x_{t8} represent locations in the cerebellum.

As is evident from [Figure 1.6](#) in Chapter 1, different areas of the brain are responding differently, and a principal component analysis may help in indicating which locations are responding with the most spectral power, and which locations do not contribute to the spectral power at the stimulus signal period. In this analysis,




Fig. 7.18. The estimated spectral density, $\hat{\lambda}_1(j/128)$, of the first principal component series in Example 7.14.

Table 7.5. Magnitudes of the PC Vector at the Stimulus Frequency

Location	1	2	3	4	5	6	7	8
$ \hat{\mathbf{e}}_1(\frac{4}{128}) $	.64	.36	.36	.22	.32	.05*	.13	.39

*Zero is in an approximate 99% confidence region for this component.

we will focus primarily on the signal period of 64 seconds, which translates to four cycles in 256 seconds or $\omega = 4/128$ cycles per time point.

Figure 7.17 shows individual periodograms of the series x_{tk} for $k = 1, \dots, 8$. As was evident from Figure 1.6, a strong response to the brush stimulus occurred in areas of the cortex. To estimate the spectral density of $\mathbf{x}_t$, we used (7.149) with $L = 5$ and $\{h_0 = 3/9, h_{\pm 1} = 2/9, h_{\pm 2} = 1/9\}$; this is a Daniell kernel with $m = 1$ passed twice. Calling the estimated spectrum $\hat{f}_{xx}(j/128)$, for $j = 0, 1, \dots, 64$, we can obtain the estimated spectrum of the first principal component series y_{t1} by calculating the largest eigenvalue, $\hat{\lambda}_1(j/128)$, of $\hat{f}_{xx}(j/128)$ for each $j = 0, 1, \dots, 64$. The result, $\hat{\lambda}_1(j/128)$, is shown in Figure 7.18. As expected, there is a large peak at the stimulus frequency $4/128$, wherein $\hat{\lambda}_1(4/128) = 2$. The total power at the stimulus frequency is $\text{tr}(\hat{f}_{xx}(4/128)) = 2.05$, so the proportion of the power at frequency $4/128$ attributed to the first principal component series is about $2/2.05$ or roughly 98%. Because the first principal component explains nearly all of the total power at the stimulus frequency, there is no need to explore the other principal component series at this frequency.

The estimated first principal component series at frequency $4/128$ is given by $\hat{y}_{t1}(4/128) = \hat{\mathbf{e}}_1^*(4/128)\mathbf{x}_t$, and the components of $\hat{\mathbf{e}}_1(4/128)$ can give insight as to which locations of the brain are responding to the brush stimulus. Table 7.5 shows

the magnitudes of $\widehat{\epsilon}_1(4/128)$. In addition, an approximate 99% confidence interval was obtained for each component using (7.152). As expected, the analysis indicates that location 6 is not contributing to the power at this frequency, but surprisingly, the analysis suggests location 5 (cerebellum 1) is responding to the stimulus.

The R code for this example is as follows.

```
n = 128; Per = abs(mvfft(fmri1[,-1]))^2/n
par(mfrow=c(2,4),mar=c(3,2,2,1),mgp = c(1.6,.6,0),oma=c(0,1,0,0))
for (i in 1:8) plot(0:20, Per[1:21,i], type="l", ylim=c(0,8),
  main=colnames(fmri1)[i+1], xlab="Cycles",ylab="",xaxp=c(0,20,5))
mtext("Periodogram", side=2, line=-.3, outer=TRUE, adj=c(.2,.8))
fxx = mvspec(fmri1[,-1], kernel("daniell", c(1,1)), taper=.5)$fxx
l.val = rep(NA,64)
for (k in 1:64) {
  u = eigen(fxx[, ,k], symmetric=TRUE, only.values = TRUE)
  l.val[k] = u$values[1] # largest e-value
  dev.new()
  plot(l.val, type="l",xaxp=c(0,64,8), xlab="Cycles (Frequency x 128)",
    ylab="First Principal Component")
  axis(1, seq(4,60,by=8), labels=FALSE)
  # At freq 4/128
  u = eigen(fxx[, ,4], symmetric=TRUE)
  lam=u$values; evec=u$vectors
  lam[1]/sum(lam) # % of variance explained
  sig.e1 = matrix(0,8,8)
  for (l in 2:5){ # last 3 evs are 0
    sig.e1= sig.e1 + lam[l]*evec[,l]%*%Conj(t(evec[,l]))/(lam[1]-lam[l])^2}
  sig.e1 = Re(sig.e1)*lam[1]*sum(kernel("daniell", c(1,1))$coef^2)
  p.val = round(pchisq(2*abs(evec[,1])^2/diag(sig.e1), 2, lower.tail=FALSE), 3)
  cbind(colnames(fmri1)[-1], abs(evec[,1]), p.val) # table values
```

FACTOR ANALYSIS

Classical factor analysis is similar to classical principal component analysis. Suppose $\mathbf{x}$ is a mean-zero, $p \times 1$, random vector with variance–covariance matrix Σ_{xx} . The factor model proposes that $\mathbf{x}$ is dependent on a few unobserved common factors, $z_1, \dots, z_q$, plus error. In this model, one hopes that q will be much smaller than p . The factor model is given by

$$\mathbf{x} = \mathcal{B}\mathbf{z} + \boldsymbol{\epsilon}, \quad (7.153)$$

where $\mathcal{B}$ is a $p \times q$ matrix of factor loadings, $\mathbf{z} = (z_1, \dots, z_q)'$ is a random $q \times 1$ vector of factors such that $E(\mathbf{z}) = \mathbf{0}$ and $E(\mathbf{z}\mathbf{z}') = I_q$, the $q \times q$ identity matrix. The $p \times 1$ unobserved error vector $\boldsymbol{\epsilon}$ is assumed to be independent of the factors, with zero mean and diagonal variance-covariance matrix $D = \text{diag}\{\delta_1^2, \dots, \delta_p^2\}$. Note, (7.153) differs from the multivariate regression model in §5.7 because the factors, $\mathbf{z}$, are unobserved. Equivalently, the factor model, (7.153), can be written in terms of the covariance structure of $\mathbf{x}$,

$$\Sigma_{xx} = \mathcal{B}\mathcal{B}' + D; \quad (7.154)$$

i.e., the variance-covariance matrix of $\mathbf{x}$ is the sum of a symmetric, nonnegative-definite rank $q \leq p$ matrix and a nonnegative-definite diagonal matrix. If $q = p$, then

Σ_{xx} can be reproduced exactly as $\mathcal{B}\mathcal{B}'$, using the fact that $\Sigma_{xx} = \lambda_1 \mathbf{e}_1 \mathbf{e}_1' + \cdots + \lambda_p \mathbf{e}_p \mathbf{e}_p'$, where $(\lambda_i, \mathbf{e}_i)$ are the eigenvalue–eigenvector pairs of Σ_{xx} . As previously indicated, however, we hope q will be much smaller than p . Unfortunately, most covariance matrices cannot be factored as (7.154) when q is much smaller than p .

To motivate factor analysis, suppose the components of $\mathbf{x}$ can be grouped into meaningful groups. Within each group, the components are highly correlated, but the correlation between variables that are not in the same group is small. A group is supposedly formed by a single construct, represented as an unobservable factor, responsible for the high correlations within a group. For example, a person competing in a decathlon performs $p = 10$ athletic events, and we may represent the outcome of the decathlon as a 10×1 vector of scores. The events in a decathlon involve running, jumping, or throwing, and it is conceivable the 10×1 vector of scores might be able to be factored into $q = 4$ factors, (1) arm strength, (2) leg strength, (3) running speed, and (4) running endurance. The model (7.153) specifies that $\text{cov}(\mathbf{x}, \mathbf{z}) = \mathcal{B}$, or $\text{cov}(x_i, z_j) = b_{ij}$ where b_{ij} is the ij -th component of the factor loading matrix $\mathcal{B}$, for $i = 1, \dots, p$ and $j = 1, \dots, q$. Thus, the elements of $\mathcal{B}$ are used to identify which hypothetical factors the components of $\mathbf{x}$ belong to, or load on.

At this point, some ambiguity is still associated with the factor model. Let $\mathcal{Q}$ be a $q \times q$ orthogonal matrix; that is $\mathcal{Q}'\mathcal{Q} = \mathcal{Q}\mathcal{Q}' = I_q$. Let $\mathcal{B}_* = \mathcal{B}\mathcal{Q}$ and $\mathbf{z}_* = \mathcal{Q}'\mathbf{z}$ so (7.153) can be written as

$$\mathbf{x} = \mathcal{B}\mathbf{z} + \boldsymbol{\epsilon} = \mathcal{B}\mathcal{Q}\mathcal{Q}'\mathbf{z} + \boldsymbol{\epsilon} = \mathcal{B}_*\mathbf{z}_* + \boldsymbol{\epsilon}. \quad (7.155)$$

The model in terms of $\mathcal{B}_*$ and $\mathbf{z}_*$ fulfills all of the factor model requirements, for example, $\text{cov}(\mathbf{z}_*) = \mathcal{Q}'\text{cov}(\mathbf{z})\mathcal{Q} = \mathcal{Q}\mathcal{Q}' = I_q$, so

$$\Sigma_{xx} = \mathcal{B}_*\text{cov}(\mathbf{z}_*)\mathcal{B}_*' + D = \mathcal{B}\mathcal{Q}\mathcal{Q}'\mathcal{B}' + D = \mathcal{B}\mathcal{B}' + D. \quad (7.156)$$

Hence, on the basis of observations on $\mathbf{x}$, we cannot distinguish between the loadings $\mathcal{B}$ and the rotated loadings $\mathcal{B}_* = \mathcal{B}\mathcal{Q}$. Typically, $\mathcal{Q}$ is chosen so the matrix $\mathcal{B}$ is easy to interpret, and this is the basis of what is called factor rotation.

Given a sample $\mathbf{x}_1, \dots, \mathbf{x}_n$, a number of methods are used to estimate the parameters of the factor model, and we discuss two of them here. The first method is the principal component method. Let S_{xx} denote the sample variance–covariance matrix, and let $(\hat{\lambda}_i, \hat{\mathbf{e}}_i)$ be the eigenvalue–eigenvector pairs of S_{xx} . The $p \times q$ matrix of estimated factor loadings is found by setting

$$\hat{\mathcal{B}} = [\hat{\lambda}_1^{1/2} \hat{\mathbf{e}}_1 \mid \hat{\lambda}_2^{1/2} \hat{\mathbf{e}}_2 \mid \cdots \mid \hat{\lambda}_q^{1/2} \hat{\mathbf{e}}_q]. \quad (7.157)$$

The argument here is that if q factors exist, then

$$S_{xx} \approx \hat{\lambda}_1 \hat{\mathbf{e}}_1 \hat{\mathbf{e}}_1' + \cdots + \hat{\lambda}_q \hat{\mathbf{e}}_q \hat{\mathbf{e}}_q' = \hat{\mathcal{B}}\hat{\mathcal{B}}', \quad (7.158)$$

because the remaining eigenvalues, $\hat{\lambda}_{q+1}, \dots, \hat{\lambda}_p$, will be negligible. The estimated diagonal matrix of error variances is then obtained by setting $\hat{D} = \text{diag}\{\hat{\delta}_1^2, \dots, \hat{\delta}_p^2\}$, where $\hat{\delta}_j^2$ is the j -th diagonal element of $S_{xx} - \hat{\mathcal{B}}\hat{\mathcal{B}}'$.

The second method, which can give answers that are considerably different from the principal component method is maximum likelihood. Upon further assumption that in (7.153), $\mathbf{z}$ and ϵ are multivariate normal, the log likelihood of $\mathcal{B}$ and D ignoring a constant is

$$-2 \ln L(\mathcal{B}, D) = n \ln |\Sigma_{xx}| + \sum_{j=1}^n \mathbf{x}'_j \Sigma_{xx}^{-1} \mathbf{x}_j. \quad (7.159)$$

The likelihood depends on $\mathcal{B}$ and D through (7.154), $\Sigma_{xx} = \mathcal{B}\mathcal{B}' + D$. As discussed in (7.155)-(7.156), the likelihood is not well defined because $\mathcal{B}$ can be rotated. Typically, restricting $\mathcal{B}D^{-1}\mathcal{B}'$ to be a diagonal matrix is a computationally convenient uniqueness condition. The actual maximization of the likelihood is accomplished using numerical methods.

One obvious method of performing maximum likelihood for the Gaussian factor model is the EM algorithm. For example, suppose the factor vector $\mathbf{z}$ is known. Then, the factor model is simply the multivariate regression model given in §5.7, that is, write $X' = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n]$ and $Z' = [\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_n]$, and note that X is $n \times p$ and Z is $n \times q$. Then, the MLE of $\mathcal{B}$ is

$$\hat{\mathcal{B}} = X'Z(Z'Z)^{-1} = \left(n^{-1} \sum_{j=1}^n \mathbf{x}_j \mathbf{z}'_j\right) \left(n^{-1} \sum_{j=1}^n \mathbf{z}_j \mathbf{z}'_j\right)^{-1} \stackrel{\text{def}}{=} C_{xz} C_{zz}^{-1} \quad (7.160)$$

and the MLE of D is

$$\hat{D} = \text{diag} \left\{ n^{-1} \sum_{j=1}^n (\mathbf{x}_j - \hat{\mathcal{B}} \mathbf{z}_j) (\mathbf{x}_j - \hat{\mathcal{B}} \mathbf{z}_j)' \right\}; \quad (7.161)$$

that is, only the diagonal elements of the right-hand side of (7.161) are used. The bracketed quantity in (7.161) reduces to

$$C_{xx} - C_{xz} C_{zz}^{-1} C'_{xz}, \quad (7.162)$$

where $C_{xx} = n^{-1} \sum_{j=1}^n \mathbf{x}_j \mathbf{x}'_j$.

Based on the derivation of the EM algorithm for the state-space model, (4.66)-(4.75), we conclude that, to employ the EM algorithm here, given the current parameter estimates, in C_{xz} , we replace $\mathbf{x}_j \mathbf{z}'_j$ by $\mathbf{x}_j \tilde{\mathbf{z}}'_j$, where $\tilde{\mathbf{z}}_j = E(\mathbf{z}_j | \mathbf{x}_j)$, and in C_{zz} , we replace $\mathbf{z}_j \mathbf{z}'_j$ by $P_z + \tilde{\mathbf{z}}_j \tilde{\mathbf{z}}'_j$, where $P_z = \text{var}(\mathbf{z}_j | \mathbf{x}_j)$. Using the fact that the $(p+q) \times 1$ vector $(\mathbf{x}'_j, \mathbf{z}'_j)'$ is multivariate normal with mean-zero, and variance-covariance matrix given by

$$\begin{pmatrix} \mathcal{B}\mathcal{B}' + D & \mathcal{B} \\ \mathcal{B}' & I_q \end{pmatrix}, \quad (7.163)$$

we have

$$\tilde{\mathbf{z}}_j \equiv E(\mathbf{z}_j | \mathbf{x}_j) = \mathcal{B}'(\mathcal{B}'\mathcal{B} + D)^{-1} \mathbf{x}_j \quad (7.164)$$

and

$$P_z \equiv \text{var}(\mathbf{z}_j | \mathbf{x}_j) = I_q - \mathcal{B}'(\mathcal{B}'\mathcal{B} + D)^{-1} \mathcal{B}. \quad (7.165)$$

For time series, suppose $\mathbf{x}_t$ is a stationary $p \times 1$ process with $p \times p$ spectral matrix $f_{xx}(\omega)$. Analogous to the classical model displayed in (7.154), we may postulate that at a given frequency of interest, ω , the spectral matrix of $\mathbf{x}_t$ satisfies

$$f_{xx}(\omega) = \mathcal{B}(\omega)\mathcal{B}(\omega)^* + D(\omega), \quad (7.166)$$

where $\mathcal{B}(\omega)$ is a complex-valued $p \times q$ matrix with $\text{rank}(\mathcal{B}(\omega)) = q \leq p$ and $D(\omega)$ is a real, nonnegative-definite, diagonal matrix. Typically, we expect q will be much smaller than p .

As an example of a model that gives rise to (7.166), let $\mathbf{x}_t = (x_{t1}, \dots, x_{tp})'$, and suppose

$$x_{tj} = c_j s_{t-\tau_j} + \epsilon_{tj}, \quad j = 1, \dots, p, \quad (7.167)$$

where $c_j \geq 0$ are individual amplitudes and s_t is a common unobserved signal (factor) with spectral density $f_{ss}(\omega)$. The values τ_j are the individual phase shifts. Assume s_t is independent of $\boldsymbol{\epsilon}_t = (\epsilon_{t1}, \dots, \epsilon_{tp})'$ and the spectral matrix of $\boldsymbol{\epsilon}_t$, $D_{\epsilon\epsilon}(\omega)$, is diagonal. The DFT of x_{tj} is given by

$$X_j(\omega) = n^{-1/2} \sum_{t=1}^n x_{tj} \exp(-2\pi i t \omega)$$

and, in terms of the model (7.167),

$$X_j(\omega) = a_j(\omega)X_s(\omega) + X_{\epsilon_j}(\omega), \quad (7.168)$$

where $a_j(\omega) = c_j \exp(-2\pi i \tau_j \omega)$, and $X_s(\omega)$ and $X_{\epsilon_j}(\omega)$ are the respective DFTs of the signal s_t and the noise ϵ_{tj} . Stacking the individual elements of (7.168), we obtain a complex version of the classical factor model with one factor,

$$\begin{pmatrix} X_1(\omega) \\ \vdots \\ X_p(\omega) \end{pmatrix} = \begin{pmatrix} a_1(\omega) \\ \vdots \\ a_p(\omega) \end{pmatrix} X_s(\omega) + \begin{pmatrix} X_{\epsilon_1}(\omega) \\ \vdots \\ X_{\epsilon_p}(\omega) \end{pmatrix},$$

or more succinctly,

$$\mathbf{X}(\omega) = \mathbf{a}(\omega)X_s(\omega) + \mathbf{X}_{\epsilon}(\omega). \quad (7.169)$$

From (7.169), we can identify the spectral components of the model; that is,

$$f_{xx}(\omega) = \mathbf{b}(\omega)\mathbf{b}(\omega)^* + D_{\epsilon\epsilon}(\omega), \quad (7.170)$$

where $\mathbf{b}(\omega)$ is a $p \times 1$ complex-valued vector, $\mathbf{b}(\omega)\mathbf{b}(\omega)^* = \mathbf{a}(\omega)f_{ss}(\omega)\mathbf{a}(\omega)^*$. Model (7.170) could be considered the one-factor model for time series. This model can be extended to more than one factor by adding other independent signals into the original model (7.167). More details regarding this and related models can be found in Stoffer (1999).

Example 7.15 Single Factor Analysis of the fMRI Data

The fMRI data analyzed in [Example 7.14](#) is well suited for a single factor analysis using the model (7.167), or, equivalently, the complex-valued, single factor model (7.169). In terms of (7.167), we can think of the signal s_t as representing the brush stimulus signal. As before, the frequency of interest is $\omega = 4/128$, which corresponds to a period of 32 time points, or 64 seconds.

A simple way to estimate the components $\mathbf{b}(\omega)$ and $D_{\epsilon\epsilon}(\omega)$, as specified in (7.170), is to use the principal components method. Let $\hat{f}_{xx}(\omega)$ denote the estimate of the spectral density of $\mathbf{x}_t = (x_{t1}, \dots, x_{t8})'$ obtained in [Example 7.14](#). Then, analogous to (7.157) and (7.158), we set

$$\hat{\mathbf{b}}(\omega) = \sqrt{\hat{\lambda}_1(\omega)} \hat{\mathbf{e}}_1(\omega),$$

where $(\hat{\lambda}_1(\omega), \hat{\mathbf{e}}_1(\omega))$ is the first eigenvalue–eigenvector pair of $\hat{f}_{xx}(\omega)$. The diagonal elements of $\hat{D}_{\epsilon\epsilon}(\omega)$ are obtained from the diagonal elements of $\hat{f}_{xx}(\omega) - \hat{\mathbf{b}}(\omega)\hat{\mathbf{b}}(\omega)^*$. The appropriateness of the model can be assessed by checking the elements of the residual matrix, $\hat{f}_{xx}(\omega) - [\hat{\mathbf{b}}(\omega)\hat{\mathbf{b}}(\omega)^* + \hat{D}_{\epsilon\epsilon}(\omega)]$, are negligible in magnitude.

Concentrating on the stimulus frequency, recall $\hat{\lambda}_1(4/128) = 2$. The magnitudes of $\hat{\mathbf{e}}_1(4/128)$ are displayed in [Table 7.5](#), indicating all locations load on the stimulus factor except for location 6, and location 7 could be considered borderline. The diagonal elements of $\hat{f}_{xx}(\omega) - \hat{\mathbf{b}}(\omega)\hat{\mathbf{b}}(\omega)^*$ yield

$$\hat{D}_{\epsilon\epsilon}(4/128) = 0.001 \times \text{diag}\{1.36, 2.04, 6.22, 11.30, 0.73, 13.26, 6.93, 5.88\}.$$

The magnitudes of the elements of the residual matrix at $\omega = 4/128$ are

$$0.001 \times \begin{bmatrix} 0.00 & 1.73 & 3.88 & 3.61 & 0.88 & 2.04 & 1.60 & 2.81 \\ 2.41 & 0.00 & 1.17 & 3.77 & 1.49 & 5.58 & 3.68 & 4.21 \\ 8.49 & 5.34 & 0.00 & 2.94 & 7.58 & 10.91 & 8.36 & 10.64 \\ 12.65 & 11.84 & 6.12 & 0.00 & 12.56 & 14.64 & 13.34 & 16.10 \\ 0.32 & 0.29 & 2.10 & 2.01 & 0.00 & 1.18 & 2.01 & 1.18 \\ 10.34 & 16.69 & 17.09 & 15.94 & 13.49 & 0.00 & 5.78 & 14.74 \\ 5.71 & 8.51 & 8.94 & 10.18 & 7.56 & 0.97 & 0.00 & 8.66 \\ 6.25 & 8.00 & 10.31 & 10.69 & 5.95 & 8.69 & 7.64 & 0.00 \end{bmatrix},$$

indicating the model fit is good. Assuming the results of the previous example are available, use the following R code.

```
bhat = sqrt(lam[1])*evec[,1]
Dhat = Re(diag(fxx[, ,4] - bhat%*%Conj(t(bhat))))
res = Mod(fxx[, ,4] - Dhat - bhat%*%Conj(t(bhat)))
```

A number of authors have considered factor analysis in the spectral domain, for example Priestley et al. (1974); Priestley and Subba Rao (1975); Geweke (1977), and Geweke and Singleton (1981), to mention a few. An obvious extension of simple model (7.167) is the factor model

$$\mathbf{x}_t = \sum_{j=-\infty}^{\infty} \Lambda_j \mathbf{s}_{t-j} + \boldsymbol{\epsilon}_t, \quad (7.171)$$

where $\{\Lambda_j\}$ is a real-valued $p \times q$ filter, $\mathbf{s}_t$ is a $q \times 1$ stationary, unobserved signal, with independent components, and $\boldsymbol{\epsilon}_t$ is white noise. We assume the signal and noise process are independent, $\mathbf{s}_t$ has $q \times q$ real, diagonal spectral matrix $f_{ss}(\omega) = \text{diag}\{f_{s1}(\omega), \dots, f_{sq}(\omega)\}$, and $\boldsymbol{\epsilon}_t$ has a real, diagonal, $p \times p$ spectral matrix given by $D_{\epsilon\epsilon}(\omega) = \text{diag}\{f_{\epsilon 1}(\omega), \dots, f_{\epsilon p}(\omega)\}$. If, in addition, $\sum \|\Lambda_j\| < \infty$, the spectral matrix of $\mathbf{x}_t$ can be written as

$$f_{xx}(\omega) = \Lambda(\omega) f_{ss}(\omega) \Lambda(\omega)^* + D_{\epsilon\epsilon}(\omega) = \mathcal{B}(\omega) \mathcal{B}(\omega)^* + D_{\epsilon\epsilon}(\omega), \quad (7.172)$$

where

$$\Lambda(\omega) = \sum_{t=-\infty}^{\infty} \Lambda_t \exp(-2\pi i t \omega) \quad (7.173)$$

and $\mathcal{B}(\omega) = \Lambda(\omega) f_{ss}^{1/2}(\omega)$. Thus, by (7.172), the model (7.171) is seen to satisfy the basic requirement of the spectral domain factor analysis model; that is, the $p \times p$ spectral density matrix of the process of interest, $f_{xx}(\omega)$, is the sum of a rank $q \leq p$ matrix, $\mathcal{B}(\omega) \mathcal{B}(\omega)^*$, and a real, diagonal matrix, $D_{\epsilon\epsilon}(\omega)$. For the purpose of identifiability we set $f_{ss}(\omega) = I_q$ for all ω ; in which case, $\mathcal{B}(\omega) = \Lambda(\omega)$. As in the classical case [see (7.156)], the model is specified only up to rotations; for details, see Bloomfield and Davis (1994).

Parameter estimation for the model (7.171), or equivalently (7.172), can be accomplished using the principal component method. Let $\hat{f}_{xx}(\omega)$ be an estimate of $f_{xx}(\omega)$, and let $(\hat{\lambda}_j(\omega), \hat{\mathbf{e}}_j(\omega))$, for $j = 1, \dots, p$, be the eigenvalue–eigenvector pairs, in the usual order, of $\hat{f}_{xx}(\omega)$. Then, as in the classical case, the $p \times q$ matrix $\mathcal{B}$ is estimated by

$$\hat{\mathcal{B}}(\omega) = [\hat{\lambda}_1(\omega)^{1/2} \hat{\mathbf{e}}_1(\omega) \mid \hat{\lambda}_2(\omega)^{1/2} \hat{\mathbf{e}}_2(\omega) \mid \dots \mid \hat{\lambda}_q(\omega)^{1/2} \hat{\mathbf{e}}_q(\omega)]. \quad (7.174)$$

The estimated diagonal spectral density matrix of errors is then obtained by setting $\hat{D}_{\epsilon\epsilon}(\omega) = \text{diag}\{\hat{f}_{\epsilon 1}(\omega), \dots, \hat{f}_{\epsilon p}(\omega)\}$, where $\hat{f}_{\epsilon j}(\omega)$ is the j -th diagonal element of $\hat{f}_{xx}(\omega) - \hat{\mathcal{B}}(\omega) \hat{\mathcal{B}}(\omega)^*$.

Alternatively, we can estimate the parameters by approximate likelihood methods. As in (7.169), let $\mathbf{X}(\omega_j)$ denote the DFT of the data $\mathbf{x}_1, \dots, \mathbf{x}_n$ at frequency $\omega_j = j/n$. Similarly, let $\mathbf{X}_s(\omega_j)$ and $\mathbf{X}_\epsilon(\omega_j)$ be the DFTs of the signal and of the noise processes, respectively. Then, under certain conditions (see Pawitan and Shumway, 1989), for $\ell = 0, \pm 1, \dots, \pm m$,

$$\mathbf{X}(\omega_j + \ell/n) = \Lambda(\omega_j) \mathbf{X}_s(\omega_j + \ell/n) + \mathbf{X}_\epsilon(\omega_j + \ell/n) + o_{as}(n^{-\alpha}), \quad (7.175)$$

where $\Lambda(\omega_j)$ is given by (7.173) and $o_{as}(n^{-\alpha}) \rightarrow 0$ almost surely for some $0 \leq \alpha < 1/2$ as $n \rightarrow \infty$. In (7.175), the $\mathbf{X}(\omega_j + \ell/n)$ are the DFTs of the data at the L odd frequencies $\{\omega_j + \ell/n; \ell = 0, \pm 1, \dots, \pm m\}$ surrounding the central frequency of interest $\omega_j = j/n$.

Under appropriate conditions $\{\mathbf{X}(\omega_j + \ell/n); \ell = 0, \pm 1, \dots, \pm m\}$ in (7.175) are approximately ($n \rightarrow \infty$) independent, complex Gaussian random vectors with variance–covariance matrix $f_{xx}(\omega_j)$. The approximate likelihood is given by

$$\begin{aligned} & -2 \ln L(\mathcal{B}(\omega_j), D_{\epsilon\epsilon}(\omega_j)) \\ &= n \ln |f_{xx}(\omega_j)| + \sum_{\ell=-m}^m \mathbf{X}^*(\omega_j + \ell/n) f_{xx}^{-1}(\omega_j) \mathbf{X}(\omega_j + \ell/n), \end{aligned} \quad (7.176)$$

with the constraint $f_{xx}(\omega_j) = \mathcal{B}(\omega_j)\mathcal{B}(\omega_j)^* + D_{\epsilon\epsilon}(\omega_j)$. As in the classical case, we can use various numerical methods to maximize $L(\mathcal{B}(\omega_j), D_{\epsilon\epsilon}(\omega_j))$ at every frequency, ω_j , of interest. For example, the EM algorithm discussed for the classical case, (7.160)–(7.165), can easily be extended to this case.

Assuming $f_{ss}(\omega) = I_q$, the estimate of $\mathcal{B}(\omega_j)$ is also the estimate of $\Lambda(\omega_j)$. Calling this estimate $\widehat{\Lambda}(\omega_j)$, the time domain filter can be estimated by

$$\widehat{\Lambda}_t^M = M^{-1} \sum_{j=0}^{M-1} \widehat{\Lambda}(\omega_j) \exp(2\pi i j t / n), \quad (7.177)$$

for some $0 < M \leq n$, which is the discrete and finite version of the inversion formula given by

$$\Lambda_t = \int_{-1/2}^{1/2} \Lambda(\omega) \exp(2\pi i \omega t) d\omega. \quad (7.178)$$

Note that we have used this approximation earlier in Chapter 4, (4.120), for estimating the time response of a frequency response function defined over a finite number of frequencies.

Example 7.16 Government Spending, Private Investment, and Unemployment

Figure 7.19 shows the seasonally adjusted, quarterly growth rate (as percentages) of five macroeconomic series, unemployment, GNP, consumption, government investment, and private investment in the United States between 1948 and 1988, $n = 160$ values. These data are analyzed in the time domain by Young and Pedregal (1998), who were investigating how government spending and private capital investment influenced the rate of unemployment.

Spectral estimation was performed on the detrended, standardized, and tapered growth rate values; see the R code at the end of this example for details. Figure 7.20 shows the individual estimated spectra of each series. We focus on three interesting frequencies. First, we note the lack of spectral power near the annual cycle ($\omega = 1$, or one cycle every four quarters), indicating the data have been seasonally adjusted. In addition, because of the seasonal adjustment, some spectral power appears near the seasonal frequency; this is a distortion apparently caused by the method of seasonally adjusting the data. Next, we note the spectral power near $\omega = .25$, or one cycle every four years, in unemployment, GNP, consumption, and, to lesser degree, in private investment. Finally, spectral power appears near $\omega = .125$, or one




Fig. 7.19. The seasonally adjusted, quarterly growth rate (as percentages) of five macroeconomic series, unemployment, GNP, consumption, government investment, and private investment in the United States between 1948 and 1988, $n = 160$ values.

cycle every eight years in government investment, and perhaps to lesser degrees in unemployment, GNP, and consumption.

Figure 7.21 shows the coherences among various series. At the frequencies of interest, $\omega = .125$ and $.25$, pairwise, GNP, Unemployment, Consumption, and Private Investment (except for Unemployment and Private Investment) are coherent. Government Investment is either not coherent or minimally coherent with the other series.

Figure 7.22 shows $\hat{\lambda}_1(\omega)$ and $\hat{\lambda}_2(\omega)$, the first and second eigenvalues of the estimated spectral matrix $\hat{f}_{xx}(\omega)$. These eigenvalues suggest the first factor is identified by the frequency of one cycle every four years, whereas the second factor is identified by the frequency of one cycle every eight years. The modulus of the corresponding eigenvectors at the frequencies of interest, $\hat{e}_1(10/160)$ and $\hat{e}_2(5/160)$, are shown in Table 7.6. These values confirm Unemployment, GNP, Consumption, and Private Investment load on the first factor, and Government Investment loads on the second factor. The remainder of the details involving the factor analysis of these data is left as an exercise.




Fig. 7.20. The individual estimated spectra (scaled by 1000) of each series show in Figure 7.19 in terms of the number of cycles in 160 quarters.

The following code was used to perform the analysis in R.

```
gr = diff(log(ts(econ5, start=1948, frequency=4))) # growth rate
plot(100*gr, main="Growth Rates (%)")
# scale each series to have variance 1
gr = ts(apply(gr,2,scale), freq=4) # scaling strips ts attributes
L = c(7,7) # degree of smoothing
gr.spec = mvspec(gr, spans=L, demean=FALSE, detrend=FALSE, taper=.25)
dev.new()
plot(kernel("modified.daniell", L)) # view the kernel - not shown
dev.new()
plot(gr.spec, log="no", col=1, main="Individual Spectra", lty=1:5, lwd=2)
legend("topright", colnames(econ5), lty=1:5, lwd=2)
dev.new()
plot.spec.coherency(gr.spec, ci=NA, main="Squared Coherencies")
# PCs
n.freq = length(gr.spec$freq)
lam = matrix(0,n.freq,5)
for (k in 1:n.freq) lam[k,] = eigen(gr.spec$fx[,k], symmetric=TRUE,
  only.values=TRUE)$values
dev.new()
par(mfrow=c(2,1), mar=c(4,2,2,1), mgp=c(1.6,.6,0))
plot(gr.spec$freq, lam[,1], type="l", ylab="", xlab="Frequency", main="First
  Eigenvalue")
```

Squared Coherencies




Fig. 7.21. The squared coherencies between the various series displayed in Figure 7.19.

Table 7.6. Magnitudes of the Eigenvectors in Example 7.16

	Unemp	GNP	Cons	G. Inv.	P. Inv.
$ \hat{e}_1(\frac{10}{160}) $	0.53	0.50	0.51	0.06	0.44
$ \hat{e}_2(\frac{5}{160}) $	0.19	0.14	0.23	0.93	0.16

```
abline(v=.25, lty=2)
plot(gr.spec$freq, lam[,2], type="l", ylab="", xlab="Frequency",
     main="Second Eigenvalue")
abline(v=.125, lty=2)
e.vec1 = eigen(gr.spec$fx[, ,10], symmetric=TRUE)$vectors[,1]
e.vec2 = eigen(gr.spec$fx[, ,5], symmetric=TRUE)$vectors[,2]
round(Mod(e.vec1), 2); round(Mod(e.vec2), 3)
```




Fig. 7.22. The first, $\hat{\lambda}_1(\omega)$, and second, $\hat{\lambda}_2(\omega)$, eigenvalues of the estimated spectral matrix $\hat{f}_{xx}(\omega)$. The vertical dashed lines are shown at $\omega = .25$ and $.125$, respectively.

7.9 The Spectral Envelope

The concept of spectral envelope for the spectral analysis and scaling of categorical time series was first introduced in Stoffer et al. (1993). Since then, the idea has been extended in various directions (not only restricted to categorical time series), and we will explore these problems as well. First, we give a brief introduction to the concept of scaling time series.

The spectral envelope was motivated by collaborations with researchers who collected categorical-valued time series with an interest in the cyclic behavior of the data. For example, [Table 7.7](#) shows the per-minute sleep state of an infant taken from a study on the effects of prenatal exposure to alcohol. Details can be found in Stoffer et al. (1988), but, briefly, an electroencephalographic (EEG) sleep recording of approximately two hours is obtained on a full-term infant 24 to 36 hours after birth, and the recording is scored by a pediatric neurologist for sleep state. Sleep state is categorized, per minute, into one of six possible states: **qt**: quiet sleep - trace alternant, **qh**: quiet sleep - high voltage, **tr**: transitional sleep, **al**: active sleep - low voltage, **ah**: active sleep - high voltage, and **aw**: awake. This particular infant was never awake during the study.

It is not difficult to notice a pattern in the data if we concentrate on active vs. quiet sleep (that is, focus on the first letter). But, it would be difficult to try to assess patterns in a longer sequence, or if more categories were present, without some graphical aid.

Table 7.7. Infant EEG Sleep-States (per minute)
(read down and across)

ah	qt	qt	al	tr	qt	al	ah
ah	qt	qt	ah	tr	qt	al	ah
ah	qt	tr	ah	tr	qt	al	ah
ah	qt	al	ah	qh	qt	al	ah
ah	qt	al	ah	qh	qt	al	ah
ah	tr	al	ah	qt	qt	al	ah
ah	qt	al	ah	qt	qt	al	ah
ah	qt	al	ah	qt	qt	al	ah
tr	qt	tr	tr	qt	qt	al	tr
ah	qt	ah	tr	qt	tr	al	
tr	qt	al	ah	qt	al	al	
ah	qt	al	ah	qt	al	al	
ah	qt	al	ah	qt	al	al	
qh	qt	al	ah	qt	al	ah	

One simple method would be to *scale* the data, that is, *assign numerical values to the categories* and then draw a time plot of the scaled series. Because the states have an order, one obvious scaling is

$$1 = \text{qt} \quad 2 = \text{qh} \quad 3 = \text{tr} \quad 4 = \text{al} \quad 5 = \text{ah} \quad 6 = \text{aw}, \quad (7.179)$$

and [Figure 7.23](#) shows the time plot using this scaling. Another interesting scaling might be to combine the quiet states and the active states:

$$1 = \text{qt} \quad 1 = \text{qh} \quad 2 = \text{tr} \quad 3 = \text{al} \quad 3 = \text{ah} \quad 4 = \text{aw}. \quad (7.180)$$

The time plot using (7.180) would be similar to [Figure 7.23](#) as far as the cyclic (in and out of quiet sleep) behavior of this infant’s sleep pattern. [Figure 7.23](#) shows the periodogram of the sleep data using the scaling in (7.179). A large peak exists at the frequency corresponding to one cycle every 60 minutes. As we might imagine, the general appearance of the periodogram using the scaling (7.180) (not shown) is similar to [Figure 7.23](#). Most of us would feel comfortable with this analysis even though we made an arbitrary and ad hoc choice about the particular scaling. It is evident from the data (without any scaling) that if the interest is in infant sleep cycling, this particular sleep study indicates an infant cycles between active and quiet sleep at a rate of about one cycle per hour.

The intuition used in the previous example is lost when we consider a long DNA sequence. Briefly, a DNA strand can be viewed as a long string of linked nucleotides. Each nucleotide is composed of a nitrogenous base, a five carbon sugar, and a phosphate group. There are four different bases, and they can be grouped by size; the pyrimidines, thymine (T) and cytosine (C), and the purines, adenine (A) and guanine (G). The nucleotides are linked together by a backbone of alternating sugar and phosphate groups with the 5’ carbon of one sugar linked to the 3’ carbon of the next, giving the string direction. DNA molecules occur naturally as a double helix




Fig. 7.23. [Top] Time plot of the EEG sleep state data in Table 7.7 using the scaling in (7.179). [Bottom] Periodogram of the EEG sleep state data in Figure 7.23 based on the scaling in (7.179). The peak corresponds to a frequency of approximately one cycle every 60 minutes.

composed of polynucleotide strands with the bases facing inwards. The two strands are complementary, so it is sufficient to represent a DNA molecule by a sequence of bases on a single strand. Thus, a strand of DNA can be represented as a sequence of letters, termed base pairs (bp), from the finite alphabet $\{A, C, G, T\}$. The order of the nucleotides contains the genetic information specific to the organism. Expression of information stored in these molecules is a complex multistage process. One important task is to translate the information stored in the protein-coding sequences (CDS) of the DNA. A common problem in analyzing long DNA sequence data is in identifying CDS dispersed throughout the sequence and separated by regions of noncoding (which makes up most of the DNA). Table 7.8 shows part of the Epstein–Barr virus (EBV) DNA sequence. The entire EBV DNA sequence consists of approximately 172,000 bp.

We could try scaling according to the purine–pyrimidine alphabet, that is $A = G = 0$ and $C = T = 1$, but this is not necessarily of interest for every CDS of EBV. Numerous possible alphabets of interest exist. For example, we might focus on the strong–weak hydrogen-bonding alphabet $C = G = 0$ and $A = T = 1$. Although model calculations as well as experimental data strongly agree that some kind of periodic

Table 7.8. Part of the Epstein–Barr Virus DNA Sequence
(read across and down)

AGAATTCGTC	TTGCTCTATT	CACCCCTACT	TTTCTTCTTG	CCCGTTCTCT	TTCTTAGTAT
GAATCCAGTA	TGCCTGCTG	TAATTGTTGC	GCCCTACCTC	TTTGGCTGG	CGGCTATTGC
CGCCTCGTGT	TTACGGCCT	CAGTTAGTAC	CGTTGTGACC	GCCACCGGCT	TGGCCCTCTC
ACTTCTACTC	TTGGCAGCAG	TGGCCAGCTC	ATATGCCGCT	GCACAAAGGA	AACTGCTGAC
ACCGGTGACA	GTGCTTACTG	CGGTTGTCAC	TTGTGAGTAC	ACACGCACCA	TTTACAATGC
ATGATGTTTC	TGAGATTGAT	CTGTCTCTAA	CAGTTCACCT	CCTCTGCTTT	TCTCCTCAGT
CTTTGCAATT	TGCCTAACAT	GGAGGATTGA	GGACCCACCT	TTTAATTCTC	TTCTGTTTGC
ATTGCTGGCC	GCAGCTGGCG	GACTACAAGG	CATTTACGGT	TAGTGTGCCT	CTGTTATGAA
ATGCAGGTTT	GACTTCATAT	GTATGCCTTG	GCATGACGTC	AACTTTACTT	TTATTTCAGT
TCCTGGTAGT	CTTGTGCTCC	TGATACTAGC	GTACAGAAGG	AGATGGCGCC	GTTTGACTGT
TTGTGGCGGC	ATCATGTGTTT	TGGCATGTGT	ACTTGTCCCT	ATCGTCGACG	CTGTTTGGCA
GCTGAGTCCC	CTCCTTGAG	CTGTAAGTGT	GGTTTCCATG	ACGCTGCTGC	TACTGGCTTT
CGTCCTCTGG	CTCTCTTGC	CAGGGGGCCT	AGGTACTCTT	GGTGCAGCCC	TTTAAACATT
GGCAGCAGGT	AAGCCACACG	TGTGACATTG	CTTGCCTTTT	TGCCACATGT	TTTCTGGACA
CAGGACTAAC	CATGCCATCT	CTGATTATAG	CTCTGGCACT	GCTAGCGTCA	CTGATTTTGG
GCACACTTAA	CTTGACTACA	ATGTTCCCTC	TCATGCTCCT	ATGGACACTT	GGTAAAGTTT
CCCTTCCTTT	AACTCATTAC	TTGTTCTTTT	GTAATCGCAG	CTCTAACTTG	GCATCTCTTT
TACAGTGTTT	CTCCTGATTT	GCTCTTCGTG	CTCTTCATGT	CCACTGAGCA	AGATCCTTCT

signal exists in certain DNA sequences, a large disagreement about the exact type of periodicity exists. In addition, a disagreement exists about which nucleotide alphabets are involved in the signals.

If we consider the naive approach of arbitrarily assigning numerical values (scales) to the categories and then proceeding with a spectral analysis, the result will depend on the particular assignment of numerical values. For example, consider the artificial sequence **ACGTACGTACGT**... . Then, setting **A** = **G** = 0 and **C** = **T** = 1 yields the numerical sequence 0101010101... , or one cycle every two base pairs. Another interesting scaling is **A** = 1, **C** = 2, **G** = 3, and **T** = 4, which results in the sequence 123412341234... , or one cycle every four bp. In this example, both scalings (that is, $\{\mathbf{A}, \mathbf{C}, \mathbf{G}, \mathbf{T}\} = \{0, 1, 0, 1\}$ and $\{\mathbf{A}, \mathbf{C}, \mathbf{G}, \mathbf{T}\} = \{1, 2, 3, 4\}$) of the nucleotides are interesting and bring out different properties of the sequence. Hence, we do not want to focus on only one scaling. Instead, the focus should be on finding all possible scalings that bring out all of the interesting features in the data. Rather than choose values arbitrarily, the spectral envelope approach selects scales that help emphasize any periodic feature that exists in a categorical time series of virtually any length in a quick and automated fashion. In addition, the technique can help in determining whether a sequence is merely a random assignment of categories.

THE SPECTRAL ENVELOPE FOR CATEGORICAL TIME SERIES

As a general description, the spectral envelope is a frequency-based principal components technique applied to a multivariate time series. First, we will focus on the basic concept and its use in the analysis of categorical time series. Technical details can be found in Stoffer et al. (1993).

Briefly, in establishing the spectral envelope for categorical time series, the basic question of how to efficiently discover periodic components in categorical time series

was addressed. This, was accomplished via nonparametric spectral analysis as follows. Let x_t , $t = 0, \pm 1, \pm 2, \dots$, be a categorical-valued time series with finite state-space $\mathcal{C} = \{c_1, c_2, \dots, c_k\}$. Assume x_t is stationary and $p_j = \Pr\{x_t = c_j\} > 0$ for $j = 1, 2, \dots, k$. For $\boldsymbol{\beta} = (\beta_1, \beta_2, \dots, \beta_k)' \in \mathbb{R}^k$, denote by $x_t(\boldsymbol{\beta})$ the real-valued stationary time series corresponding to the scaling that assigns the category c_j the numerical value β_j , $j = 1, 2, \dots, k$. The spectral density of $x_t(\boldsymbol{\beta})$ will be denoted by $f_{xx}(\omega; \boldsymbol{\beta})$. The goal is to find scalings $\boldsymbol{\beta}$, so the spectral density is in some sense interesting, and to summarize the spectral information by what is called the spectral envelope.

In particular, $\boldsymbol{\beta}$ is chosen to maximize the power at each frequency, ω , of interest, relative to the total power $\sigma^2(\boldsymbol{\beta}) = \text{var}\{x_t(\boldsymbol{\beta})\}$. That is, we chose $\boldsymbol{\beta}(\omega)$, at each ω of interest, so

$$\lambda(\omega) = \max_{\boldsymbol{\beta}} \left\{ \frac{f_{xx}(\omega; \boldsymbol{\beta})}{\sigma^2(\boldsymbol{\beta})} \right\}, \quad (7.181)$$

over all $\boldsymbol{\beta}$ not proportional to $\mathbf{1}_k$, the $k \times 1$ vector of ones. Note, $\lambda(\omega)$ is not defined if $\boldsymbol{\beta} = a\mathbf{1}_k$ for $a \in \mathbb{R}$ because such a scaling corresponds to assigning each category the same value a ; in this case, $f_{xx}(\omega; \boldsymbol{\beta}) \equiv 0$ and $\sigma^2(\boldsymbol{\beta}) = 0$. The optimality criterion $\lambda(\omega)$ possesses the desirable property of being invariant under location and scale changes of $\boldsymbol{\beta}$.

As in most scaling problems for categorical data, it is useful to represent the categories in terms of the unit vectors $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_k$, where $\mathbf{u}_j$ represents the $k \times 1$ vector with a one in the j -th row, and zeros elsewhere. We then define a k -dimensional stationary time series $\mathbf{y}_t$ by $\mathbf{y}_t = \mathbf{u}_j$ when $x_t = c_j$. The time series $x_t(\boldsymbol{\beta})$ can be obtained from the $\mathbf{y}_t$ time series by the relationship $x_t(\boldsymbol{\beta}) = \boldsymbol{\beta}'\mathbf{y}_t$. Assume the vector process $\mathbf{y}_t$ has a continuous spectral density denoted by $f_{yy}(\omega)$. For each ω , $f_{yy}(\omega)$ is, of course, a $k \times k$ complex-valued Hermitian matrix. The relationship $x_t(\boldsymbol{\beta}) = \boldsymbol{\beta}'\mathbf{y}_t$ implies $f_{xx}(\omega; \boldsymbol{\beta}) = \boldsymbol{\beta}'f_{yy}(\omega)\boldsymbol{\beta} = \boldsymbol{\beta}'f_{yy}^{re}(\omega)\boldsymbol{\beta}$, where $f_{yy}^{re}(\omega)$ denotes the real part² of $f_{yy}(\omega)$. The imaginary part disappears from the expression because it is skew-symmetric, that is, $f_{yy}^{im}(\omega)' = -f_{yy}^{im}(\omega)$. The optimality criterion can thus be expressed as

$$\lambda(\omega) = \max_{\boldsymbol{\beta}} \left\{ \frac{\boldsymbol{\beta}'f_{yy}^{re}(\omega)\boldsymbol{\beta}}{\boldsymbol{\beta}'V\boldsymbol{\beta}} \right\}, \quad (7.182)$$

where V is the variance–covariance matrix of $\mathbf{y}_t$. The resulting scaling $\boldsymbol{\beta}(\omega)$ is called the optimal scaling.

The $\mathbf{y}_t$ process is a multivariate point process, and any particular component of $\mathbf{y}_t$ is the individual point process for the corresponding state (for example, the first component of $\mathbf{y}_t$ indicates whether the process is in state c_1 at time t). For any fixed t , $\mathbf{y}_t$ represents a single observation from a simple multinomial sampling scheme. It readily follows that $V = D - p p'$, where $p = (p_1, \dots, p_k)'$, and D is the $k \times k$ diagonal matrix $D = \text{diag}\{p_1, \dots, p_k\}$. Because, by assumption, $p_j > 0$ for $j = 1, 2, \dots, k$, it follows that $\text{rank}(V) = k - 1$ with the null space of V being spanned by $\mathbf{1}_k$. For any $k \times (k - 1)$ full rank matrix Q whose columns are linearly independent of $\mathbf{1}_k$, $Q'VQ$ is a $(k - 1) \times (k - 1)$ positive-definite symmetric matrix.

² In this section, it is more convenient to write complex values in the form $z = z^{re} + iz^{im}$, which represents a change from the notation used previously.

With the matrix Q as previously defined, define $\lambda(\omega)$ to be the largest eigenvalue of the determinantal equation

$$|Q' f_{yy}^{re}(\omega) Q - \lambda(\omega) Q' V Q| = 0,$$

and let $\mathbf{b}(\omega) \in \mathbb{R}^{k-1}$ be any corresponding eigenvector, that is,

$$Q' f_{yy}^{re}(\omega) Q \mathbf{b}(\omega) = \lambda(\omega) Q' V Q \mathbf{b}(\omega).$$

The eigenvalue $\lambda(\omega) \geq 0$ does not depend on the choice of Q . Although the eigenvector $\mathbf{b}(\omega)$ depends on the particular choice of Q , the equivalence class of scalings associated with $\boldsymbol{\beta}(\omega) = Q \mathbf{b}(\omega)$ does not depend on Q . A convenient choice of Q is $Q = [I_{k-1} \mid \mathbf{0}]'$, where I_{k-1} is the $(k-1) \times (k-1)$ identity matrix and $\mathbf{0}$ is the $(k-1) \times 1$ vector of zeros. For this choice, $Q' f_{yy}^{re}(\omega) Q$ and $Q' V Q$ are the upper $(k-1) \times (k-1)$ blocks of $f_{yy}^{re}(\omega)$ and V , respectively. This choice corresponds to setting the last component of $\boldsymbol{\beta}(\omega)$ to zero.

The value $\lambda(\omega)$ itself has a useful interpretation; specifically, $\lambda(\omega)d\omega$ represents the largest proportion of the total power that can be attributed to the frequencies $(\omega, \omega+d\omega)$ for any particular scaled process $x_t(\boldsymbol{\beta})$, with the maximum being achieved by the scaling $\boldsymbol{\beta}(\omega)$. Because of its central role, $\lambda(\omega)$ is defined to be the spectral envelope of a stationary categorical time series.

The name spectral envelope is appropriate since $\lambda(\omega)$ envelopes the standardized spectrum of any scaled process. That is, given any $\boldsymbol{\beta}$ normalized so that $x_t(\boldsymbol{\beta})$ has total power one, $f_{xx}(\omega; \boldsymbol{\beta}) \leq \lambda(\omega)$ with equality if and only if $\boldsymbol{\beta}$ is proportional to $\boldsymbol{\beta}(\omega)$.

Given observations x_t , for $t = 1, \dots, n$, on a categorical time series, we form the multinomial point process $\mathbf{y}_t$, for $t = 1, \dots, n$. Then, the theory for estimating the spectral density of a multivariate, real-valued time series can be applied to estimating $f_{yy}(\omega)$, the $k \times k$ spectral density of $\mathbf{y}_t$. Given an estimate $\hat{f}_{yy}(\omega)$ of $f_{yy}(\omega)$, estimates $\hat{\lambda}(\omega)$ and $\hat{\boldsymbol{\beta}}(\omega)$ of the spectral envelope, $\lambda(\omega)$, and the corresponding scalings, $\boldsymbol{\beta}(\omega)$, can then be obtained. Details on estimation and inference for the sample spectral envelope and the optimal scalings can be found in Stoffer et al. (1993), but the main result of that paper is as follows: If $\hat{f}_{yy}(\omega)$ is a consistent spectral estimator and if for each $j = 1, \dots, J$, the largest root of $f_{yy}^{re}(\omega_j)$ is distinct, then

$$\left\{ \eta_n [\hat{\lambda}(\omega_j) - \lambda(\omega_j)] / \lambda(\omega_j), \eta_n [\hat{\boldsymbol{\beta}}(\omega_j) - \boldsymbol{\beta}(\omega_j)]; j = 1, \dots, J \right\} \quad (7.183)$$

converges ($n \rightarrow \infty$) jointly in distribution to independent zero-mean, normal, distributions, the first of which is standard normal; the asymptotic covariance structure of $\hat{\boldsymbol{\beta}}(\omega_j)$ is discussed in Stoffer et al. (1993). Result (7.183) is similar to (7.150), but in this case, $\boldsymbol{\beta}(\omega)$ and $\hat{\boldsymbol{\beta}}(\omega)$ are real. The term η_n is the same as in (7.183), and its value depends on the type of estimator being used. Based on these results, asymptotic normal confidence intervals and tests for $\lambda(\omega)$ can be readily constructed. Similarly, for $\boldsymbol{\beta}(\omega)$, asymptotic confidence ellipsoids and chi-square tests can be constructed; details can be found in Stoffer et al. (1993, Theorems 3.1 – 3.3).

Peak searching for the smoothed spectral envelope estimate can be aided using the following approximations. Using a first-order Taylor expansion, we have

$$\log \widehat{\lambda}(\omega) \approx \log \lambda(\omega) + \frac{\widehat{\lambda}(\omega) - \lambda(\omega)}{\lambda(\omega)}, \quad (7.184)$$

so $\eta_n[\log \widehat{\lambda}(\omega) - \log \lambda(\omega)]$ is approximately standard normal. It follows that $E[\log \widehat{\lambda}(\omega)] \approx \log \lambda(\omega)$ and $\text{var}[\log \widehat{\lambda}(\omega)] \approx \eta_n^{-2}$. If no signal is present in a sequence of length n , we expect $\lambda(j/n) \approx 2/n$ for $1 < j < n/2$, and hence approximately $(1 - \alpha) \times 100\%$ of the time, $\log \widehat{\lambda}(\omega)$ will be less than $\log(2/n) + (z_\alpha/\eta_n)$, where z_α is the $(1 - \alpha)$ upper tail cutoff of the standard normal distribution. Exponentiating, the α critical value for $\widehat{\lambda}(\omega)$ becomes $(2/n) \exp(z_\alpha/\eta_n)$. Useful values of z_α are $z_{.001} = 3.09$, $z_{.0001} = 3.71$, and $z_{.00001} = 4.26$, and from our experience, thresholding at these levels works well.

Example 7.17 Spectral Analysis of DNA Sequences

We give explicit instructions for the calculations involved in estimating the spectral envelope of a DNA sequence, x_t , for $t = 1, \dots, n$, using the nucleotide alphabet.

- (i) In this example, we hold the scale for **T** fixed at zero. In this case, we form the 3×1 data vectors $\mathbf{y}_t$:

$$\begin{aligned} \mathbf{y}_t &= (1, 0, 0)' \text{ if } x_t = \mathbf{A}; & \mathbf{y}_t &= (0, 1, 0)' \text{ if } x_t = \mathbf{C}; \\ \mathbf{y}_t &= (0, 0, 1)' \text{ if } x_t = \mathbf{G}; & \mathbf{y}_t &= (0, 0, 0)' \text{ if } x_t = \mathbf{T}. \end{aligned}$$

The scaling vector is $\boldsymbol{\beta} = (\beta_1, \beta_2, \beta_3)'$, and the scaled process is $x_t(\boldsymbol{\beta}) = \boldsymbol{\beta}'\mathbf{y}_t$.

- (ii) Calculate the DFT of the data

$$\mathbf{Y}(j/n) = n^{-1/2} \sum_{t=1}^n \mathbf{y}_t \exp(-2\pi i t j/n).$$

Note $\mathbf{Y}(j/n)$ is a 3×1 complex-valued vector. Calculate the periodogram, $I(j/n) = \mathbf{Y}(j/n)\mathbf{Y}^*(j/n)$, for $j = 1, \dots, [n/2]$, and retain only the real part, say, $I^{re}(j/n)$.

- (iii) Smooth the $I^{re}(j/n)$ to obtain an estimate of $f_{yy}^{re}(j/n)$. For example, using (7.149) with $L = 3$ and triangular weighting, we would calculate

$$\widehat{f}_{yy}^{re}(j/n) = \frac{1}{4} I^{re}\left(\frac{j-1}{n}\right) + \frac{1}{2} I^{re}\left(\frac{j}{n}\right) + \frac{1}{4} I^{re}\left(\frac{j+1}{n}\right).$$

- (iv) Calculate the 3×3 sample variance–covariance matrix,

$$S_{yy} = n^{-1} \sum_{t=1}^n (\mathbf{y}_t - \bar{\mathbf{y}})(\mathbf{y}_t - \bar{\mathbf{y}})',$$

where $\bar{\mathbf{y}} = n^{-1} \sum_{t=1}^n \mathbf{y}_t$ is the sample mean of the data.




Fig. 7.24. Smoothed sample spectral envelope of the BNRF1 gene from the Epstein–Barr virus.

- (v) For each $\omega_j = j/n$, $j = 0, 1, \dots, [n/2]$, determine the largest eigenvalue and the corresponding eigenvector of the matrix $2n^{-1} S_{yy}^{-1/2} \widehat{f}_{yy}(\omega_j) S_{yy}^{-1/2}$. Note, $S_{yy}^{1/2}$ is the unique square root matrix of S_{yy} .
- (vi) The sample spectral envelope $\widehat{\lambda}(\omega_j)$ is the eigenvalue obtained in the previous step. If $\mathbf{b}(\omega_j)$ denotes the eigenvector obtained in the previous step, the optimal sample scaling is $\widehat{\mathbf{b}}(\omega_j) = S_{yy}^{-1/2} \mathbf{b}(\omega_j)$; this will result in three values, the value corresponding to the fourth category, $\mathbf{T}$ being held fixed at zero.

Example 7.18 Dynamic Analysis of the Gene Labeled BNRF1 of the Epstein–Barr Virus

In this example, we focus on a dynamic (or sliding-window) analysis of the gene labeled BNRF1 (bp 1736–5689) of Epstein–Barr. **Figure 7.24** shows the spectral envelope estimate of the entire coding sequence (3954 bp long). The figure also shows a strong signal at frequency $1/3$; the corresponding optimal scaling was $\mathbf{A} = .10, \mathbf{C} = .61, \mathbf{G} = .78, \mathbf{T} = 0$, which indicates the signal is in the strong–weak bonding alphabet, $S = \{\mathbf{C}, \mathbf{G}\}$ and $W = \{\mathbf{A}, \mathbf{T}\}$.

Figure 7.25 shows the result of computing the spectral envelope over three nonoverlapping 1000-bp windows and one window of 954 bp, across the CDS, namely, the first, second, third, and fourth quarters of BNRF1. An approximate 0.0001 significance threshold is .69%. The first three quarters contain the signal at the frequency $1/3$ (**Figure 7.25a–c**); the corresponding sample optimal scalings for the first three windows were (a) $\mathbf{A} = .01, \mathbf{C} = .71, \mathbf{G} = .71, \mathbf{T} = 0$; (b) $\mathbf{A} = .08, \mathbf{C} =$




Fig. 7.25. Smoothed sample spectral envelope of the BNRF1 gene from the Epstein–Barr virus: (a) first 1000 bp, (b) second 1000 bp, (c) third 1000 bp, and (d) last 954 bp.

0.71, $G = .70$, $T = 0$; (c) $A = .20$, $C = .58$, $G = .79$, $T = 0$. The first two windows are consistent with the overall analysis. The third section, however, shows some minor departure from the strong-weak bonding alphabet. The most interesting outcome is that the fourth window shows that no signal is present. This leads to the conjecture that the fourth quarter of BNRF1 of Epstein–Barr is actually noncoding.

The R code for this example is as follows.

```
u = factor(bnrf1ebv) # first, input the data as factors and then
x = model.matrix(~u-1)[,1:3] # make an indicator matrix
# x = x[1:1000,] # select subsequence if desired
Var = var(x) # var-cov matrix
xspe = mvspec(x, spans=c(7,7))
fxr = Re(xspe$fx) # fxr is real(fxx)
# compute  $Q = Var^{-1/2}$ 
ev = eigen(Var)
Q = ev$vectors %*% diag(1/sqrt(ev$values)) %*% t(ev$vectors)
# compute spec envelope and scale vectors
num = xspe$n.used # sample size used for FFT
```

```

nfreq = length(xspec$freq) # number of freqs used
specenv = matrix(0,nfreq,1) # initialize the spec envelope
beta = matrix(0,nfreq,3) # initialize the scale vectors
for (k in 1:nfreq){
  ev = eigen(2*Q%%fxxr[,k]%%Q/num, symmetric=TRUE)
  specenv[k] = ev$values[1] # spec env at freq k/n is max evalue
  b = Q%%ev$vectors[,1] # beta at freq k/n
  beta[k,] = b/sqrt(sum(b^2)) } # helps to normalize beta
# output and graphics
frequency = xspec$freq
plot(frequency, 100*specenv, type="l", ylab="Spectral Envelope (%)")
# add significance threshold to plot
m = xspec$kernel$m
etainv = sqrt(sum(xspec$kernel[-m:m]^2))
thresh = 100*(2/num)*exp(qnorm(.9999)*etainv)*rep(1,nfreq)
lines(frequency, thresh, lty="dashed", col="blue")
# details
output = cbind(frequency, specenv, beta)
colnames(output) = c("freq", "specenv", "A", "C", "G")
round(output,3)

```

THE SPECTRAL ENVELOPE FOR REAL-VALUED TIME SERIES

The concept of the spectral envelope for categorical time series was extended to real-valued time series, $\{x_t; t = 0, \pm 1, \pm 2, \dots\}$, in McDougall et al. (1997). The process x_t can be vector-valued, but here we will concentrate on the univariate case. Further details can be found in McDougall et al. (1997). The concept is similar to projection pursuit (Friedman and Stuetzle, 1981). Let $\mathcal{G}$ denote a k -dimensional vector space of continuous real-valued transformations with $\{g_1, \dots, g_k\}$ being a set of basis functions satisfying $E[g_i(x_t)^2] < \infty, i = 1, \dots, k$. Analogous to the categorical time series case, define the scaled time series with respect to the set $\mathcal{G}$ to be the real-valued process

$$x_t(\boldsymbol{\beta}) = \boldsymbol{\beta}'\mathbf{y}_t = \beta_1 g_1(x_t) + \dots + \beta_k g_k(x_t)$$

obtained from the vector process

$$\mathbf{y}_t = (g_1(X_t), \dots, g_k(X_t))',$$

where $\boldsymbol{\beta} = (\beta_1, \dots, \beta_k)' \in \mathbf{R}^k$. If the vector process, $\mathbf{y}_t$, is assumed to have a continuous spectral density, say, $f_{yy}(\omega)$, then $x_t(\boldsymbol{\beta})$ will have a continuous spectral density $f_{xx}(\omega; \boldsymbol{\beta})$ for all $\boldsymbol{\beta} \neq \mathbf{0}$. Noting, $f_{xx}(\omega; \boldsymbol{\beta}) = \boldsymbol{\beta}' f_{yy}(\omega) \boldsymbol{\beta} = \boldsymbol{\beta}' f_{yy}^{re}(\omega) \boldsymbol{\beta}$, and $\sigma^2(\boldsymbol{\beta}) = \text{var}[x_t(\boldsymbol{\beta})] = \boldsymbol{\beta}' V \boldsymbol{\beta}$, where $V = \text{var}(\mathbf{y}_t)$ is assumed to be positive definite, the optimality criterion

$$\lambda(\omega) = \sup_{\boldsymbol{\beta} \neq \mathbf{0}} \left\{ \frac{\boldsymbol{\beta}' f_{yy}^{re}(\omega) \boldsymbol{\beta}}{\boldsymbol{\beta}' V \boldsymbol{\beta}} \right\}, \quad (7.185)$$

is well defined and represents the largest proportion of the total power that can be attributed to the frequency ω for any particular scaled process $x_t(\boldsymbol{\beta})$. This interpretation of $\lambda(\omega)$ is consistent with the notion of the spectral envelope introduced in the




Fig. 7.26. Spectral envelope with respect to $\mathcal{G} = \{x, |x|, x^2\}$ of the residuals from an MA(2) fit to the U.S. GNP growth rate data.

previous section and provides the following working definition: *The spectral envelope of a time series with respect to the space $\mathcal{G}$ is defined to be $\lambda(\omega)$.*

The solution to this problem, as in the categorical case, is attained by finding the largest scalar $\lambda(\omega)$ such that

$$f_{yy}^{re}(\omega)\boldsymbol{\beta}(\omega) = \lambda(\omega)V\boldsymbol{\beta}(\omega) \quad (7.186)$$

for $\boldsymbol{\beta}(\omega) \neq \mathbf{0}$. That is, $\lambda(\omega)$ is the largest eigenvalue of $f_{yy}^{re}(\omega)$ in the metric of V , and the optimal scaling, $\boldsymbol{\beta}(\omega)$, is the corresponding eigenvector.

If x_t is a categorical time series taking values in the finite state-space $\mathcal{S} = \{c_1, c_2, \dots, c_k\}$, where c_j represents a particular category, an appropriate choice for $\mathcal{G}$ is the set of indicator functions $g_j(x_t) = I(x_t = c_j)$. Hence, this is a natural generalization of the categorical case. In the categorical case, $\mathcal{G}$ does not consist of linearly independent g 's, but it was easy to overcome this problem by reducing the dimension by one. In the vector-valued case, $\mathbf{x}_t = (x_{1t}, \dots, x_{pt})'$, we consider $\mathcal{G}$ to be the class of transformations from $\mathbb{R}^p$ into $\mathbb{R}$ such that the spectral density of $g(\mathbf{x}_t)$ exists. One class of transformations of interest are linear combinations of $\mathbf{x}_t$. In Tiao et al. (1993), for example, linear transformations of this type are used in a time domain approach to investigate contemporaneous relationships among the components of multivariate time series. Estimation and inference for the real-valued case are analogous to the methods described in the previous section for the categorical case. We focus on two examples here; numerous other examples can be found in McDougall et al. (1997).

Example 7.19 Residual Analysis

A relevant situation may be when x_t is the residual process obtained from some modeling procedure. If the fitted model is appropriate, the residuals should exhibit properties similar to an iid sequence. Departures of the data from the fitted model may suggest model misspecification, non-Gaussian data, or the existence of a non-linear structure, and the spectral envelope would provide a simple diagnostic tool to aid in a residual analysis.

The series considered here is the quarterly U.S. real GNP which was analyzed in Chapter 3, [Example 3.38](#) and [Example 3.39](#). Recall an MA(2) model was fit to the growth rate, and the residuals from this fit are plotted in [Figure 3.19](#). As discussed in [Example 3.39](#), the residuals from the model fit appear to be uncorrelated; there appears to be one or two outliers, but their magnitudes are not that extreme. In addition, the standard residual analyses showed no obvious structure among the residuals.

Although the MA(2) model appears to be appropriate, Tiao and Tsay (1994) investigated the possibility of nonlinearities in GNP growth rate. Their overall conclusion was that there is subtle nonlinear behavior in the data because the economy behaves differently during expansion periods than during recession periods.

The spectral envelope, used as a diagnostic tool on the residuals, clearly indicates the MA(2) model is not adequate, and that further analysis is warranted. Here, the generating set $\mathcal{G} = \{x, |x|, x^2\}$ —which seems natural for a residual analysis—was used to estimate the spectral envelope for the residuals from the MA(2) fit, and the result is plotted in [Figure 7.26](#). Clearly, the residuals are not iid, and considerable power is present at the low frequencies. The presence of spectral power at very low frequencies in detrended economic series has been frequently reported and is typically associated with long-range dependence. In fact, our choice of $\mathcal{G}$ was partly influenced by the work of Ding et al. (1993) who applied transformations of the form $|x_t|^d$, for $d \in (0, 3]$, to the S&P 500 stock market series. The estimated optimal transformation at the first nonzero frequency, $\omega = .018$, was $\hat{\beta}(.018) = (1, 22, -478)'$, which leads to the transformation

$$g(x) = x + 22|x| - 478x^2. \quad (7.187)$$

This transformation is plotted in [Figure 7.27](#). The transformation given in (7.187) is basically the absolute value (with some slight curvature and asymmetry) for most of the residual values, but the effect of extreme-valued residuals (outliers) is dampened.

The following R code, which is nearly identical to the previous example, was used in this example.

```
u = arima(diff(log(gnp)), order=c(0,0,2))$resid # residuals
x = cbind(u, abs(u), u^2) # transformation set
Var = var(x)
xspect = mvspec(x, spans=c(5,3), taper=.1)
fxr = Re(xspect$fx); ev = eigen(Var)
Q = ev$vectors%*%diag(1/sqrt(ev$values))%*%t(ev$vectors)
num = xspect$n.used; nfreq = length(xspect$freq)
specenv = matrix(0, nfreq, 1); beta = matrix(0,nfreq,3)
```




Fig. 7.27. Estimated optimal transformation, (7.187), for the GNP residuals at $\omega = .018$.

```
for (k in 1:nfreq){
  ev = eigen(2*Q%%fxxr[,k]%%Q/num)
  specenv[k] = ev$values[1]
  b = Q%%ev$vectors[,1]
  beta[k,] = b/b[1] }
# output and graphics
frequency = xspec$freq
plot(frequency, 100*specenv, type="l", ylab="Spectral Envelope (%)")
output = cbind(frequency, specenv, beta)
colnames(output) = c("freq", "specenv", "x", "|x|", "x^2")
round(output,4) # results (not shown)
# plot transformation
b = output[1, 3:5]; g = function(x) b[1]*x+b[2]*abs(x)+b[3]*x^2
curve(g, -.04, .04)
```

Example 7.20 Optimal Transformations

In this example, we consider a contrived data set, in which we know the optimal transformation, say, g_0 , and we determine whether the technology can find the transformation when g_0 is not in $\mathcal{G}$. The data, x_t , are generated by the nonlinear model

$$x_t = \exp\{3 \sin(2\pi t\omega_0) + \epsilon_t\}, \quad t = 1, \dots, 500, \quad (7.188)$$

where $\omega_0 = .1$ and ϵ_t is white Gaussian noise with a variance of 16. This example is adapted from Breiman and Friedman (1985), where the ACE algorithm is introduced. The optimal transformation in this case is $g_0(x_t) = \ln(x_t)$, wherein the data are generated from a sinusoid plus noise. Of the 500 generated data, about 96%




Fig. 7.28. Periodogram, in decibels, of the data generated from (7.188) after tapering by a cosine bell.

were less than 4000. Occasionally, the data values were extremely large (the data exceeded 100,000 five times). The spectral estimate of data is shown in Figure 7.28 and provides no evidence of any dominant frequency, including ω_0 . In contrast, the sample spectral envelope (Figure 7.29) computed with respect to $\mathcal{G} = \{x, \sqrt{x}, \sqrt[3]{x}\}$ has no difficulty in isolating ω_0 .

Figure 7.30 compares the estimated optimal transformation with respect to $\mathcal{G}$ with the log transformation for values less than 4000. The estimated transformation at ω_0 is given by

$$g(x) = 4.5 + .0001x - 0.1187\sqrt{x} + .6887\sqrt[3]{x}; \quad (7.189)$$

that is, $\hat{\beta}(\omega_0) = (.0001, -.1187, .6887)'$ after rescaling so (7.189) can be compared directly with $y = \ln(x)$.

Finally, it is worth mentioning the result obtained when the rather inappropriate basis, $\{x, x^2, x^3\}$, was used. Surprisingly, the spectral envelope in this case looks similar to that of Figure 7.29. Also, the resulting estimated optimal transformation at ω_0 , is close to the log transformation for large values of the data. Details may be found in McDougall et al. (1997).

The R code for this example is as follows.

```
set.seed(90210)
u = exp(3*sin(2*pi*1:500*.1) + rnorm(500,0,4)) # the data
spec.pgram(u, spans=c(5,3), taper=.5, log="dB")
dev.new()
x = cbind(u, sqrt(u), u^(1/3)) # transformation set
Var = var(x)
xspect = mvspec(x, spans=c(5,3), taper=.5)
fxr = Re(xspect$fx)
```





Fig. 7.29. Spectral envelope with respect to $\mathcal{G} = \{x, \sqrt{x}, \sqrt[3]{x}\}$ of data generated from (7.188).

```

ev = eigen(Var)
Q = ev$vectors%%diag(1/sqrt(ev$values))%%t(ev$vectors)
num = xspec$n.used
nfreq = length(xspec$freq)
specenv = matrix(0,nfreq,1)
beta=matrix(0,nfreq,3)
for (k in 1:nfreq){
  ev = eigen(2*Q%%fxxr[,k]%%Q/num)
  specenv[k] = ev$values[1]
  b = Q%%ev$vectors[,1]
  beta[k,] = b/sign(b[1]) }
# Output and Graphics
frequency = xspec$freq
plot(frequency, 100*specenv, type="l", ylab="Spectral Envelope (%)")
m = xspec$kernel$m
etainv = sqrt(sum(xspec$kernel[-m:m]^2))
thresh = 100*(2/num)*exp(qnorm(.999)*etainv)*rep(1,nfreq)
lines(frequency,thresh, lty="dashed", col="blue")
output = cbind(frequency, specenv, beta)
colnames(output) = c("freq","specenv","x", "sqrt(x)", "x^(1/3)")
round(output,4)
# Plot Transform
dev.new()
b = output[50,3:5]
g = function(x) 4.5 + b[1]*x+b[2]*sqrt(x)+b[3]*x^(1/3)
curve(g, 1, 4000)
lines(log(1:4000), lty=2)

```




Fig. 7.30. The estimated optimal transformation at ω_0 as given in (7.189) (solid line) and the log transformation, $g(x) = \ln(x)$ (dashed line).

Problems

Section 7.2

7.1 Consider the complex Gaussian distribution for the random variable $\mathbf{X} = \mathbf{X}_c - i\mathbf{X}_s$, as defined in (7.1)-(7.3), where the argument ω_k has been suppressed. Now, the $2p \times 1$ real random variable $\mathbf{Z} = (\mathbf{X}'_c, \mathbf{X}'_s)'$ has a multivariate normal distribution with density

$$p(\mathbf{Z}) = (2\pi)^{-p} |\Sigma|^{-1/2} \exp\left\{-\frac{1}{2}(\mathbf{Z} - \boldsymbol{\mu})' \Sigma^{-1} (\mathbf{Z} - \boldsymbol{\mu})\right\},$$

where $\boldsymbol{\mu} = (\mathbf{M}'_c, \mathbf{M}'_s)'$ is the mean vector. Prove

$$|\Sigma| = \left(\frac{1}{2}\right)^{2p} |C - iQ|^2,$$

using the result that the eigenvectors and eigenvalues of Σ occur in pairs, i.e., $(\mathbf{v}'_c, \mathbf{v}'_s)'$ and $(\mathbf{v}'_s, -\mathbf{v}'_c)'$, where $\mathbf{v}_c - i\mathbf{v}_s$ denotes the eigenvector of f_{xx} . Show that

$$\frac{1}{2}(\mathbf{Z} - \boldsymbol{\mu})' \Sigma^{-1} (\mathbf{Z} - \boldsymbol{\mu}) = (\mathbf{X} - \mathbf{M})^* f^{-1} (\mathbf{X} - \mathbf{M})$$

so $p(\mathbf{X}) = p(\mathbf{Z})$ and we can identify the density of the complex multivariate normal variable $\mathbf{X}$ with that of the real multivariate normal $\mathbf{Z}$.

7.2 Prove $\hat{f}$ in (7.6) maximizes the log likelihood (7.5) by minimizing the negative of the log likelihood

$$L \ln |f| + L \operatorname{tr} \{\hat{f} f^{-1}\}$$

in the form

$$L \sum_i (\lambda_i - \ln \lambda_i - 1) + Lp + L \ln |\hat{f}|,$$

where the λ_i values correspond to the eigenvalues in a simultaneous diagonalization of the matrices f and $\hat{f}$; i.e., there exists a matrix P such that $P^* f P = I$ and $P^* \hat{f} P = \operatorname{diag}(\lambda_1, \dots, \lambda_p) = \Lambda$. Note, $\lambda_i - \ln \lambda_i - 1 \geq 0$ with equality if and only if $\lambda_i = 1$, implying $\Lambda = I$ maximizes the log likelihood and $f = \hat{f}$ is the maximizing value.

Section 7.3

7.3 Verify (7.18) and (7.19) for the mean-squared prediction error MSE in (7.11). Use the orthogonality principle, which implies

$$MSE = E \left[(y_t - \sum_{r=-\infty}^{\infty} \beta'_r x_{t-r}) y_t \right]$$

and gives a set of equations involving the autocovariance functions. Then, use the spectral representations and Fourier transform results to get the final result. Next, consider the predicted series

$$\hat{y}_t = \sum_{r=-\infty}^{\infty} \beta'_r x_{t-r},$$

where β_r satisfies (7.13). Show the ordinary coherence between y_t and $\hat{y}_t$ is exactly the multiple coherence (7.20).

7.4 Consider the complex regression model (7.28) in the form

$$\mathbf{Y} = \mathbf{X}\mathbf{B} + \mathbf{V},$$

where $\mathbf{Y} = (Y_1, Y_2, \dots, Y_L)'$ denotes the observed DFTs after they have been re-indexed and $\mathbf{X} = (\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_L)'$ is a matrix containing the reindexed input vectors. The model is a complex regression model with $\mathbf{Y} = \mathbf{Y}_c - i\mathbf{Y}_s$, $\mathbf{X} = \mathbf{X}_c - i\mathbf{X}_s$, $\mathbf{B} = \mathbf{B}_c - i\mathbf{B}_s$, and $\mathbf{V} = \mathbf{V}_c - i\mathbf{V}_s$ denoting the representation in terms of the usual cosine and sine transforms. Show the partitioned real regression model involving the $2L \times 1$ vector of cosine and sine transforms, say,

$$\begin{pmatrix} \mathbf{Y}_c \\ \mathbf{Y}_s \end{pmatrix} = \begin{pmatrix} \mathbf{X}_c & -\mathbf{X}_s \\ \mathbf{X}_s & \mathbf{X}_c \end{pmatrix} \begin{pmatrix} \mathbf{B}_c \\ \mathbf{B}_s \end{pmatrix} + \begin{pmatrix} \mathbf{V}_c \\ \mathbf{V}_s \end{pmatrix},$$

is isomorphic to the complex regression model in the sense that the real and imaginary parts of the complex model appear as components of the vectors in

the real regression model. Use the usual regression theory to verify (7.27) holds. For example, writing the real regression model as

$$\mathbf{y} = \mathbf{x}\mathbf{b} + \mathbf{v},$$

the isomorphism would imply

$$\begin{aligned} L(\widehat{f}_{yy} - \widehat{f}_{xy}^* \widehat{f}_{xx}^{-1} \widehat{f}_{xy}) &= \mathbf{Y}^* \mathbf{Y} - \mathbf{Y}^* \mathbf{X} (\mathbf{X}^* \mathbf{X})^{-1} \mathbf{X}^* \mathbf{Y} \\ &= \mathbf{y}' \mathbf{y} - \mathbf{y}' \mathbf{x} (\mathbf{x}' \mathbf{x})^{-1} \mathbf{x}' \mathbf{y}. \end{aligned}$$

Section 7.4

7.5 Consider estimating the function

$$\psi_t = \sum_{r=-\infty}^{\infty} \mathbf{a}'_r \boldsymbol{\beta}_{t-r}$$

by a linear filter estimator of the form

$$\widehat{\psi}_t = \sum_{r=-\infty}^{\infty} \mathbf{a}'_r \widehat{\boldsymbol{\beta}}_{t-r},$$

where $\widehat{\boldsymbol{\beta}}_t$ is defined by (7.42). Show a sufficient condition for $\widehat{\psi}_t$ to be an unbiased estimator; i.e., $E \widehat{\psi}_t = \psi_t$, is

$$H(\omega)Z(\omega) = I$$

for all ω . Similarly, show any other unbiased estimator satisfying the above condition has minimum variance (see Shumway and Dean, 1968), so the estimator given is a best linear unbiased (BLUE) estimator.

7.6 Consider a linear model with mean value function μ_t and a signal α_t delayed by an amount τ_j on each sensor, i.e.,

$$y_{jt} = \mu_t + \alpha_{t-\tau_j} + v_{jt}.$$

Show the estimators (7.42) for the mean and the signal are the Fourier transforms of

$$\widehat{M}(\omega) = \frac{Y(\omega) - \overline{\phi(\omega)} B_w(\omega)}{1 - |\phi(\omega)|^2}$$

and

$$\widehat{A}(\omega) = \frac{B_w(\omega) - \phi(\omega) Y(\omega)}{1 - |\phi(\omega)|^2},$$

where

$$\phi(\omega) = \frac{1}{N} \sum_{j=1}^N e^{2\pi i \omega \tau_j}$$

and $B_w(\omega)$ is defined in (7.64).

Section 7.5

7.7 Consider the estimator (7.67) as applied in the context of the random coefficient model (7.65). Prove the filter coefficients for the minimum mean square estimator can be determined from (7.68) and the mean square covariance is given by (7.71).

7.8 For the random coefficient model, verify the expected mean square of the regression power component is

$$\begin{aligned} E[SSR(\omega_k)] &= E[Y^*(\omega_k)Z(\omega_k)S_z^{-1}(\omega_k)Z^*(\omega_k)Y(\omega_k)] \\ &= Lf_\beta(\omega_k)\text{tr}\{S_z(\omega_k)\} + Lqf_v(\omega_k). \end{aligned}$$

Recall, the underlying frequency domain model is

$$\mathbf{Y}(\omega_k) = Z(\omega_k)\mathbf{B}(\omega_k) + \mathbf{V}(\omega_k),$$

where $\mathbf{B}(\omega_k)$ has spectrum $f_\beta(\omega_k)I_q$ and $\mathbf{V}(\omega_k)$ has spectrum $f_v(\omega_k)I_N$ and the two processes are uncorrelated.

Section 7.6

7.9 Suppose we have $I = 2$ groups and the models

$$y_{1jt} = \mu_t + \alpha_{1t} + v_{1jt}$$

for the $j = 1, \dots, N$ observations in group 1 and

$$y_{2jt} = \mu_t + \alpha_{2t} + v_{2jt}$$

for the $j = 1, \dots, N$ observations in group 2, with $\alpha_{1t} + \alpha_{2t} = 0$. Suppose we want to test equality of the two group means; i.e.,

$$y_{ijt} = \mu_t + v_{ijt}, \quad i = 1, 2.$$

- Derive the residual and error power components corresponding to (7.83) and (7.84) for this particular case.
- Verify the forms of the linear compounds involving the mean given in (7.90) and (7.91), using (7.88) and (7.89).
- Show the ratio of the two smoothed spectra in (7.103) has the indicated F -distribution when $f_1(\omega) = f_2(\omega)$. When the spectra are not equal, show the variable is proportional to an F -distribution, where the proportionality constant depends on the ratio of the spectra.

Section 7.7

7.10 The problem of detecting a signal in noise can be considered using the model

$$x_t = s_t + w_t, \quad t = 1, \dots, n,$$

for $p_1(\mathbf{x})$ when a signal is present and the model

$$x_t = w_t, \quad t = 1, \dots, n,$$

for $p_2(\mathbf{x})$ when no signal is present. Under multivariate normality, we might specialize even further by assuming the vector $\mathbf{w} = (w_1, \dots, w_n)'$ has a multivariate normal distribution with mean $\mathbf{0}$ and covariance matrix $\Sigma = \sigma_w^2 I_n$, corresponding to white noise. Assuming the signal vector $\mathbf{s} = (s_1, \dots, s_n)'$ is fixed and known, show the discriminant function (7.112) becomes the matched filter

$$\frac{1}{\sigma_w^2} \sum_{t=1}^n s_t x_t - \frac{1}{2} \left(\frac{S}{N} \right) + \ln \frac{\pi_1}{\pi_2},$$

where

$$\left(\frac{S}{N} \right) = \frac{\sum_{t=1}^n s_t^2}{\sigma_w^2}$$

denotes the signal-to-noise ratio. Give the decision criterion if the prior probabilities are assumed to be the same. Express the false alarm and missed signal probabilities in terms of the normal cdf and the signal-to-noise ratio.

7.11 Assume the same additive signal plus noise representations as in the previous problem, except, the signal is now a random process with a zero mean and covariance matrix $\sigma_s^2 I$. Derive the comparable version of (7.115) as a quadratic detector, and characterize its performance under both hypotheses in terms of constant multiples of the chi-squared distribution.

Section 7.8

7.12 Perform principal component analyses on the stimulus conditions (i) awake-heat and (ii) awake-shock, and compare your results to the results of [Example 7.14](#). Use the data in [fmri](#) and average across subjects.

7.13 For this problem, consider the first three earthquake series (EQ1, EQ2, EQ3) listed in [eqexp](#).

- Estimate and compare the spectral density of the P component and then of the S component for each individual earthquake.
- Estimate and compare the squared coherency between the P and S components of each individual earthquake. Comment on the strength of the coherence.

- (c) Let x_{ti} be the P component of earthquake $i = 1, 2, 3$, and let $\mathbf{x}_t = (x_{t1}, x_{t2}, x_{t3})'$ be the 3×1 vector of P components. Estimate the spectral density, $\lambda_1(\omega)$, of the first principal component series of $\mathbf{x}_t$. Compare this to the corresponding spectra calculated in (a).
- (d) Analogous to part (c), let $\mathbf{y}_t$ denote the 3×1 vector series of S components of the first three earthquakes. Repeat the analysis of part (c) on $\mathbf{y}_t$.

7.14 In the factor analysis model (7.154), let $p = 3$, $q = 1$, and

$$\Sigma_{xx} = \begin{bmatrix} 1 & .4 & .9 \\ .4 & 1 & .7 \\ .9 & .7 & 1 \end{bmatrix}.$$

Show there is a unique choice for $\mathcal{B}$ and D , but $\delta_3^2 < 0$, so the choice is not valid.

7.15 Extend the EM algorithm for classical factor analysis, (7.160)–(7.165), to the time series case of maximizing $\ln L(\mathcal{B}(\omega_j), D_{\epsilon\epsilon}(\omega_j))$ in (7.176). Then, for the data used in Example 7.16, find the approximate maximum likelihood estimates of $\mathcal{B}(\omega_j)$ and $D_{\epsilon\epsilon}(\omega_j)$, and, consequently, Λ_t .

Section 7.9

7.16 Verify, as stated in (7.181), the imaginary part of a $k \times k$ spectral matrix, $f^{im}(\omega)$, is skew symmetric, and then show $\boldsymbol{\beta}' f_{yy}^{im}(\omega) \boldsymbol{\beta} = 0$ for a real $k \times 1$ vector, $\boldsymbol{\beta}$.

7.17 Repeat the analysis of Example 7.18 on BNRF1 of herpesvirus saimiri (the data file is [bnrf1hvs](#)), and compare the results with the results obtained for Epstein–Barr.

7.18 For the NYSE returns, say, r_t , analyzed in Chapter 5, Example 5.5:

- (a) Estimate the spectrum of the r_t . Does the spectral estimate appear to support the hypothesis that the returns are white?
- (b) Examine the possibility of spectral power near the zero frequency for a transformation of the returns, say, $g(r_t)$, using the spectral envelope with Example 7.19 as your guide. Compare the optimal transformation near or at the zero frequency with the usual transformation $y_t = r_t^2$.

Appendix A

Large Sample Theory

A.1 Convergence Modes

The study of the optimality properties of various estimators (such as the sample autocorrelation function) depends, in part, on being able to assess the large-sample behavior of these estimators. We summarize briefly here the kinds of convergence useful in this setting, namely, mean square convergence, convergence in probability, and convergence in distribution.

We consider first a particular class of random variables that plays an important role in the study of second-order time series, namely, the class of random variables belonging to the space L^2 , satisfying $E|x|^2 < \infty$. In proving certain properties of the class L^2 we will often use, for random variables $x, y \in L^2$, the Cauchy–Schwarz inequality,

$$|E(xy)|^2 \leq E(|x|^2)E(|y|^2), \quad (\text{A.1})$$

and the Tchebycheff inequality,

$$P\{|x| \geq a\} \leq \frac{E(|x|^2)}{a^2}, \quad (\text{A.2})$$

for $a > 0$.

Next, we investigate the properties of mean square convergence of random variables in L^2 .

Definition A.1 A sequence of L^2 random variables $\{x_n\}$ is said to converge in **mean square** to a random variable $x \in L^2$, denoted by

$$x_n \xrightarrow{ms} x, \quad (\text{A.3})$$

if and only if

$$E|x_n - x|^2 \rightarrow 0 \quad (\text{A.4})$$

as $n \rightarrow \infty$.

Example A.1 Mean Square Convergence of the Sample Mean

Consider the white noise sequence w_t and the *signal plus noise* series

$$x_t = \mu + w_t.$$

Then, because

$$E|\bar{x}_n - \mu|^2 = \frac{\sigma_w^2}{n} \rightarrow 0$$

as $n \rightarrow \infty$, where $\bar{x}_n = n^{-1} \sum_{t=1}^n x_t$ is the sample mean, we have $\bar{x}_n \xrightarrow{ms} \mu$.

We summarize some of the properties of mean square convergence as follows. If $x_n \xrightarrow{ms} x$, and $y_n \xrightarrow{ms} y$, then, as $n \rightarrow \infty$,

$$E(x_n) \rightarrow E(x); \quad (\text{A.5})$$

$$E(|x_n|^2) \rightarrow E(|x|^2); \quad (\text{A.6})$$

$$E(x_n y_n) \rightarrow E(xy). \quad (\text{A.7})$$

We also note the L^2 completeness theorem known as the *Riesz–Fisher Theorem*.

Theorem A.1 *Let $\{x_n\}$ be a sequence in L^2 . Then, there exists an x in L^2 such that $x_n \xrightarrow{ms} x$ if and only if*

$$E|x_n - x_m|^2 \rightarrow 0 \quad (\text{A.8})$$

for $m, n \rightarrow \infty$.

Often the condition of **Theorem A.1** is easier to verify to establish that a mean square limit x exists without knowing what it is. Sequences that satisfy (A.8) are said to be Cauchy sequences in L^2 and (A.8) is also known as the Cauchy criterion for L^2 .

Example A.2 Time Invariant Linear Filter

As an important example of the use of the Riesz–Fisher Theorem and the properties of mean square convergent series given in (A.5)–(A.7), a time-invariant linear filter is defined as a convolution of the form

$$y_t = \sum_{j=-\infty}^{\infty} a_j x_{t-j} \quad (\text{A.9})$$

for each $t = 0, \pm 1, \pm 2, \dots$, where x_t is a weakly stationary input series with mean μ_x and autocovariance function $\gamma_x(h)$, and a_j , for $j = 0, \pm 1, \pm 2, \dots$ are constants satisfying

$$\sum_{j=-\infty}^{\infty} |a_j| < \infty. \quad (\text{A.10})$$

The output series y_t defines a filtering or smoothing of the input series that changes the character of the time series in a predictable way. We need to know the conditions under which the outputs y_t in (A.9) and the linear process (1.29) exist.

Considering the sequence

$$y_t^n = \sum_{j=-n}^n a_j x_{t-j}, \quad (\text{A.11})$$

$n = 1, 2, \dots$, we need to show first that y_t^n has a mean square limit. By [Theorem A.1](#), it is enough to show that

$$E |y_t^n - y_t^m|^2 \rightarrow 0$$

as $m, n \rightarrow \infty$. For $n > m > 0$,

$$\begin{aligned} E |y_t^n - y_t^m|^2 &= E \left| \sum_{m < |j| \leq n} a_j x_{t-j} \right|^2 \\ &= \sum_{m < |j| \leq n} \sum_{m \leq |k| \leq n} a_j a_k E(x_{t-j} x_{t-k}) \\ &\leq \sum_{m < |j| \leq n} \sum_{m \leq |k| \leq n} |a_j| |a_k| |E(x_{t-j} x_{t-k})| \\ &\leq \sum_{m < |j| \leq n} \sum_{m \leq |k| \leq n} |a_j| |a_k| (E|x_{t-j}|^2)^{1/2} (E|x_{t-k}|^2)^{1/2} \\ &= [\gamma_x(0) + \mu_x^2] \left(\sum_{m \leq |j| \leq n} |a_j| \right)^2 \rightarrow 0 \end{aligned}$$

as $m, n \rightarrow \infty$, because $\gamma_x(0)$ is a constant and $\{a_j\}$ is absolutely summable (the second inequality follows from the Cauchy–Schwarz inequality).

Although we know that the sequence $\{y_t^n\}$ given by (A.11) converges in mean square, we have not established its mean square limit. It should be obvious, however, that $y_t^n \xrightarrow{ms} y_t$ as $n \rightarrow \infty$, where y_t is given by (A.9).¹

Finally, we may use (A.5) and (A.7) to establish the mean, μ_y and autocovariance function, $\gamma_y(h)$ of y_t . In particular we have,

$$\mu_y = \mu_x \sum_{j=-\infty}^{\infty} a_j, \quad (\text{A.12})$$

and

$$\begin{aligned} \gamma_y(h) &= E \sum_{j=-\infty}^{\infty} \sum_{k=-\infty}^{\infty} a_j (x_{t+h-j} - \mu_x) a_k (x_{t-k} - \mu_x) \\ &= \sum_{j=-\infty}^{\infty} \sum_{k=-\infty}^{\infty} a_j \gamma_x(h - j + k) a_k. \end{aligned} \quad (\text{A.13})$$

A second important kind of convergence is convergence in probability.

¹ If S denotes the mean square limit of y_t^n , then using Fatou's Lemma, $E |S - y_t|^2 = E \liminf_{n \rightarrow \infty} |S - y_t^n|^2 \leq \liminf_{n \rightarrow \infty} E |S - y_t^n|^2 = 0$, which establishes that y_t is the mean square limit of y_t^n .

Definition A.2 The sequence $\{x_n\}$, for $n = 1, 2, \dots$, **converges in probability** to a random variable x , denoted by

$$x_n \xrightarrow{P} x, \quad (\text{A.14})$$

if and only if

$$P\{|x_n - x| > \epsilon\} \rightarrow 0 \quad (\text{A.15})$$

for all $\epsilon > 0$, as $n \rightarrow \infty$.

An immediate consequence of the Tchebycheff inequality, (A.2), is that

$$P\{|x_n - x| \geq \epsilon\} \leq \frac{E(|x_n - x|^2)}{\epsilon^2},$$

so convergence in mean square implies convergence in probability, i.e.,

$$x_n \xrightarrow{ms} x \Rightarrow x_n \xrightarrow{P} x. \quad (\text{A.16})$$

This result implies, for example, that the filter (A.9) exists as a limit in probability because it converges in mean square [it is also easily established that (A.9) exists with probability one]. We mention, at this point, the useful Weak Law of Large Numbers which states that, for an independent identically distributed sequence x_n of random variables with mean μ , we have

$$\bar{x}_n \xrightarrow{P} \mu \quad (\text{A.17})$$

as $n \rightarrow \infty$, where $\bar{x}_n = n^{-1} \sum_{t=1}^n x_t$ is the usual sample mean.

We also will make use of the following concepts.

Definition A.3 For order in probability we write

$$x_n = o_p(a_n) \quad (\text{A.18})$$

if and only if

$$\frac{x_n}{a_n} \xrightarrow{P} 0. \quad (\text{A.19})$$

The term **boundedness in probability**, written $x_n = O_p(a_n)$, means that for every $\epsilon > 0$, there exists a $\delta(\epsilon) > 0$ such that

$$P\left\{\left|\frac{x_n}{a_n}\right| > \delta(\epsilon)\right\} \leq \epsilon \quad (\text{A.20})$$

for all n .

Under this convention, e.g., the notation for $x_n \xrightarrow{P} x$ becomes $x_n - x = o_p(1)$. The definitions can be compared with their nonrandom counterparts, namely, for a fixed sequence $x_n = o(1)$ if $x_n \rightarrow 0$ and $x_n = O(1)$ if x_n , for $n = 1, 2, \dots$ is bounded. Some handy properties of $o_p(\cdot)$ and $O_p(\cdot)$ are as follows.

- (i) If $x_n = o_p(a_n)$ and $y_n = o_p(b_n)$, then $x_n y_n = o_p(a_n b_n)$ and $x_n + y_n = o_p(\max(a_n, b_n))$.

- (ii) If $x_n = o_p(a_n)$ and $y_n = O_p(b_n)$, then $x_n y_n = o_p(a_n b_n)$.
 (iii) Statement (i) is true if $O_p(\cdot)$ replaces $o_p(\cdot)$.

Example A.3 Convergence and Order in Probability for the Sample Mean

For the sample mean, $\bar{x}_n$, of iid random variables with mean μ and variance σ^2 , by the Tchebycheff inequality,

$$\begin{aligned} P\{|\bar{x}_n - \mu| > \epsilon\} &\leq \frac{E[(\bar{x}_n - \mu)^2]}{\epsilon^2} \\ &= \frac{\sigma^2}{n\epsilon^2} \rightarrow 0, \end{aligned}$$

as $n \rightarrow \infty$. It follows that $\bar{x}_n \xrightarrow{P} \mu$, or $\bar{x}_n - \mu = o_p(1)$. To find the rate, it follows that, for $\delta(\epsilon) > 0$,

$$P\left\{\sqrt{n} |\bar{x}_n - \mu| > \delta(\epsilon)\right\} \leq \frac{\sigma^2/n}{\delta^2(\epsilon)/n} = \frac{\sigma^2}{\delta^2(\epsilon)}$$

by Tchebycheff's inequality, so taking $\epsilon = \sigma^2/\delta^2(\epsilon)$ shows that $\delta(\epsilon) = \sigma/\sqrt{\epsilon}$ does the job and

$$\bar{x}_n - \mu = O_p(n^{-1/2}).$$

For $k \times 1$ random vectors $\mathbf{x}_n$, convergence in probability, written $\mathbf{x}_n \xrightarrow{P} \mathbf{x}$ or $\mathbf{x}_n - \mathbf{x} = o_p(1)$ is defined as element-by-element convergence in probability, or equivalently, as convergence in terms of the Euclidean distance

$$\|\mathbf{x}_n - \mathbf{x}\| \xrightarrow{P} 0, \quad (\text{A.21})$$

where $\|\mathbf{a}\| = \sum_j a_j^2$ for any vector $\mathbf{a}$. In this context, we note the result that if $\mathbf{x}_n \xrightarrow{P} \mathbf{x}$ and $g(\mathbf{x}_n)$ is a continuous mapping,

$$g(\mathbf{x}_n) \xrightarrow{P} g(\mathbf{x}). \quad (\text{A.22})$$

Furthermore, if $\mathbf{x}_n - \mathbf{a} = O_p(\delta_n)$ with $\delta_n \rightarrow 0$ and $g(\cdot)$ is a function with continuous first derivatives continuous in a neighborhood of $\mathbf{a} = (a_1, a_2, \dots, a_k)'$, we have the Taylor series expansion in probability

$$g(\mathbf{x}_n) = g(\mathbf{a}) + \left. \frac{\partial g(\mathbf{x})}{\partial \mathbf{x}} \right|_{\mathbf{x}=\mathbf{a}}' (\mathbf{x}_n - \mathbf{a}) + O_p(\delta_n), \quad (\text{A.23})$$

where

$$\left. \frac{\partial g(\mathbf{x})}{\partial \mathbf{x}} \right|_{\mathbf{x}=\mathbf{a}} = \left(\left. \frac{\partial g(\mathbf{x})}{\partial x_1} \right|_{\mathbf{x}=\mathbf{a}}, \dots, \left. \frac{\partial g(\mathbf{x})}{\partial x_k} \right|_{\mathbf{x}=\mathbf{a}} \right)'$$

denotes the vector of partial derivatives with respect to $x_1, x_2, \dots, x_k$, evaluated at $\mathbf{a}$. This result remains true if $O_p(\delta_n)$ is replaced everywhere by $o_p(\delta_n)$.

Example A.4 Expansion for the Logarithm of the Sample Mean

With the same conditions as **Example A.3**, consider $g(\bar{x}_n) = \log \bar{x}_n$, which has a derivative at μ , for $\mu > 0$. Then, because $\bar{x}_n - \mu = O_p(n^{-1/2})$ from **Example A.3**, the conditions for the Taylor expansion in probability, (A.23), are satisfied and we have

$$\log \bar{x}_n = \log \mu + \mu^{-1}(\bar{x}_n - \mu) + O_p(n^{-1/2}).$$

The large sample distributions of sample mean and sample autocorrelation functions defined earlier can be developed using the notion of convergence in distribution.

Definition A.4 A sequence of $k \times 1$ random vectors $\{\mathbf{x}_n\}$ is said to **converge in distribution**, written

$$\mathbf{x}_n \xrightarrow{d} \mathbf{x} \quad (\text{A.24})$$

if and only if

$$F_n(\mathbf{x}) \rightarrow F(\mathbf{x}) \quad (\text{A.25})$$

at the continuity points of distribution function $F(\cdot)$.

Example A.5 Convergence in Distribution

Consider a sequence $\{x_n\}$ of iid normal random variables with mean zero and variance $1/n$. Using the standard normal cdf, $\Phi(z) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^z \exp\{-\frac{1}{2}u^2\} du$, we have $F_n(z) = \Phi(\sqrt{n}z)$, so

$$F_n(z) \rightarrow \begin{cases} 0 & z < 0, \\ 1/2 & z = 0 \\ 1 & z > 0 \end{cases}$$

and we may take

$$F(z) = \begin{cases} 0 & z < 0, \\ 1 & z \geq 0, \end{cases}$$

because the point where the two functions differ is not a continuity point of $F(z)$.

The distribution function relates uniquely to the characteristic function through the Fourier transform, defined as a function with vector argument $\boldsymbol{\lambda} = (\lambda_1, \lambda_2, \dots, \lambda_k)'$, say

$$\phi(\boldsymbol{\lambda}) = E(\exp\{i\boldsymbol{\lambda}'\mathbf{x}\}) = \int \exp\{i\boldsymbol{\lambda}'\mathbf{x}\} dF(\mathbf{x}). \quad (\text{A.26})$$

Hence, for a sequence $\{\mathbf{x}_n\}$ we may characterize convergence in distribution of $F_n(\cdot)$ in terms of convergence of the sequence of characteristic functions $\phi_n(\cdot)$, i.e.,

$$\phi_n(\boldsymbol{\lambda}) \rightarrow \phi(\boldsymbol{\lambda}) \Leftrightarrow F_n(\mathbf{x}) \xrightarrow{d} F(\mathbf{x}), \quad (\text{A.27})$$

where $\Leftrightarrow$ means that the implication goes both directions. In this connection, we have

Proposition A.1 The Cramér–Wold device. *Let $\{\mathbf{x}_n\}$ be a sequence of $k \times 1$ random vectors. Then, for every $\mathbf{c} = (c_1, c_2, \dots, c_k)' \in \mathbb{R}^k$*

$$\mathbf{c}'\mathbf{x}_n \xrightarrow{d} \mathbf{c}'\mathbf{x} \Leftrightarrow \mathbf{x}_n \xrightarrow{d} \mathbf{x}. \quad (\text{A.28})$$

Proposition A.1 can be useful because sometimes it is easier to show the convergence in distribution of $\mathbf{c}'\mathbf{x}_n$ than $\mathbf{x}_n$ directly.

Convergence in probability implies convergence in distribution, namely,

$$\mathbf{x}_n \xrightarrow{p} \mathbf{x} \Rightarrow \mathbf{x}_n \xrightarrow{d} \mathbf{x}, \quad (\text{A.29})$$

but the converse is only true when $\mathbf{x}_n \xrightarrow{d} \mathbf{c}$, where $\mathbf{c}$ is a constant vector. If $\mathbf{x}_n \xrightarrow{d} \mathbf{x}$ and $\mathbf{y}_n \xrightarrow{d} \mathbf{c}$ are two sequences of random vectors and $\mathbf{c}$ is a constant vector,

$$\mathbf{x}_n + \mathbf{y}_n \xrightarrow{d} \mathbf{x} + \mathbf{c} \quad \text{and} \quad \mathbf{y}_n'\mathbf{x}_n \xrightarrow{d} \mathbf{c}'\mathbf{x}. \quad (\text{A.30})$$

For a continuous mapping $h(\mathbf{x})$,

$$\mathbf{x}_n \xrightarrow{d} \mathbf{x} \Rightarrow h(\mathbf{x}_n) \xrightarrow{d} h(\mathbf{x}). \quad (\text{A.31})$$

A number of results in time series depend on making a series of approximations to prove convergence in distribution. For example, we have that if $\mathbf{x}_n \xrightarrow{d} \mathbf{x}$ can be approximated by the sequence $\mathbf{y}_n$ in the sense that

$$\mathbf{y}_n - \mathbf{x}_n = o_p(1), \quad (\text{A.32})$$

then we have that $\mathbf{y}_n \xrightarrow{d} \mathbf{x}$, so the approximating sequence $\mathbf{y}_n$ has the same limiting distribution as $\mathbf{x}$. We present the following Basic Approximation Theorem (BAT) that will be used later to derive asymptotic distributions for the sample mean and ACF.

Theorem A.2 [Basic Approximation Theorem (BAT)] *Let $\mathbf{x}_n$ for $n = 1, 2, \dots$, and $\mathbf{y}_{mn}$ for $m = 1, 2, \dots$, be random $k \times 1$ vectors such that*

- (i) $\mathbf{y}_{mn} \xrightarrow{d} \mathbf{y}_m$ as $n \rightarrow \infty$ for each m ;
- (ii) $\mathbf{y}_m \xrightarrow{d} \mathbf{y}$ as $m \rightarrow \infty$;
- (iii) $\lim_{m \rightarrow \infty} \limsup_{n \rightarrow \infty} P\{|\mathbf{x}_n - \mathbf{y}_{mn}| > \epsilon\} = 0$ for every $\epsilon > 0$.

Then, $\mathbf{x}_n \xrightarrow{d} \mathbf{y}$.

As a practical matter, the BAT condition (iii) is implied by the Tchebycheff inequality if

$$(iii') \quad E\{|\mathbf{x}_n - \mathbf{y}_{mn}|^2\} \rightarrow 0 \quad (\text{A.33})$$

as $m, n \rightarrow \infty$, and (iii') is often much easier to establish than (iii).

The theorem allows approximation of the underlying sequence in two steps, through the intermediary sequence $\mathbf{y}_{mn}$, depending on two arguments. In the time series case, n is generally the sample length and m is generally the number of terms in an approximation to the linear process of the form (A.11).

Proof: The proof of the theorem is a simple exercise in using the characteristic functions and appealing to (A.27). We need to show

$$|\phi_{\mathbf{x}_n} - \phi_{\mathbf{y}}| \rightarrow 0,$$

where we use the shorthand notation $\phi \equiv \phi(\boldsymbol{\lambda})$ for ease. First,

$$|\phi_{\mathbf{x}_n} - \phi_{\mathbf{y}}| \leq |\phi_{\mathbf{x}_n} - \phi_{\mathbf{y}_{mn}}| + |\phi_{\mathbf{y}_{mn}} - \phi_{\mathbf{y}_m}| + |\phi_{\mathbf{y}_m} - \phi_{\mathbf{y}}|. \quad (\text{A.34})$$

By the condition (ii) and (A.27), the last term converges to zero, and by condition (i) and (A.27), the second term converges to zero and we only need consider the first term in (A.34). Now, write

$$\begin{aligned} |\phi_{\mathbf{x}_n} - \phi_{\mathbf{y}_{mn}}| &= \left| E(e^{i\boldsymbol{\lambda}'\mathbf{x}_n} - e^{i\boldsymbol{\lambda}'\mathbf{y}_{mn}}) \right| \\ &\leq E \left| e^{i\boldsymbol{\lambda}'\mathbf{x}_n} (1 - e^{i\boldsymbol{\lambda}'(\mathbf{y}_{mn} - \mathbf{x}_n)}) \right| \\ &= E \left| 1 - e^{i\boldsymbol{\lambda}'(\mathbf{y}_{mn} - \mathbf{x}_n)} \right| \\ &= E \left\{ \left| 1 - e^{i\boldsymbol{\lambda}'(\mathbf{y}_{mn} - \mathbf{x}_n)} \right| I\{|\mathbf{y}_{mn} - \mathbf{x}_n| < \delta\} \right\} \\ &\quad + E \left\{ \left| 1 - e^{i\boldsymbol{\lambda}'(\mathbf{y}_{mn} - \mathbf{x}_n)} \right| I\{|\mathbf{y}_{mn} - \mathbf{x}_n| \geq \delta\} \right\}, \end{aligned}$$

where $\delta > 0$ and $I\{A\}$ denotes the indicator function of the set A . Then, given $\boldsymbol{\lambda}$ and $\epsilon > 0$, choose $\delta(\epsilon) > 0$ such that

$$\left| 1 - e^{i\boldsymbol{\lambda}'(\mathbf{y}_{mn} - \mathbf{x}_n)} \right| < \epsilon$$

if $|\mathbf{y}_{mn} - \mathbf{x}_n| < \delta$, and the first term is less than ϵ , an arbitrarily small constant. For the second term, note that

$$\left| 1 - e^{i\boldsymbol{\lambda}'(\mathbf{y}_{mn} - \mathbf{x}_n)} \right| \leq 2$$

and we have

$$E \left\{ \left| 1 - e^{i\boldsymbol{\lambda}'(\mathbf{y}_{mn} - \mathbf{x}_n)} \right| I\{|\mathbf{y}_{mn} - \mathbf{x}_n| \geq \delta\} \right\} \leq 2P\{|\mathbf{y}_{mn} - \mathbf{x}_n| \geq \delta\},$$

which converges to zero as $n \rightarrow \infty$ by property (iii). $\square$

A.2 Central Limit Theorems

We will generally be concerned with the large-sample properties of estimators that turn out to be normally distributed as $n \rightarrow \infty$.

Definition A.5 A sequence of random variables $\{x_n\}$ is said to be **asymptotically normal** with mean μ_n and variance σ_n^2 if, as $n \rightarrow \infty$,

$$\sigma_n^{-1}(x_n - \mu_n) \xrightarrow{d} z,$$

where z has the standard normal distribution. We shall abbreviate this as

$$x_n \sim AN(\mu_n, \sigma_n^2), \quad (\text{A.35})$$

where $\sim$ will denote is distributed as.

We state the important Central Limit Theorem, as follows.

Theorem A.3 Let $x_1, \dots, x_n$ be independent and identically distributed with mean μ and variance σ^2 . If $\bar{x}_n = (x_1 + \dots + x_n)/n$ denotes the sample mean, then

$$\bar{x}_n \sim AN(\mu, \sigma^2/n). \quad (\text{A.36})$$

Often, we will be concerned with a sequence of $k \times 1$ vectors $\{\mathbf{x}_n\}$. The following property is motivated by the Cramér–Wold device, [Proposition A.1](#).

Proposition A.2 A sequence of random vectors is asymptotically normal, i.e.,

$$\mathbf{x}_n \sim AN(\boldsymbol{\mu}_n, \boldsymbol{\Sigma}_n), \quad (\text{A.37})$$

if and only if

$$\mathbf{c}'\mathbf{x}_n \sim AN(\mathbf{c}'\boldsymbol{\mu}_n, \mathbf{c}'\boldsymbol{\Sigma}_n\mathbf{c}) \quad (\text{A.38})$$

for all $\mathbf{c} \in \mathbb{R}^k$ and $\boldsymbol{\Sigma}_n$ is positive definite.

In order to begin to consider what happens for dependent data in the limiting case, it is necessary to define, first of all, a particular kind of dependence known as M -dependence. We say that a time series x_t is M -dependent if the set of values $x_s, s \leq t$ is independent of the set of values $x_s, s \geq t + M + 1$, so time points separated by more than M units are independent. A central limit theorem for such dependent processes, used in conjunction with the Basic Approximation Theorem, will allow us to develop large-sample distributional results for the sample mean $\bar{x}$ and the sample ACF $\hat{\rho}_x(h)$ in the stationary case.

In the arguments that follow, we often make use of the formula for the variance of $\bar{x}_n$ in the stationary case, namely,

$$\text{var } \bar{x}_n = n^{-1} \sum_{u=-(n-1)}^{(n-1)} \left(1 - \frac{|u|}{n}\right) \gamma(u), \quad (\text{A.39})$$

which was established in (1.33) on page 27. We shall also use the fact that, for

$$\sum_{u=-\infty}^{\infty} |\gamma(u)| < \infty,$$

we would have, by dominated convergence,²

$$n \operatorname{var} \bar{x}_n \rightarrow \sum_{u=-\infty}^{\infty} \gamma(u), \quad (\text{A.40})$$

because $|(1 - |u|/n)\gamma(u)| \leq |\gamma(u)|$ and $(1 - |u|/n)\gamma(u) \rightarrow \gamma(u)$. We may now state the M-Dependent Central Limit Theorem as follows.

Theorem A.4 *If x_t is a strictly stationary M-dependent sequence of random variables with mean zero and autocovariance function $\gamma(\cdot)$ and if*

$$V_M = \sum_{u=-M}^M \gamma(u), \quad (\text{A.41})$$

where $V_M \neq 0$,

$$\bar{x}_n \sim AN(0, V_M/n). \quad (\text{A.42})$$

Proof: To prove the theorem, using **Theorem A.2**, the Basic Approximation Theorem, we may construct a sequence of variables y_{mn} approximating

$$n^{1/2} \bar{x}_n = n^{-1/2} \sum_{t=1}^n x_t$$

in the dependent case and then simply verify conditions (i), (ii), and (iii) of **Theorem A.2**. For $m > 2M$, we may first consider the approximation

$$\begin{aligned} y_{mn} &= n^{-1/2}[(x_1 + \cdots + x_{m-M}) + (x_{m+1} + \cdots + x_{2m-M}) \\ &\quad + (x_{2m+1} + \cdots + x_{3m-M}) + \cdots + (x_{(r-1)m+1} + \cdots + x_{rm-M})] \\ &= n^{-1/2}(z_1 + z_2 + \cdots + z_r), \end{aligned}$$

where $r = [n/m]$, with $[n/m]$ denoting the greatest integer less than or equal to n/m . This approximation contains only part of $n^{1/2} \bar{x}_n$, but the random variables $z_1, z_2, \dots, z_r$ are independent because they are separated by more than M time points, e.g., $m+1 - (m-M) = M+1$ points separate z_1 and z_2 . Because of strict stationarity, $z_1, z_2, \dots, z_r$ are identically distributed with zero means and variances

$$S_{m-M} = \sum_{|u| \leq M} (m - M - |u|)\gamma(u)$$

by a computation similar to that producing (A.39). We now verify the conditions of the Basic Approximation Theorem hold.

² Dominated convergence technically relates to convergent sequences (with respect to a sigma-additive measure μ) of measurable functions $f_n \rightarrow f$ bounded by an integrable function g , $\int g \, d\mu < \infty$. For such a sequence,

$$\int f_n \, d\mu \rightarrow \int f \, d\mu.$$

For the case in point, take $f_n(u) = (1 - |u|/n)\gamma(u)$ for $|u| < n$ and as zero for $|u| \geq n$. Take $\mu(u) = 1, u = \pm 1, \pm 2, \dots$ to be counting measure.

(i) Applying the Central Limit Theorem to the sum y_{mn} gives

$$y_{mn} = n^{-1/2} \sum_{i=1}^r z_i = (n/r)^{-1/2} r^{-1/2} \sum_{i=1}^r z_i.$$

Because $(n/r)^{-1/2} \rightarrow m^{1/2}$ and

$$r^{-1/2} \sum_{i=1}^r z_i \xrightarrow{d} N(0, S_{m-M}),$$

it follows from (A.30) that

$$y_{mn} \xrightarrow{d} y_m \sim N(0, S_{m-M}/m).$$

as $n \rightarrow \infty$, for a fixed m .

(ii) Note that as $m \rightarrow \infty$, $S_{m-M}/m \rightarrow V_M$ using dominated convergence, where V_M is defined in (A.41). Hence, the characteristic function of y_m , say,

$$\phi_m(\lambda) = \exp\left\{-\frac{1}{2}\lambda^2 \frac{S_{m-M}}{m}\right\} \rightarrow \exp\left\{-\frac{1}{2}\lambda^2 V_M\right\},$$

as $m \rightarrow \infty$, which is the characteristic function of a random variable $y \sim N(0, V_M)$ and the result follows because of (A.27).

(iii) To verify the last condition of the BAT theorem,

$$\begin{aligned} n^{1/2}\bar{x}_n - y_{mn} &= n^{-1/2}[(x_{m-M+1} + \cdots + x_m) \\ &\quad + (x_{2m-M+1} + \cdots + x_{2m}) \\ &\quad + (x_{(r-1)m-M+1} + \cdots + x_{(r-1)m}) \\ &\quad \vdots \\ &\quad + (x_{rm-M+1} + \cdots + x_n)] \\ &= n^{-1/2}(w_1 + w_2 + \cdots + w_r), \end{aligned}$$

so the error is expressed as a scaled sum of iid variables with variance S_M for the first $r-1$ variables and

$$\begin{aligned} \text{var}(w_r) &= \sum_{|u| \leq m-M} \left(n - [n/m]m + M - |u| \right) \gamma(u) \\ &\leq \sum_{|u| \leq m-M} (m + M - |u|) \gamma(u). \end{aligned}$$

Hence,

$$\text{var}[n^{1/2}\bar{x} - y_{mn}] = n^{-1}[(r-1)S_M + \text{var } w_r],$$

which converges to $m^{-1}S_M$ as $n \rightarrow \infty$. Because $m^{-1}S_M \rightarrow 0$ as $m \rightarrow \infty$, the condition of (iii) holds by the Tchebycheff inequality.

□

A.3 The Mean and Autocorrelation Functions

The background material in the previous two sections can be used to develop the asymptotic properties of the sample mean and ACF used to evaluate statistical significance. In particular, we are interested in verifying **Property 1.1**.

We begin with the distribution of the sample mean $\bar{x}_n$, noting that (A.40) suggests a form for the limiting variance. In all of the asymptotics, we will use the assumption that x_t is a linear process, as defined in **Definition 1.12**, but with the added condition that $\{w_t\}$ is iid. That is, throughout this section we assume

$$x_t = \mu_x + \sum_{j=-\infty}^{\infty} \psi_j w_{t-j} \quad (\text{A.43})$$

where $w_t \sim \text{iid}(0, \sigma_w^2)$, and the coefficients satisfy

$$\sum_{j=-\infty}^{\infty} |\psi_j| < \infty. \quad (\text{A.44})$$

Before proceeding further, we should note that the exact sampling distribution of $\bar{x}_n$ is available if the distribution of the underlying vector $\mathbf{x} = (x_1, x_2, \dots, x_n)'$ is multivariate normal. Then, $\bar{x}_n$ is just a linear combination of jointly normal variables that will have the normal distribution

$$\bar{x}_n \sim N\left(\mu_x, n^{-1} \sum_{|u| < n} \left(1 - \frac{|u|}{n}\right) \gamma_x(u)\right), \quad (\text{A.45})$$

by (A.39). In the case where x_t are not jointly normally distributed, we have the following theorem.

Theorem A.5 *If x_t is a linear process of the form (A.43) and $\sum_j \psi_j \neq 0$, then*

$$\bar{x}_n \sim AN(\mu_x, n^{-1}V), \quad (\text{A.46})$$

where

$$V = \sum_{h=-\infty}^{\infty} \gamma_x(h) = \sigma_w^2 \left(\sum_{j=-\infty}^{\infty} \psi_j \right)^2 \quad (\text{A.47})$$

and $\gamma_x(\cdot)$ is the autocovariance function of x_t .

Proof: To prove the above, we can again use the Basic Approximation Theorem, **Theorem A.2**, by first defining the strictly stationary $2m$ -dependent linear process with finite limits

$$x_t^m = \sum_{j=-m}^m \psi_j w_{t-j}$$

as an approximation to x_t to use in the approximating mean

$$\bar{x}_{n,m} = n^{-1} \sum_{t=1}^n x_t^m.$$

Then, take

$$y_{mn} = n^{1/2}(\bar{x}_{n,m} - \mu_x)$$

as an approximation to $n^{1/2}(\bar{x}_n - \mu_x)$.

(i) Applying **Theorem A.4**, we have

$$y_{mn} \xrightarrow{d} y_m \sim N(0, V_m),$$

as $n \rightarrow \infty$, where

$$V_m = \sum_{h=-2m}^{2m} \gamma_x(h) = \sigma_w^2 \left(\sum_{j=-m}^m \psi_j \right)^2.$$

To verify the above, we note that for the general linear process with infinite limits, (1.30) implies that

$$\sum_{h=-\infty}^{\infty} \gamma_x(h) = \sigma_w^2 \sum_{h=-\infty}^{\infty} \sum_{j=-\infty}^{\infty} \psi_{j+h} \psi_j = \sigma_w^2 \left(\sum_{j=-\infty}^{\infty} \psi_j \right)^2,$$

so taking the special case $\psi_j = 0$, for $|j| > m$, we obtain V_m .

(ii) Because $V_m \rightarrow V$ in (A.47) as $m \rightarrow \infty$, we may use the same characteristic function argument as under (ii) in the proof of **Theorem A.4** to note that

$$y_m \xrightarrow{d} y \sim N(0, V),$$

where V is given by (A.47).

(iii) Finally,

$$\begin{aligned} \text{var} \left[n^{1/2}(\bar{x}_n - \mu_x) - y_{mn} \right] &= n \text{var} \left[n^{-1} \sum_{t=1}^n \sum_{|j|>m} \psi_j w_{t-j} \right] \\ &= \sigma_w^2 \left(\sum_{|j|>m} \psi_j \right)^2 \rightarrow 0 \end{aligned}$$

as $m \rightarrow \infty$.

□

In order to develop the sampling distribution of the sample autocovariance function, $\hat{\gamma}_x(h)$, and the sample autocorrelation function, $\hat{\rho}_x(h)$, we need to develop some idea as to the mean and variance of $\hat{\gamma}_x(h)$ under some reasonable assumptions. These computations for $\hat{\gamma}_x(h)$ are messy, and we consider a comparable quantity

$$\tilde{\gamma}_x(h) = n^{-1} \sum_{t=1}^n (x_{t+h} - \mu_x)(x_t - \mu_x) \quad (\text{A.48})$$

as an approximation. By [Problem 1.30](#),

$$n^{1/2}[\tilde{\gamma}_x(h) - \widehat{\gamma}_x(h)] = o_p(1),$$

so that limiting distributional results proved for $n^{1/2}\tilde{\gamma}_x(h)$ will hold for $n^{1/2}\widehat{\gamma}_x(h)$ by [\(A.32\)](#).

We begin by proving formulas for the variance and for the limiting variance of $\tilde{\gamma}_x(h)$ under the assumptions that x_t is a linear process of the form [\(A.43\)](#), satisfying [\(A.44\)](#) with the white noise variates w_t having variance σ_w^2 as before, but also required to have fourth moments satisfying

$$E(w_t^4) = \eta\sigma_w^4 < \infty, \quad (\text{A.49})$$

where η is some constant. We seek results comparable with [\(A.39\)](#) and [\(A.40\)](#) for $\tilde{\gamma}_x(h)$. To ease the notation, we will henceforth drop the subscript x from the notation.

Using [\(A.48\)](#), $E[\tilde{\gamma}(h)] = \gamma(h)$. Under the above assumptions, we show now that, for $p, q = 0, 1, 2, \dots$,

$$\text{cov} [\tilde{\gamma}(p), \tilde{\gamma}(q)] = n^{-1} \sum_{u=-(n-1)}^{(n-1)} \left(1 - \frac{|u|}{n}\right) V_u, \quad (\text{A.50})$$

where

$$\begin{aligned} V_u &= \gamma(u)\gamma(u+p-q) + \gamma(u+p)\gamma(u-q) \\ &\quad + (\eta-3)\sigma_w^4 \sum_i \psi_{i+u+q}\psi_{i+u}\psi_{i+p}\psi_i. \end{aligned} \quad (\text{A.51})$$

The absolute summability of the ψ_j can then be shown to imply the absolute summability of the V_u .³ Thus, the dominated convergence theorem implies

$$\begin{aligned} n \text{cov} [\tilde{\gamma}(p), \tilde{\gamma}(q)] &\rightarrow \sum_{u=-\infty}^{\infty} V_u \\ &= (\eta-3)\gamma(p)\gamma(q) \\ &\quad + \sum_{u=-\infty}^{\infty} \left[\gamma(u)\gamma(u+p-q) + \gamma(u+p)\gamma(u-q) \right]. \end{aligned} \quad (\text{A.52})$$

To verify [\(A.50\)](#) is somewhat tedious, so we only go partially through the calculations, leaving the repetitive details to the reader. First, rewrite [\(A.43\)](#) as

$$x_t = \mu + \sum_{i=-\infty}^{\infty} \psi_{t-i} w_i,$$

³ Note: $\sum_{j=-\infty}^{\infty} |a_j| < \infty$ and $\sum_{j=-\infty}^{\infty} |b_j| < \infty$ implies $\sum_{j=-\infty}^{\infty} |a_j b_j| < \infty$.

so that

$$E[\tilde{\gamma}(p)\tilde{\gamma}(q)] = n^{-2} \sum_{s,t} \sum_{i,j,k,\ell} \psi_{s+p-i} \psi_{s-j} \psi_{t+q-k} \psi_{t-\ell} E(w_i w_j w_k w_\ell).$$

Then, evaluate, using the easily verified properties of the w_t series

$$E(w_i w_j w_k w_\ell) = \begin{cases} \eta \sigma_w^4 & \text{if } i = j = k = \ell \\ \sigma_w^4 & \text{if } i = j \neq k = \ell \\ 0 & \text{if } i \neq j, i \neq k \text{ and } i \neq \ell. \end{cases}$$

To apply the rules, we break the sum over the subscripts i, j, k, ℓ into four terms, namely,

$$\sum_{i,j,k,\ell} = \sum_{i=j=k=\ell} + \sum_{i=j \neq k=\ell} + \sum_{i=k \neq j=\ell} + \sum_{i=\ell \neq j=k} = S_1 + S_2 + S_3 + S_4.$$

Now,

$$S_1 = \eta \sigma_w^4 \sum_i \psi_{s+p-i} \psi_{s-i} \psi_{t+q-i} \psi_{t-i} = \eta \sigma_w^4 \sum_i \psi_{i+s-t+p} \psi_{i+s-t} \psi_{i+q} \psi_i,$$

where we have let $i' = t - i$ to get the final form. For the second term,

$$\begin{aligned} S_2 &= \sum_{i=j \neq k=\ell} \psi_{s+p-i} \psi_{s-j} \psi_{t+q-k} \psi_{t-\ell} E(w_i w_j w_k w_\ell) \\ &= \sum_{i \neq k} \psi_{s+p-i} \psi_{s-i} \psi_{t+q-k} \psi_{t-k} E(w_i^2) E(w_k^2). \end{aligned}$$

Then, using the fact that

$$\sum_{i \neq k} = \sum_{i,k} - \sum_{i=k},$$

we have

$$\begin{aligned} S_2 &= \sigma_w^4 \sum_{i,k} \psi_{s+p-i} \psi_{s-i} \psi_{t+q-k} \psi_{t-k} - \sigma_w^4 \sum_i \psi_{s+p-i} \psi_{s-i} \psi_{t+q-i} \psi_{t-i} \\ &= \gamma(p)\gamma(q) - \sigma_w^4 \sum_i \psi_{i+s-t+p} \psi_{i+s-t} \psi_{i+q} \psi_i, \end{aligned}$$

letting $i' = s - i, k' = t - k$ in the first term and $i' = s - i$ in the second term. Repeating the argument for S_3 and S_4 and substituting into the covariance expression yields

$$\begin{aligned} E[\tilde{\gamma}(p)\tilde{\gamma}(q)] &= n^{-2} \sum_{s,t} \left[\gamma(p)\gamma(q) + \gamma(s-t)\gamma(s-t+p-q) \right. \\ &\quad + \gamma(s-t+p)\gamma(s-t-q) \\ &\quad \left. + (\eta - 3)\sigma_w^4 \sum_i \psi_{i+s-t+p} \psi_{i+s-t} \psi_{i+q} \psi_i \right]. \end{aligned}$$

Then, letting $u = s - t$ and subtracting $E[\tilde{\gamma}(p)]E[\tilde{\gamma}(q)] = \gamma(p)\gamma(q)$ from the summation leads to the result (A.51). Summing (A.51) over u and applying dominated convergence leads to (A.52).

The above results for the variances and covariances of the approximating statistics $\tilde{\gamma}(\cdot)$ enable proving the following central limit theorem for the autocovariance functions $\hat{\gamma}(\cdot)$.

Theorem A.6 *If x_t is a stationary linear process of the form (A.43) satisfying the fourth moment condition (A.49), then, for fixed K ,*

$$\begin{pmatrix} \hat{\gamma}(0) \\ \hat{\gamma}(1) \\ \vdots \\ \hat{\gamma}(K) \end{pmatrix} \sim AN \left[\begin{pmatrix} \gamma(0) \\ \gamma(1) \\ \vdots \\ \gamma(K) \end{pmatrix}, n^{-1}V \right],$$

where V is the matrix with elements given by

$$v_{pq} = (\eta - 3)\gamma(p)\gamma(q) + \sum_{u=-\infty}^{\infty} [\gamma(u)\gamma(u-p+q) + \gamma(u+q)\gamma(u-p)]. \quad (\text{A.53})$$

Proof: It suffices to show the result for the approximate autocovariance (A.48) for $\tilde{\gamma}(\cdot)$ by the remark given below it (see also Problem 1.30). First, define the strictly stationary $(2m + K)$ -dependent $(K + 1) \times 1$ vector

$$\mathbf{y}_t^m = \begin{pmatrix} (x_t^m - \mu)^2 \\ (x_{t+1}^m - \mu)(x_t^m - \mu) \\ \vdots \\ (x_{t+K}^m - \mu)(x_t^m - \mu) \end{pmatrix},$$

where

$$x_t^m = \mu + \sum_{j=-m}^m \psi_j w_{t-j}$$

is the usual approximation. The sample mean of the above vector is

$$\bar{\mathbf{y}}_{mn} = n^{-1} \sum_{t=1}^n \mathbf{y}_t^m = \begin{pmatrix} \bar{\gamma}^{mn}(0) \\ \bar{\gamma}^{mn}(1) \\ \vdots \\ \bar{\gamma}^{mn}(K) \end{pmatrix},$$

where

$$\bar{\gamma}^{mn}(h) = n^{-1} \sum_{t=1}^n (x_{t+h}^m - \mu)(x_t^m - \mu)$$

denotes the sample autocovariance of the approximating series. Also,

$$E\mathbf{y}_t^m = \begin{pmatrix} \gamma^m(0) \\ \gamma^m(1) \\ \vdots \\ \gamma^m(K) \end{pmatrix},$$

where $\gamma^m(h)$ is the theoretical covariance function of the series x_t^m . Then, consider the vector

$$\mathbf{y}_{mn} = n^{1/2}[\bar{\mathbf{y}}_{mn} - E(\bar{\mathbf{y}}_{mn})]$$

as an approximation to

$$\mathbf{y}_n = n^{1/2} \left[\begin{pmatrix} \tilde{\gamma}(0) \\ \tilde{\gamma}(1) \\ \vdots \\ \tilde{\gamma}(K) \end{pmatrix} - \begin{pmatrix} \gamma(0) \\ \gamma(1) \\ \vdots \\ \gamma(K) \end{pmatrix} \right],$$

where $E(\bar{\mathbf{y}}_{mn})$ is the same as $E(\mathbf{y}_t^m)$ given above. The elements of the vector approximation $\mathbf{y}_{mn}$ are clearly $n^{1/2}(\tilde{\gamma}^{mn}(h) - \tilde{\gamma}^m(h))$. Note that the elements of $\mathbf{y}_n$ are based on the linear process x_t , whereas the elements of $\mathbf{y}_{mn}$ are based on the m -dependent linear process x_t^m . To obtain a limiting distribution for $\mathbf{y}_n$, we apply the Basic Approximation Theorem, [Theorem A.2](#), using $\mathbf{y}_{mn}$ as our approximation. We now verify (i), (ii), and (iii) of [Theorem A.2](#).

- (i) First, let $\mathbf{c}$ be a $(K + 1) \times 1$ vector of constants, and apply the central limit theorem to the $(2m + K)$ -dependent series $\mathbf{c}'\mathbf{y}_{mn}$ using the Cramér–Wold device ([A.28](#)). We obtain

$$\mathbf{c}'\mathbf{y}_{mn} = n^{1/2}\mathbf{c}'[\bar{\mathbf{y}}_{mn} - E(\bar{\mathbf{y}}_{mn})] \xrightarrow{d} \mathbf{c}'\mathbf{y}_m \sim N(0, \mathbf{c}'V_m\mathbf{c}),$$

as $n \rightarrow \infty$, where V_m is a matrix containing the finite analogs of the elements v_{pq} defined in ([A.53](#)).

- (ii) Note that, since $V_m \rightarrow V$ as $m \rightarrow \infty$, it follows that

$$\mathbf{c}'\mathbf{y}_m \xrightarrow{d} \mathbf{c}'\mathbf{y} \sim N(0, \mathbf{c}'V\mathbf{c}),$$

so, by the Cramér–Wold device, the limiting $(K + 1) \times 1$ multivariate normal variable is $N(\mathbf{0}, V)$.

- (iii) For this condition, we can focus on the element-by-element components of

$$P\{|\mathbf{y}_n - \mathbf{y}_{mn}| > \epsilon\}.$$

For example, using the Tchebycheff inequality, the h -th element of the probability statement can be bounded by

$$\begin{aligned} n\epsilon^{-2} \text{var}(\tilde{\gamma}(h) - \tilde{\gamma}^m(h)) \\ = \epsilon^{-2} \{n \text{var} \tilde{\gamma}(h) + n \text{var} \tilde{\gamma}^m(h) - 2n \text{cov}[\tilde{\gamma}(h), \tilde{\gamma}^m(h)]\}. \end{aligned}$$

Using the results that led to (A.52), we see that the preceding expression approaches

$$(v_{hh} + v_{hh} - 2v_{hh})/\epsilon^2 = 0,$$

as $m, n \rightarrow \infty$.

□

To obtain a result comparable to **Theorem A.6** for the autocorrelation function ACF, we note the following theorem.

Theorem A.7 *If x_t is a stationary linear process of the form (1.29) satisfying the fourth moment condition (A.49), then for fixed K ,*

$$\begin{pmatrix} \widehat{\rho}(1) \\ \vdots \\ \widehat{\rho}(K) \end{pmatrix} \sim AN \left[\begin{pmatrix} \rho(1) \\ \vdots \\ \rho(K) \end{pmatrix}, n^{-1}W \right],$$

where W is the matrix with elements given by

$$\begin{aligned} w_{pq} &= \sum_{u=-\infty}^{\infty} \left[\rho(u+p)\rho(u+q) + \rho(u-p)\rho(u+q) + 2\rho(p)\rho(q)\rho^2(u) \right. \\ &\quad \left. - 2\rho(p)\rho(u)\rho(u+q) - 2\rho(q)\rho(u)\rho(u+p) \right] \\ &= \sum_{u=1}^{\infty} [\rho(u+p) + \rho(u-p) - 2\rho(p)\rho(u)] \\ &\quad \times [\rho(u+q) + \rho(u-q) - 2\rho(q)\rho(u)], \end{aligned} \quad (\text{A.54})$$

where the last form is more convenient.

Proof: To prove the theorem, we use the delta method⁴ for the limiting distribution of a function of the form

$$\mathbf{g}(x_0, x_1, \dots, x_K) = (x_1/x_0, \dots, x_K/x_0)',$$

where $x_h = \widehat{\gamma}(h)$, for $h = 0, 1, \dots, K$. Hence, using the delta method and **Theorem A.6**,

$$\mathbf{g}(\widehat{\gamma}(0), \widehat{\gamma}(1), \dots, \widehat{\gamma}(K)) = (\widehat{\rho}(1), \dots, \widehat{\rho}(K))'$$

is asymptotically normal with mean vector $(\rho(1), \dots, \rho(K))'$ and covariance matrix

$$n^{-1}W = n^{-1}DVD',$$

where V is defined by (A.53) and D is the $(K+1) \times K$ matrix of partial derivatives

⁴ The delta method states that if a k -dimensional vector sequence $\mathbf{x}_n \sim AN(\boldsymbol{\mu}, a_n^2 \Sigma)$, with $a_n \rightarrow 0$, and $\mathbf{g}(\mathbf{x})$ is an $r \times 1$ continuously differentiable vector function of $\mathbf{x}$, then $\mathbf{g}(\mathbf{x}_n) \sim AN(\mathbf{g}(\boldsymbol{\mu}), a_n^2 D \Sigma D')$ where D is the $r \times k$ matrix with elements $d_{ij} = \frac{\partial g_i(\mathbf{x})}{\partial x_j} \big|_{\boldsymbol{\mu}}$.

$$D = \frac{1}{x_0^2} \begin{pmatrix} -x_1 & x_0 & 0 & \dots & 0 \\ -x_2 & 0 & x_0 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ -x_K & 0 & 0 & \dots & x_0 \end{pmatrix}$$

Substituting $\gamma(h)$ for x_h , we note that D can be written as the patterned matrix

$$D = \frac{1}{\gamma(0)} \begin{pmatrix} -\boldsymbol{\rho} & I_K \end{pmatrix},$$

where $\boldsymbol{\rho} = (\rho(1), \rho(2), \dots, \rho(K))'$ is the $K \times 1$ matrix of autocorrelations and I_K is the $K \times K$ identity matrix. Then, it follows from writing the matrix V in the partitioned form

$$V = \begin{pmatrix} v_{00} & \mathbf{v}'_1 \\ \mathbf{v}_1 & V_{22} \end{pmatrix}$$

that

$$W = \gamma^{-2}(0) [v_{00}\boldsymbol{\rho}\boldsymbol{\rho}' - \boldsymbol{\rho}\mathbf{v}'_1 - \mathbf{v}_1\boldsymbol{\rho}' + V_{22}],$$

where $\mathbf{v}_1 = (v_{10}, v_{20}, \dots, v_{K0})'$ and $V_{22} = \{v_{pq}; p, q = 1, \dots, K\}$. Hence,

$$\begin{aligned} w_{pq} &= \gamma^{-2}(0) [v_{pq} - \rho(p)v_{0q} - \rho(q)v_{p0} + \rho(p)\rho(q)v_{00}] \\ &= \sum_{u=-\infty}^{\infty} \left[\rho(u)\rho(u-p+q) + \rho(u-p)\rho(u+q) + 2\rho(p)\rho(q)\rho^2(u) \right. \\ &\quad \left. - 2\rho(p)\rho(u)\rho(u+q) - 2\rho(q)\rho(u)\rho(u-p) \right]. \end{aligned}$$

Interchanging the summations, we get the w_{pq} specified in the statement of the theorem, finishing the proof. $\square$

Specializing the theorem to the case of interest in this chapter, we note that if $\{x_t\}$ is iid with finite fourth moment, then $w_{pq} = 1$ for $p = q$ and is zero otherwise. In this case, for $h = 1, \dots, K$, the $\widehat{\rho}(h)$ are asymptotically independent and jointly normal with

$$\widehat{\rho}(h) \sim AN(0, n^{-1}). \quad (\text{A.55})$$

This justifies the use of (1.36) and the discussion below it as a method for testing whether a series is white noise.

For the cross-correlation, it has been noted that the same kind of approximation holds and we quote the following theorem for the bivariate case, which can be proved using similar arguments (see Brockwell and Davis, 1991, p. 410).

Theorem A.8 *If*

$$x_t = \sum_{j=-\infty}^{\infty} \alpha_j w_{t-j,1}$$

and

$$y_t = \sum_{j=-\infty}^{\infty} \beta_j w_{t-j,2}$$

are two linear processes with absolutely summable coefficients and the two white noise sequences are iid and independent of each other with variances σ_1^2 and σ_2^2 , then for $h \geq 0$,

$$\widehat{\rho}_{xy}(h) \sim AN\left(\rho_{xy}(h), n^{-1} \sum_j \rho_x(j) \rho_y(j)\right) \quad (\text{A.56})$$

and the joint distribution of $(\widehat{\rho}_{xy}(h), \widehat{\rho}_{xy}(k))'$ is asymptotically normal with mean vector zero and

$$\text{cov}\left(\widehat{\rho}_{xy}(h), \widehat{\rho}_{xy}(k)\right) = n^{-1} \sum_j \rho_x(j) \rho_y(j+k-h). \quad (\text{A.57})$$

Again, specializing to the case of interest in this chapter, as long as at least one of the two series is white (iid) noise, we obtain

$$\widehat{\rho}_{xy}(h) \sim AN(0, n^{-1}), \quad (\text{A.58})$$

which justifies [Property 1.2](#).

Appendix B

Time Domain Theory

B.1 Hilbert Spaces and the Projection Theorem

Most of the material on mean square estimation and regression can be embedded in a more general setting involving an inner product space that is also complete (that is, satisfies the Cauchy condition). Two examples of inner products are $E(xy^*)$, where the elements are random variables, and $\sum x_i y_i^*$, where the elements are sequences. These examples include the possibility of complex elements, in which case, $*$ denotes the conjugation. We denote an inner product, in general, by the notation $\langle x, y \rangle$. Now, define an inner product space by its properties, namely,

- (i) $\langle x, y \rangle = \langle y, x \rangle^*$
- (ii) $\langle x + y, z \rangle = \langle x, z \rangle + \langle y, z \rangle$
- (iii) $\langle \alpha x, y \rangle = \alpha \langle x, y \rangle$
- (iv) $\langle x, x \rangle = \|x\|^2 \geq 0$
- (v) $\langle x, x \rangle = 0$ iff $x = 0$.

We introduced the notation $\|\cdot\|$ for the norm or distance in property (iv). The norm satisfies the triangle inequality

$$\|x + y\| \leq \|x\| + \|y\| \quad (\text{B.1})$$

and the Cauchy–Schwarz inequality

$$|\langle x, y \rangle|^2 \leq \|x\|^2 \|y\|^2, \quad (\text{B.2})$$

which we have seen before for random variables in (A.35). Now, a Hilbert space, $\mathcal{H}$, is defined as an inner product space with the Cauchy property. In other words, $\mathcal{H}$ is a complete inner product space. This means that every Cauchy sequence converges in norm; that is, $x_n \rightarrow x \in \mathcal{H}$ if and only if $\|x_n - x_m\| \rightarrow 0$ as $m, n \rightarrow \infty$. This is just the L^2 completeness **Theorem A.1** for random variables.

For a broad overview of Hilbert space techniques that are useful in statistical inference and in probability, see Small and McLeish (1994). Also, Brockwell and Davis (1991, Chapter 2) is a nice summary of Hilbert space techniques that are useful

in time series analysis. In our discussions, we mainly use the *projection theorem* (**Theorem B.1**) and the associated orthogonality principle as a means for solving various kinds of linear estimation problems.

Theorem B.1 (Projection Theorem) *Let $\mathcal{M}$ be a closed subspace of the Hilbert space $\mathcal{H}$ and let y be an element in $\mathcal{H}$. Then, y can be uniquely represented as*

$$y = \hat{y} + z, \quad (\text{B.3})$$

where $\hat{y}$ belongs to $\mathcal{M}$ and z is orthogonal to $\mathcal{M}$; that is, $\langle z, w \rangle = 0$ for all w in $\mathcal{M}$. Furthermore, the point $\hat{y}$ is closest to y in the sense that, for any w in $\mathcal{M}$, $\|y - w\| \geq \|y - \hat{y}\|$, where equality holds if and only if $w = \hat{y}$.

We note that (B.3) and the statement following it yield the *orthogonality property*

$$\langle y - \hat{y}, w \rangle = 0 \quad (\text{B.4})$$

for any w belonging to $\mathcal{M}$, which can sometimes be used easily to find an expression for the projection. The norm of the error can be written as

$$\begin{aligned} \|y - \hat{y}\|^2 &= \langle y - \hat{y}, y - \hat{y} \rangle \\ &= \langle y - \hat{y}, y \rangle - \langle y - \hat{y}, \hat{y} \rangle \\ &= \langle y - \hat{y}, y \rangle \end{aligned} \quad (\text{B.5})$$

because of orthogonality.

Using the notation of **Theorem B.1**, we call the mapping $P_{\mathcal{M}}y = \hat{y}$, for $y \in \mathcal{H}$, the projection mapping of $\mathcal{H}$ onto $\mathcal{M}$. In addition, the closed span of a finite set $\{x_1, \dots, x_n\}$ of elements in a Hilbert space, $\mathcal{H}$, is defined to be the set of all linear combinations $w = a_1x_1 + \dots + a_nx_n$, where $a_1, \dots, a_n$ are scalars. This subspace of $\mathcal{H}$ is denoted by $\mathcal{M} = \overline{\text{sp}}\{x_1, \dots, x_n\}$. By the projection theorem, the projection of $y \in \mathcal{H}$ onto $\mathcal{M}$ is unique and given by

$$P_{\mathcal{M}}y = a_1x_1 + \dots + a_nx_n,$$

where $\{a_1, \dots, a_n\}$ are found using the orthogonality principle

$$\langle y - P_{\mathcal{M}}y, x_j \rangle = 0 \quad j = 1, \dots, n.$$

Evidently, $\{a_1, \dots, a_n\}$ can be obtained by solving

$$\sum_{i=1}^n a_i \langle x_i, x_j \rangle = \langle y, x_j \rangle \quad j = 1, \dots, n. \quad (\text{B.6})$$

When the elements of $\mathcal{H}$ are vectors, this problem is the linear regression problem.

Example B.1 Linear Regression Analysis

For the regression model introduced in §2.2, we want to find the regression coefficients β_i that minimize the residual sum of squares. Consider the vectors $\mathbf{y} = (y_1, \dots, y_n)'$ and $\mathbf{z}_i = (z_{1i}, \dots, z_{ni})'$, for $i = 1, \dots, q$ and the inner product

$$\langle \mathbf{z}_i, \mathbf{y} \rangle = \sum_{t=1}^n z_{ti} y_t = \mathbf{z}_i' \mathbf{y}.$$

We solve the problem of finding a projection of the observed $\mathbf{y}$ on the linear space spanned by $\beta_1 \mathbf{z}_1 + \dots + \beta_q \mathbf{z}_q$, that is, linear combinations of the $\mathbf{z}_i$. The orthogonality principle gives

$$\left\langle \mathbf{y} - \sum_{i=1}^q \beta_i \mathbf{z}_i, \mathbf{z}_j \right\rangle = 0$$

for $j = 1, \dots, q$. Writing the orthogonality condition, as in (B.6), in vector form gives

$$\mathbf{y}' \mathbf{z}_j = \sum_{i=1}^q \beta_i \mathbf{z}_i' \mathbf{z}_j \quad j = 1, \dots, q, \quad (\text{B.7})$$

which can be written in the usual matrix form by letting $\mathbf{Z} = (\mathbf{z}_1, \dots, \mathbf{z}_q)$, which is assumed to be full rank. That is, (B.7) can be written as

$$\mathbf{y}' \mathbf{Z} = \boldsymbol{\beta}' (\mathbf{Z}' \mathbf{Z}), \quad (\text{B.8})$$

where $\boldsymbol{\beta} = (\beta_1, \dots, \beta_q)'$. Transposing both sides of (B.8) provides the solution for the coefficients,

$$\widehat{\boldsymbol{\beta}} = (\mathbf{Z}' \mathbf{Z})^{-1} \mathbf{Z}' \mathbf{y}.$$

The mean-square error in this case would be

$$\left\| \mathbf{y} - \sum_{i=1}^q \widehat{\beta}_i \mathbf{z}_i \right\|^2 = \left\langle \mathbf{y} - \sum_{i=1}^q \widehat{\beta}_i \mathbf{z}_i, \mathbf{y} \right\rangle = \langle \mathbf{y}, \mathbf{y} \rangle - \sum_{i=1}^q \widehat{\beta}_i \langle \mathbf{z}_i, \mathbf{y} \rangle = \mathbf{y}' \mathbf{y} - \widehat{\boldsymbol{\beta}}' \mathbf{Z}' \mathbf{y},$$

which is in agreement with §2.1.

The extra generality in the above approach hardly seems necessary in the finite dimensional case, where differentiation works perfectly well. It is convenient, however, in many cases to regard the elements of $\mathcal{H}$ as infinite dimensional, so that the orthogonality principle becomes of use. For example, the projection of the process $\{x_t; t = 0 \pm 1, \pm 2, \dots\}$ on the linear manifold spanned by all filtered convolutions of the form

$$\widehat{x}_t = \sum_{k=-\infty}^{\infty} a_k x_{t-k}$$

would be in this form.

There are some useful results, which we state without proof, pertaining to projection mappings.

Theorem B.2 *Under established notation and conditions:*

- (i) $P_{\mathcal{M}}(ax + by) = aP_{\mathcal{M}}x + bP_{\mathcal{M}}y$, for $x, y \in \mathcal{H}$, where a and b are scalars.
- (ii) If $\|y_n - y\| \rightarrow 0$, then $P_{\mathcal{M}}y_n \rightarrow P_{\mathcal{M}}y$, as $n \rightarrow \infty$.
- (iii) $w \in \mathcal{M}$ if and only if $P_{\mathcal{M}}w = w$. Consequently, a projection mapping can be characterized by the property that $P_{\mathcal{M}}^2 = P_{\mathcal{M}}$, in the sense that, for any $y \in \mathcal{H}$, $P_{\mathcal{M}}(P_{\mathcal{M}}y) = P_{\mathcal{M}}y$.
- (iv) Let $\mathcal{M}_1$ and $\mathcal{M}_2$ be closed subspaces of $\mathcal{H}$. Then, $\mathcal{M}_1 \subseteq \mathcal{M}_2$ if and only if $P_{\mathcal{M}_1}(P_{\mathcal{M}_2}y) = P_{\mathcal{M}_1}y$ for all $y \in \mathcal{H}$.
- (v) Let $\mathcal{M}$ be a closed subspace of $\mathcal{H}$ and let $\mathcal{M}_{\perp}$ denote the orthogonal complement of $\mathcal{M}$. Then, $\mathcal{M}_{\perp}$ is also a closed subspace of $\mathcal{H}$, and for any $y \in \mathcal{H}$, $y = P_{\mathcal{M}}y + P_{\mathcal{M}_{\perp}}y$.

Part (iii) of **Theorem B.2** leads to the well-known result, often used in linear models, that a square matrix M is a projection matrix if and only if it is symmetric and idempotent (that is, $M^2 = M$). For example, using the notation of **Example B.1** for linear regression, the projection of $\mathbf{y}$ onto $\text{sp}\{\mathbf{z}_1, \dots, \mathbf{z}_q\}$, the space generated by the columns of Z , is $P_Z(\mathbf{y}) = Z\hat{\boldsymbol{\beta}} = Z(Z'Z)^{-1}Z'\mathbf{y}$. The matrix $M = Z(Z'Z)^{-1}Z'$ is an $n \times n$, symmetric and idempotent matrix of rank q (which is the dimension of the space that M projects $\mathbf{y}$ onto). Parts (iv) and (v) of **Theorem B.2** are useful for establishing recursive solutions for estimation and prediction.

By imposing extra structure, conditional expectation can be defined as a projection mapping for random variables in L^2 with the equivalence relation that, for $x, y \in L^2$, $x = y$ if $\Pr(x = y) = 1$. In particular, for $y \in L^2$, if $\mathcal{M}$ is a closed subspace of L^2 containing 1, the conditional expectation of y given $\mathcal{M}$ is defined to be the projection of y onto $\mathcal{M}$, namely, $E_{\mathcal{M}}y = P_{\mathcal{M}}y$. This means that conditional expectation, $E_{\mathcal{M}}$, must satisfy the orthogonality principle of the Projection Theorem and that the results of **Theorem B.2** remain valid (the most widely used tool in this case is item (iv) of the theorem). If we let $\mathcal{M}(x)$ denote the closed subspace of all random variables in L^2 that can be written as a (measurable) function of x , then we may define, for $x, y \in L^2$, the conditional expectation of y given x as $E(y|x) = E_{\mathcal{M}(x)}y$. This idea may be generalized in an obvious way to define the conditional expectation of y given $\mathbf{x} = (x_1, \dots, x_n)$; that is $E(y|\mathbf{x}) = E_{\mathcal{M}(\mathbf{x})}y$. Of particular interest to us is the following result which states that, in the Gaussian case, conditional expectation and linear prediction are equivalent.

Theorem B.3 *Under established notation and conditions, if $(y, x_1, \dots, x_n)$ is multivariate normal, then*

$$E(y \mid x_1, \dots, x_n) = P_{\text{sp}\{1, x_1, \dots, x_n\}}y.$$

Proof: First, by the projection theorem, the conditional expectation of y given $\mathbf{x} = \{x_1, \dots, x_n\}$ is the unique element $E_{\mathcal{M}(\mathbf{x})}y$ that satisfies the orthogonality principle,

$$E\{(y - E_{\mathcal{M}(\mathbf{x})}y) w\} = 0 \quad \text{for all } w \in \mathcal{M}(\mathbf{x}).$$

We will show that $\hat{y} = P_{\text{sp}\{1, x_1, \dots, x_n\}}y$ is that element. In fact, by the projection theorem, $\hat{y}$ satisfies

$$\langle y - \hat{y}, x_i \rangle = 0 \quad \text{for } i = 0, 1, \dots, n,$$

where we have set $x_0 = 1$. But $\langle y - \hat{y}, x_i \rangle = \text{cov}(y - \hat{y}, x_i) = 0$, implying that $y - \hat{y}$ and $(x_1, \dots, x_n)$ are independent because the vector $(y - \hat{y}, x_1, \dots, x_n)'$ is multivariate normal. Thus, if $w \in \mathcal{M}(\mathbf{x})$, then w and $y - \hat{y}$ are independent and, hence, $\langle y - \hat{y}, w \rangle = E\{(y - \hat{y})w\} = E(y - \hat{y})E(w) = 0$, recalling that $0 = \langle y - \hat{y}, 1 \rangle = E(y - \hat{y})$. $\square$

In the Gaussian case, conditional expectation has an explicit form. Let $\mathbf{y} = (y_1, \dots, y_m)'$, $\mathbf{x} = (x_1, \dots, x_n)'$, and suppose the $(m + n) \times 1$ vector $(\mathbf{y}', \mathbf{x}')'$ is normal:

$$\begin{pmatrix} \mathbf{y} \\ \mathbf{x} \end{pmatrix} \sim N \left[\begin{pmatrix} \boldsymbol{\mu}_y \\ \boldsymbol{\mu}_x \end{pmatrix}, \begin{pmatrix} \Sigma_{yy} & \Sigma_{yx} \\ \Sigma_{xy} & \Sigma_{xx} \end{pmatrix} \right],$$

then $\mathbf{y}|\mathbf{x}$ is normal with

$$\boldsymbol{\mu}_{y|\mathbf{x}} = \boldsymbol{\mu}_y + \Sigma_{yx} \Sigma_{xx}^{-1} (\mathbf{x} - \boldsymbol{\mu}_x) \quad (\text{B.9})$$

$$\Sigma_{y|\mathbf{x}} = \Sigma_{yy} - \Sigma_{yx} \Sigma_{xx}^{-1} \Sigma_{xy}, \quad (\text{B.10})$$

where Σ_{xx} is assumed to be nonsingular.

B.2 Causal Conditions for ARMA Models

In this section, we prove **Property 3.1** of §3.2 pertaining to the causality of ARMA models. The proof of **Property 3.2**, which pertains to invertibility of ARMA models, is similar.

Proof of Property 3.1. Suppose first that the roots of $\phi(z)$, say, $z_1, \dots, z_p$, lie outside the unit circle. We write the roots in the following order, $1 < |z_1| \leq |z_2| \leq \dots \leq |z_p|$, noting that $z_1, \dots, z_p$ are not necessarily unique, and put $|z_1| = 1 + \epsilon$, for some $\epsilon > 0$. Thus, $\phi(z) \neq 0$ as long as $|z| < |z_1| = 1 + \epsilon$ and, hence, $\phi^{-1}(z)$ exists and has a power series expansion,

$$\frac{1}{\phi(z)} = \sum_{j=0}^{\infty} a_j z^j, \quad |z| < 1 + \epsilon.$$

Now, choose a value δ such that $0 < \delta < \epsilon$, and set $z = 1 + \delta$, which is inside the radius of convergence. It then follows that

$$\phi^{-1}(1 + \delta) = \sum_{j=0}^{\infty} a_j (1 + \delta)^j < \infty. \quad (\text{B.11})$$

Thus, we can bound each of the terms in the sum in (B.11) by a constant, say, $|a_j (1 + \delta)^j| < c$, for $c > 0$. In turn, $|a_j| < c(1 + \delta)^{-j}$, from which it follows that

$$\sum_{j=0}^{\infty} |a_j| < \infty. \quad (\text{B.12})$$

Hence, $\phi^{-1}(B)$ exists and we may apply it to both sides of the ARMA model, $\phi(B)x_t = \theta(B)w_t$, to obtain

$$x_t = \phi^{-1}(B)\phi(B)x_t = \phi^{-1}(B)\theta(B)w_t.$$

Thus, putting $\psi(B) = \phi^{-1}(B)\theta(B)$, we have

$$x_t = \psi(B)w_t = \sum_{j=0}^{\infty} \psi_j w_{t-j},$$

where the ψ -weights, which are absolutely summable, can be evaluated by $\psi(z) = \phi^{-1}(z)\theta(z)$, for $|z| \leq 1$.

Now, suppose x_t is a causal process; that is, it has the representation

$$x_t = \sum_{j=0}^{\infty} \psi_j w_{t-j}, \quad \sum_{j=0}^{\infty} |\psi_j| < \infty.$$

In this case, we write

$$x_t = \psi(B)w_t,$$

and premultiplying by $\phi(B)$ yields

$$\phi(B)x_t = \phi(B)\psi(B)w_t. \quad (\text{B.13})$$

In addition to (B.13), the model is ARMA, and can be written as

$$\phi(B)x_t = \theta(B)w_t. \quad (\text{B.14})$$

From (B.13) and (B.14), we see that

$$\phi(B)\psi(B)w_t = \theta(B)w_t. \quad (\text{B.15})$$

Now, let

$$a(z) = \phi(z)\psi(z) = \sum_{j=0}^{\infty} a_j z^j \quad |z| \leq 1$$

and, hence, we can write (B.15) as

$$\sum_{j=0}^{\infty} a_j w_{t-j} = \sum_{j=0}^q \theta_j w_{t-j}. \quad (\text{B.16})$$

Next, multiply both sides of (B.16) by w_{t-h} , for $h = 0, 1, 2, \dots$, and take expectation. In doing this, we obtain

$$\begin{aligned} a_h &= \theta_h, & h &= 0, 1, \dots, q \\ a_h &= 0, & h &> q. \end{aligned} \quad (\text{B.17})$$

From (B.17), we conclude that

$$\phi(z)\psi(z) = a(z) = \theta(z), \quad |z| \leq 1. \quad (\text{B.18})$$

If there is a complex number in the unit circle, say z_0 , for which $\phi(z_0) = 0$, then by (B.18), $\theta(z_0) = 0$. But, if there is such a z_0 , then $\phi(z)$ and $\theta(z)$ have a common factor which is not allowed. Thus, we may write $\psi(z) = \theta(z)/\phi(z)$. In addition, by hypothesis, we have that $|\psi(z)| < \infty$ for $|z| \leq 1$, and hence

$$|\psi(z)| = \left| \frac{\theta(z)}{\phi(z)} \right| < \infty, \quad \text{for } |z| \leq 1. \quad (\text{B.19})$$

Finally, (B.19) implies $\phi(z) \neq 0$ for $|z| \leq 1$; that is, the roots of $\phi(z)$ lie outside the unit circle. $\square$

B.3 Large Sample Distribution of the AR(p) Conditional Least Squares Estimators

In §3.6 we discussed the conditional least squares procedure for estimating the parameters $\phi_1, \phi_2, \dots, \phi_p$ and σ_w^2 in the AR(p) model

$$x_t = \sum_{k=1}^p \phi_k x_{t-k} + w_t,$$

where we assume $\mu = 0$, for convenience. Write the model as

$$x_t = \boldsymbol{\phi}' \mathbf{x}_{t-1} + w_t, \quad (\text{B.20})$$

where $\mathbf{x}_{t-1} = (x_{t-1}, x_{t-2}, \dots, x_{t-p})'$ is a $p \times 1$ vector of lagged values, and $\boldsymbol{\phi} = (\phi_1, \phi_2, \dots, \phi_p)'$ is the $p \times 1$ vector of regression coefficients. Assuming observations are available at $x_1, \dots, x_n$, the conditional least squares procedure is to minimize

$$S_c(\boldsymbol{\phi}) = \sum_{t=p+1}^n (x_t - \boldsymbol{\phi}' \mathbf{x}_{t-1})^2$$

with respect to $\boldsymbol{\phi}$. The solution is

$$\hat{\boldsymbol{\phi}} = \left(\sum_{t=p+1}^n \mathbf{x}_{t-1} \mathbf{x}_{t-1}' \right)^{-1} \sum_{t=p+1}^n \mathbf{x}_{t-1} x_t \quad (\text{B.21})$$

for the regression vector $\boldsymbol{\phi}$; the conditional least squares estimate of σ_w^2 is

$$\hat{\sigma}_w^2 = \frac{1}{n-p} \sum_{t=p+1}^n (x_t - \hat{\boldsymbol{\phi}}' \mathbf{x}_{t-1})^2. \quad (\text{B.22})$$

As pointed out following (3.115), Yule–Walker estimators and least squares estimators are approximately the same in that the estimators differ only by inclusion or exclusion of terms involving the endpoints of the data. Hence, it is easy to show the asymptotic equivalence of the two estimators; this is why, for $AR(p)$ models, (3.103) and (3.131), are equivalent. Details on the asymptotic equivalence can be found in Brockwell and Davis (1991, Chapter 8).

Here, we use the same approach as in Appendix A, replacing the lower limits of the sums in (B.21) and (B.22) by one and noting the asymptotic equivalence of the estimators

$$\tilde{\boldsymbol{\phi}} = \left(\sum_{t=1}^n \mathbf{x}_{t-1} \mathbf{x}_{t-1}' \right)^{-1} \sum_{t=1}^n \mathbf{x}_{t-1} x_t \quad (\text{B.23})$$

and

$$\tilde{\sigma}_w^2 = \frac{1}{n} \sum_{t=1}^n (x_t - \tilde{\boldsymbol{\phi}}' \mathbf{x}_{t-1})^2 \quad (\text{B.24})$$

to those two estimators. In (B.23) and (B.24), we are acting as if we are able to observe $x_{1-p}, \dots, x_0$ in addition to $x_1, \dots, x_n$. The asymptotic equivalence is then seen by arguing that for n sufficiently large, it makes no difference whether or not we observe $x_{1-p}, \dots, x_0$. In the case of (B.23) and (B.24), we obtain the following theorem.

Theorem B.4 *Let x_t be a causal $AR(p)$ series with white (iid) noise w_t satisfying $E(w_t^4) = \eta \sigma_w^4$. Then,*

$$\tilde{\boldsymbol{\phi}} \sim \text{AN} \left(\boldsymbol{\phi}, n^{-1} \sigma_w^2 \Gamma_p^{-1} \right), \quad (\text{B.25})$$

where $\Gamma_p = \{\gamma(i-j)\}_{i,j=1}^p$ is the $p \times p$ autocovariance matrix of the vector $\mathbf{x}_{t-1}$. We also have, as $n \rightarrow \infty$,

$$n^{-1} \sum_{t=1}^n \mathbf{x}_{t-1} \mathbf{x}_{t-1}' \xrightarrow{P} \Gamma_p \quad \text{and} \quad \tilde{\sigma}_w^2 \xrightarrow{P} \sigma_w^2. \quad (\text{B.26})$$

Proof: First, (B.26) follows from the fact that $E(\mathbf{x}_{t-1} \mathbf{x}_{t-1}') = \Gamma_p$, recalling that from Theorem A.6, second-order sample moments converge in probability to their population moments for linear processes in which w_t has a finite fourth moment. To show (B.25), we can write

$$\begin{aligned} \tilde{\boldsymbol{\phi}} &= \left(\sum_{t=1}^n \mathbf{x}_{t-1} \mathbf{x}_{t-1}' \right)^{-1} \sum_{t=1}^n \mathbf{x}_{t-1} (\mathbf{x}_{t-1}' \boldsymbol{\phi} + w_t) \\ &= \boldsymbol{\phi} + \left(\sum_{t=1}^n \mathbf{x}_{t-1} \mathbf{x}_{t-1}' \right)^{-1} \sum_{t=1}^n \mathbf{x}_{t-1} w_t, \end{aligned}$$

so that

$$n^{1/2}(\tilde{\boldsymbol{\phi}} - \boldsymbol{\phi}) = \left(n^{-1} \sum_{t=1}^n \mathbf{x}_{t-1} \mathbf{x}_{t-1}' \right)^{-1} n^{-1/2} \sum_{t=1}^n \mathbf{x}_{t-1} w_t$$

$$= \left(n^{-1} \sum_{t=1}^n \mathbf{x}_{t-1} \mathbf{x}'_{t-1} \right)^{-1} n^{-1/2} \sum_{t=1}^n \mathbf{u}_t,$$

where $\mathbf{u}_t = \mathbf{x}_{t-1} w_t$. We use the fact that w_t and $\mathbf{x}_{t-1}$ are independent to write $E\mathbf{u}_t = E(\mathbf{x}_{t-1})E(w_t) = \mathbf{0}$, because the errors have zero means. Also,

$$E\mathbf{u}_t \mathbf{u}'_t = E\mathbf{x}_{t-1} w_t w_t \mathbf{x}'_{t-1} = E\mathbf{x}_{t-1} \mathbf{x}'_{t-1} E w_t^2 = \sigma_w^2 \Gamma_p.$$

In addition, we have, for $h > 0$,

$$E\mathbf{u}_{t+h} \mathbf{u}'_t = E\mathbf{x}_{t+h-1} w_{t+h} w_t \mathbf{x}'_{t-1} = E\mathbf{x}_{t+h-1} w_t \mathbf{x}'_{t-1} E w_{t+h} = 0.$$

A similar computation works for $h < 0$.

Next, consider the mean square convergent approximation

$$x_t^m = \sum_{j=0}^m \psi_j w_{t-j}$$

for x_t , and define the $(m+p)$ -dependent process $\mathbf{u}_t^m = w_t (x_{t-1}^m, x_{t-2}^m, \dots, x_{t-p}^m)'$. Note that we need only look at a central limit theorem for the sum

$$y_{nm} = n^{-1/2} \sum_{t=1}^n \boldsymbol{\lambda}' \mathbf{u}_t^m,$$

for arbitrary vectors $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_p)'$, where y_{nm} is used as an approximation to

$$S_n = n^{-1/2} \sum_{t=1}^n \boldsymbol{\lambda}' \mathbf{u}_t.$$

First, apply the m -dependent central limit theorem to y_{nm} as $n \rightarrow \infty$ for fixed m to establish (i) of [Theorem A.2](#). This result shows $y_{nm} \xrightarrow{d} y_m$, where y_m is asymptotically normal with covariance $\boldsymbol{\lambda}' \Gamma_p^{(m)} \boldsymbol{\lambda}$, where $\Gamma_p^{(m)}$ is the covariance matrix of $\mathbf{u}_t^m$. Then, we have $\Gamma_p^{(m)} \rightarrow \Gamma_p$, so that y_m converges in distribution to a normal random variable with mean zero and variance $\boldsymbol{\lambda}' \Gamma_p \boldsymbol{\lambda}$ and we have verified part (ii) of [Theorem A.2](#). We verify part (iii) of [Theorem A.2](#) by noting that

$$E[(S_n - y_{nm})^2] = n^{-1} \sum_{t=1}^n \boldsymbol{\lambda}' E[(\mathbf{u}_t - \mathbf{u}_t^m)(\mathbf{u}_t - \mathbf{u}_t^m)'] \boldsymbol{\lambda}$$

clearly converges to zero as $n, m \rightarrow \infty$ because

$$x_t - x_t^m = \sum_{j=m+1}^{\infty} \psi_j w_{t-j}$$

form the components of $\mathbf{u}_t - \mathbf{u}_t^m$.

Now, the form for $\sqrt{n}(\tilde{\boldsymbol{\phi}} - \boldsymbol{\phi})$ contains the premultiplying matrix

$$\left(n^{-1} \sum_{t=1}^n \mathbf{x}_{t-1} \mathbf{x}'_{t-1} \right)^{-1} \xrightarrow{P} \Gamma_p^{-1},$$

because (A.22) can be applied to the function that defines the inverse of the matrix. Then, applying (A.30), shows that

$$n^{1/2} (\tilde{\boldsymbol{\phi}} - \boldsymbol{\phi}) \xrightarrow{d} N(0, \sigma_w^2 \Gamma_p^{-1} \Gamma_p \Gamma_p^{-1}),$$

so we may regard it as being multivariate normal with mean zero and covariance matrix $\sigma_w^2 \Gamma_p^{-1}$.

To investigate $\tilde{\sigma}_w^2$, note

$$\begin{aligned} \tilde{\sigma}_w^2 &= n^{-1} \sum_{t=1}^n (x_t - \tilde{\boldsymbol{\phi}}' \mathbf{x}_{t-1})^2 \\ &= n^{-1} \sum_{t=1}^n x_t^2 - n^{-1} \sum_{t=1}^n \mathbf{x}'_{t-1} x_t \left(n^{-1} \sum_{t=1}^n \mathbf{x}_{t-1} \mathbf{x}'_{t-1} \right)^{-1} n^{-1} \sum_{t=1}^n \mathbf{x}_{t-1} x_t \\ &\xrightarrow{P} \gamma(0) - \boldsymbol{\gamma}'_p \Gamma_p^{-1} \boldsymbol{\gamma}_p \\ &= \sigma_w^2, \end{aligned}$$

and we have that the sample estimator converges in probability to σ_w^2 , which is written in the form of (3.66). $\square$

The arguments above imply that, for sufficiently large n , we may consider the estimator $\hat{\boldsymbol{\phi}}$ in (B.21) as being approximately multivariate normal with mean $\boldsymbol{\phi}$ and variance–covariance matrix $\sigma_w^2 \Gamma_p^{-1}/n$. Inferences about the parameter $\boldsymbol{\phi}$ are obtained by replacing the σ_w^2 and Γ_p by their estimates given by (B.22) and

$$\hat{\Gamma}_p = n^{-1} \sum_{t=p+1}^n \mathbf{x}_{t-1} \mathbf{x}'_{t-1},$$

respectively. In the case of a nonzero mean, the data x_t are replaced by $x_t - \bar{x}$ in the estimates and the results of Theorem A.2 remain valid.

B.4 The Wold Decomposition

The ARMA approach to modeling time series is generally implied by the assumption that the dependence between adjacent values in time is best explained in terms of a regression of the current values on the past values. This assumption is partially justified, in theory, by the Wold decomposition.

In this section we assume that $\{x_t; t = 0, \pm 1, \pm 2, \dots\}$ is a stationary, mean-zero process. Using the notation of §B.1, we define

$$\mathcal{M}_n^x = \overline{\text{sp}}\{x_t, -\infty < t \leq n\}, \quad \text{with} \quad \mathcal{M}_{-\infty}^x = \bigcap_{n=-\infty}^{\infty} \mathcal{M}_n^x,$$

and

$$\sigma_x^2 = E \left(x_{n+1} - P_{\mathcal{M}_n^x} x_{n+1} \right)^2.$$

We say that x_t is a deterministic process if and only if $\sigma_x^2 = 0$. That is, a deterministic process is one in which its future is perfectly predictable from its past; a simple example is the process given in (4.1). We are now ready to present the decomposition.

Theorem B.5 (The Wold Decomposition) *Under the conditions and notation of this section, if $\sigma_x^2 > 0$, then x_t can be expressed as*

$$x_t = \sum_{j=0}^{\infty} \psi_j w_{t-j} + v_t$$

where

- (i) $\sum_{j=0}^{\infty} \psi_j^2 < \infty$ ($\psi_0 = 1$)
- (ii) $\{w_t\}$ is white noise with variance σ_w^2
- (iii) $w_t \in \mathcal{M}_t^x$
- (iv) $\text{cov}(w_s, v_t) = 0$ for all $s, t = 0, \pm 1, \pm 2, \dots$
- (v) $v_t \in \mathcal{M}_{-\infty}^x$
- (vi) $\{v_t\}$ is deterministic.

The proof of the decomposition follows from the theory of §B.1 by defining the unique sequences:

$$\begin{aligned} w_t &= x_t - P_{\mathcal{M}_{t-1}^x} x_t, \\ \psi_j &= \sigma_w^{-2} \langle x_t, w_{t-j} \rangle = \sigma_w^{-2} E(x_t w_{t-j}), \\ v_t &= x_t - \sum_{j=0}^{\infty} \psi_j w_{t-j}. \end{aligned}$$

Although every stationary process can be represented by the Wold decomposition, it does not mean that the decomposition is the best way to describe the process. In addition, there may be some dependence structure among the $\{w_t\}$; we are only guaranteed that the sequence is an uncorrelated sequence. The theorem, in its generality, falls short of our needs because we would prefer the noise process, $\{w_t\}$, to be white independent noise. But, the decomposition does give us the confidence that we will not be completely off the mark by fitting ARMA models to time series data.

Appendix C

Spectral Domain Theory

C.1 Spectral Representation Theorem

In this section, we present a spectral representation for the process x_t itself, which allows us to think of a stationary process as a random sum of sines and cosines as described in (4.3). In addition, we present results that justify representing the autocovariance function $\gamma_x(h)$ of the weakly stationary process x_t in terms of a non-negative spectral density function. The spectral density function essentially measures the variance or power in a particular kind of periodic oscillation in the function. We denote this spectral density of variance function by $f(\omega)$, where the variance is measured as a function of the frequency of oscillation ω , measured in cycles per unit time.

First, we consider developing a representation for the autocovariance function of a stationary, possibly complex, series x_t with zero mean and autocovariance function $\gamma_x(h) = E(x_{t+h}x_t^*)$. We prove the representation for arbitrary non-negative definite functions $\gamma(h)$ and then simply note the autocovariance function is non-negative definite, because, for any set of complex constants, $a_t, t = 0 \pm 1, \pm 2, \dots$, we may write, for any finite subset,

$$E \left| \sum_{s=1}^n a_s^* x_s \right|^2 = \sum_{s=1}^n \sum_{t=1}^n a_s^* \gamma(s-t) a_t \geq 0.$$

The representation is stated in terms of non-negative definite functions and a spectral distribution function $F(\omega)$ that is monotone nondecreasing, and continuous from the right, taking the values $F(-1/2) = 0$ and $F(1/2) = \sigma^2 = \gamma_x(0)$ at $\omega = -1/2$ and $1/2$, respectively.

Theorem C.1 *A function $\gamma(h)$, for $h = 0, \pm 1, \pm 2, \dots$ is non-negative definite if and only if it can be expressed as*

$$\gamma(h) = \int_{-1/2}^{1/2} \exp\{2\pi i \omega h\} dF(\omega), \quad (\text{C.1})$$

where $F(\cdot)$ is nondecreasing. The function $F(\cdot)$ is right continuous, bounded in $[-1/2, 1/2]$, and uniquely determined by the conditions $F(-1/2) = 0$, $F(1/2) = \gamma(0)$.

Proof: If $\gamma(h)$ has the representation (C.1), then

$$\begin{aligned} \sum_{s=1}^n \sum_{t=1}^n a_s^* \gamma(s-t) a_t &= \int_{-1/2}^{1/2} \sum_{s=1}^n \sum_{t=1}^n a_s^* a_t e^{2\pi i \omega (s-t)} dF(\omega) \\ &= \int_{-1/2}^{1/2} \left| \sum_{t=1}^n a_t e^{-2\pi i \omega t} \right|^2 dF(\omega) \geq 0 \end{aligned}$$

and $\gamma(h)$ is non-negative definite.

Conversely, suppose $\gamma(h)$ is a non-negative definite function. Define the non-negative function

$$\begin{aligned} f_n(\omega) &= n^{-1} \sum_{s=1}^n \sum_{t=1}^n e^{-2\pi i \omega s} \gamma(s-t) e^{2\pi i \omega t} \\ &= n^{-1} \sum_{h=-(n-1)}^{(n-1)} (n - |h|) e^{-2\pi i \omega h} \gamma(h) \geq 0 \end{aligned} \tag{C.2}$$

Now, let $F_n(\omega)$ be the distribution function corresponding to $f_n(\omega)I_{(-1/2, 1/2]}$, where $I_{(\cdot)}$ denotes the indicator function of the interval in the subscript. Note that $F_n(\omega) = 0$, $\omega \leq -1/2$ and $F_n(\omega) = F_n(1/2)$ for $\omega \geq 1/2$. Then,

$$\begin{aligned} \int_{-1/2}^{1/2} e^{2\pi i \omega h} dF_n(\omega) &= \int_{-1/2}^{1/2} e^{2\pi i \omega h} f_n(\omega) d\omega \\ &= \begin{cases} (1 - |h|/n) \gamma(h), & |h| < n \\ 0, & \text{elsewhere.} \end{cases} \end{aligned}$$

We also have

$$\begin{aligned} F_n(1/2) &= \int_{-1/2}^{1/2} f_n(\omega) d\omega \\ &= \int_{-1/2}^{1/2} \sum_{|h| < n} (1 - |h|/n) \gamma(h) e^{-2\pi i \omega h} d\omega = \gamma(0). \end{aligned}$$

Now, by Helly's first convergence theorem (Bhat, 1985, p. 157), there exists a subsequence F_{n_k} converging to F , and by the Helly-Bray Lemma (see Bhat, p. 157), this implies

$$\int_{-1/2}^{1/2} e^{2\pi i \omega h} dF_{n_k}(\omega) \rightarrow \int_{-1/2}^{1/2} e^{2\pi i \omega h} dF(\omega)$$

and, from the right-hand side of the earlier equation,

$$(1 - |h|/n_k) \gamma(h) \rightarrow \gamma(h)$$

as $n_k \rightarrow \infty$, and the required result follows. $\square$

Next, present the version of the Spectral Representation Theorem in terms of a mean-zero, stationary process, x_t . We refer the reader to Hannan (1970, §2.3) for details. This version allows us to think of a stationary process as being generated (approximately) by a random sum of sines and cosines such as described in (4.3).

Theorem C.2 *If x_t is a mean-zero stationary process, with spectral distribution $F(\omega)$ as given in Theorem C.1, then there exists a complex-valued stochastic process $Z(\omega)$, on the interval $\omega \in [-1/2, 1/2]$, having stationary uncorrelated increments, such that x_t can be written as the stochastic integral*

$$x_t = \int_{-1/2}^{1/2} \exp(-2\pi i t \omega) dZ(\omega),$$

where, for $-1/2 \leq \omega_1 \leq \omega_2 \leq 1/2$,

$$\text{var} \{Z(\omega_2) - Z(\omega_1)\} = F(\omega_2) - F(\omega_1).$$

An uncorrelated increment process such as $Z(\omega)$ is a mean-zero, finite variance, continuous-time stochastic process for which events occurring in non-overlapping intervals are uncorrelated; for example, consider Brownian motion defined in Definition 5.1. The integral in this representation is a stochastic integral. To understand its meaning, let $\omega_0, \omega_1, \dots, \omega_n$ be a partition of the interval $[-1/2, 1/2]$. Define

$$I_n = \sum_{j=1}^n \exp(-2\pi i t \omega_j) [Z(\omega_j) - Z(\omega_{j-1})].$$

Then, $I = \int_{-1/2}^{1/2} \exp(-2\pi i t \omega) dZ(\omega)$ is defined to be the mean square limit of I_n as $n \rightarrow \infty$, assuming it exists.

In general, the spectral distribution function can be a mixture of discrete and continuous distributions. The special case of greatest interest is the absolutely continuous case, namely, when $dF(\omega) = f(\omega)d\omega$, and the resulting function is the spectral density considered in §4.3. What made the proof of Theorem C.1 difficult was that, after we defined

$$f_n(\omega) = \sum_{h=-(n-1)}^{(n-1)} \left(1 - \frac{|h|}{n}\right) \gamma(h) e^{-2\pi i \omega h}$$

in (C.2), we could not simply allow $n \rightarrow \infty$ because $\gamma(h)$ may not be absolutely summable. If, however, $\gamma(h)$ is absolutely summable we may define $f(\omega) = \lim_{n \rightarrow \infty} f_n(\omega)$, and we have the following result.

Theorem C.3 *If $\gamma(h)$ is the autocovariance function of a stationary process, x_t , with*

$$\sum_{h=-\infty}^{\infty} |\gamma(h)| < \infty, \tag{C.3}$$

then the spectral density of x_t is given by

$$f(\omega) = \sum_{h=-\infty}^{\infty} \gamma(h) e^{-2\pi i \omega h}. \quad (\text{C.4})$$

We may extend the representation to the vector case $\mathbf{x}_t = (x_{t1}, \dots, x_{tp})'$ by considering linear combinations of the form

$$y_t = \sum_{j=1}^p a_j^* x_{tj},$$

which will be stationary with autocovariance functions of the form

$$\gamma_y(h) = \sum_{j=1}^p \sum_{k=1}^p a_j^* \gamma_{jk}(h) a_k,$$

where $\gamma_{jk}(h)$ is the usual cross-covariance function between x_{tj} and x_{tk} . To develop the spectral representation of $\gamma_{jk}(h)$ from the representations of the univariate series, consider the linear combinations

$$y_{t1} = x_{tj} + x_{tk} \quad \text{and} \quad y_{t2} = x_{tj} + i x_{tk},$$

which are both stationary series with respective covariance functions

$$\begin{aligned} \gamma_1(h) &= \gamma_{jj}(h) + \gamma_{jk}(h) + \gamma_{kj}(h) + \gamma_{kk}(h) \\ &= \int_{-1/2}^{1/2} e^{2\pi i \omega h} dG_1(\omega), \end{aligned}$$

$$\begin{aligned} \gamma_2(h) &= \gamma_{jj}(h) + i\gamma_{kj}(h) - i\gamma_{jk}(h) + \gamma_{kk}(h) \\ &= \int_{-1/2}^{1/2} e^{2\pi i \omega h} dG_2(\omega). \end{aligned}$$

Introducing the spectral representations for $\gamma_{jj}(h)$ and $\gamma_{kk}(h)$ yields

$$\gamma_{jk}(h) = \int_{-1/2}^{1/2} e^{2\pi i \omega h} dF_{jk}(\omega),$$

with

$$F_{jk}(\omega) = \frac{1}{2} \left[G_1(\omega) + iG_2(\omega) - (1+i)(F_{jj}(\omega) + F_{kk}(\omega)) \right].$$

Now, under the summability condition

$$\sum_{h=-\infty}^{\infty} |\gamma_{jk}(h)| < \infty,$$

we have the representation

$$\gamma_{jk}(h) = \int_{-1/2}^{1/2} e^{2\pi i \omega h} f_{jk}(\omega) d\omega,$$

where the cross-spectral density function has the inverse Fourier representation

$$f_{jk}(\omega) = \sum_{h=-\infty}^{\infty} \gamma_{jk}(h) e^{-2\pi i \omega h}.$$

The cross-covariance function satisfies $\gamma_{jk}(h) = \gamma_{kj}(-h)$, which implies $f_{jk}(\omega) = f_{kj}(-\omega)$ using the above representation.

Then, defining the autocovariance function of the general vector process $\mathbf{x}_t$ as the $p \times p$ matrix

$$\Gamma(h) = E[(\mathbf{x}_{t+h} - \boldsymbol{\mu}_x)(\mathbf{x}_t - \boldsymbol{\mu}_x)'],$$

and the $p \times p$ spectral matrix as $f(\omega) = \{f_{jk}(\omega); j, k = 1, \dots, p\}$, we have the representation in matrix form, written as

$$\Gamma(h) = \int_{-1/2}^{1/2} e^{2\pi i \omega h} f(\omega) d\omega, \quad (\text{C.5})$$

and the inverse result

$$f(\omega) = \sum_{h=-\infty}^{\infty} \Gamma(h) e^{-2\pi i \omega h}. \quad (\text{C.6})$$

which appears as **Property 4.7** in **§4.6. Theorem C.2** can also be extended to the multivariate case.

C.2 Large Sample Distribution of the DFT and Smoothed Periodogram

We have previously introduced the DFT, for the stationary zero-mean process x_t , observed at $t = 1, \dots, n$ as

$$d(\omega) = n^{-1/2} \sum_{t=1}^n x_t e^{-2\pi i \omega t}, \quad (\text{C.7})$$

as the result of matching sines and cosines of frequency ω against the series x_t . We will suppose now that x_t has an absolutely continuous spectrum $f(\omega)$ corresponding to the absolutely summable autocovariance function $\gamma(h)$. Our purpose in this section is to examine the statistical properties of the complex random variables $d(\omega_k)$, for $\omega_k = k/n$, $k = 0, 1, \dots, n-1$ in providing a basis for the estimation of $f(\omega)$. To develop the statistical properties, we examine the behavior of

$$\begin{aligned}
S_n(\omega, \omega) &= E |d(\omega)|^2 \\
&= n^{-1} E \left[\sum_{s=1}^n x_s e^{-2\pi i \omega s} \sum_{t=1}^n x_t e^{2\pi i \omega t} \right] \\
&= n^{-1} \sum_{s=1}^n \sum_{t=1}^n e^{-2\pi i \omega s} e^{2\pi i \omega t} \gamma(s-t) \\
&= \sum_{h=-(n-1)}^{n-1} (1 - |h|/n) \gamma(h) e^{-2\pi i \omega h}, \tag{C.8}
\end{aligned}$$

where we have let $h = s - t$. Using dominated convergence,

$$S_n(\omega, \omega) \rightarrow \sum_{h=-\infty}^{\infty} \gamma(h) e^{-2\pi i \omega h} = f(\omega),$$

as $n \rightarrow \infty$, making the large sample variance of the Fourier transform equal to the spectrum evaluated at ω . We have already seen this result in [Theorem C.3](#). For exact bounds it is also convenient to add an absolute summability assumption for the autocovariance function, namely,

$$\theta = \sum_{h=-\infty}^{\infty} |h| |\gamma(h)| < \infty. \tag{C.9}$$

Example C.1 Condition (C.9) Verified for an AR(1)

For a causal AR(1) series, $x_t = \phi x_{t-1} + w_t$, we may write θ in (C.9) as

$$\theta = \frac{\sigma_w^2}{1 - \phi^2} \sum_{h=-\infty}^{\infty} |h| \phi^{|h|}.$$

The condition that θ is finite is equivalent to summability of

$$\sum_{h=1}^{\infty} h \phi^h = \phi \sum_{h=1}^{\infty} h \phi^{h-1} = \phi \frac{\partial}{\partial \phi} \sum_{h=1}^{\infty} \phi^h = \frac{\phi}{(1 - \phi)^2},$$

and hence,

$$\theta = \frac{2\sigma_w^2 \phi}{(1 - \phi)^3 (1 + \phi)} < \infty.$$

To elaborate further, we derive two approximation lemmas.

Lemma C.1 For $S_n(\omega, \omega)$ as defined in (C.8) and θ in (C.9) finite, we have

$$|S_n(\omega, \omega) - f(\omega)| \leq \frac{\theta}{n} \tag{C.10}$$

or

$$S_n(\omega, \omega) = f(\omega) + O(n^{-1}). \tag{C.11}$$

Proof: To prove the lemma, write

$$\begin{aligned}
 n|S_n(\omega, \omega) - f_X(\omega)| &= \left| \sum_{|u| < n} (n - |u|)\gamma(u)e^{-2\pi i \omega u} - n \sum_{u=-\infty}^{\infty} \gamma(u)e^{-2\pi i \omega u} \right| \\
 &= \left| -n \sum_{|u| \geq n} \gamma(u)e^{-2\pi i \omega u} - \sum_{|u| < n} |u|\gamma(u)e^{-2\pi i \omega u} \right| \\
 &\leq \sum_{|u| \geq n} |u||\gamma(u)| + \sum_{|u| < n} |u||\gamma(u)| \\
 &= \theta,
 \end{aligned}$$

which establishes the lemma. $\square$

Lemma C.2 For $\omega_k = k/n$, $\omega_\ell = \ell/n$, $\omega_k - \omega_\ell \neq 0, \pm 1, \pm 2, \pm 3, \dots$, and θ in (C.9), we have

$$|S_n(\omega_k, \omega_\ell)| \leq \frac{\theta}{n} = O(n^{-1}), \quad (\text{C.12})$$

where

$$S_n(\omega_k, \omega_\ell) = E\{d(\omega_k)\overline{d(\omega_\ell)}\}. \quad (\text{C.13})$$

Proof: Write

$$\begin{aligned}
 n|S_n(\omega_k, \omega_\ell)| &= \sum_{u=-(n-1)}^{-1} \gamma(u) \sum_{v=-(u-1)}^n e^{-2\pi i(\omega_k - \omega_\ell)v} e^{-2\pi i \omega_k u} \\
 &\quad + \sum_{u=0}^{n-1} \gamma(u) \sum_{v=1}^{n-u} e^{-2\pi i(\omega_k - \omega_\ell)v} e^{-2\pi i \omega_k u}.
 \end{aligned}$$

Now, for the first term, with $u < 0$,

$$\begin{aligned}
 \sum_{v=-(u-1)}^n e^{-2\pi i(\omega_k - \omega_\ell)v} &= \left(\sum_{v=1}^n - \sum_{v=1}^{-u} \right) e^{-2\pi i(\omega_k - \omega_\ell)v} \\
 &= 0 - \sum_{v=1}^{-u} e^{-2\pi i(\omega_k - \omega_\ell)v}.
 \end{aligned}$$

For the second term with $u \geq 0$,

$$\begin{aligned}
 \sum_{v=1}^{n-u} e^{-2\pi i(\omega_k - \omega_\ell)v} &= \left(\sum_{v=1}^n - \sum_{v=n-u+1}^n \right) e^{-2\pi i(\omega_k - \omega_\ell)v} \\
 &= 0 - \sum_{v=n-u+1}^n e^{-2\pi i(\omega_k - \omega_\ell)v}.
 \end{aligned}$$

Consequently,

$$\begin{aligned}
n|S_n(\omega_k, \omega_\ell)| &= \left| - \sum_{u=-(n-1)}^{-1} \gamma(u) \sum_{v=1}^{-u} e^{-2\pi i(\omega_k - \omega_\ell)v} e^{-2\pi i\omega_k u} \right. \\
&\quad \left. - \sum_{u=1}^{n-1} \gamma(u) \sum_{v=n-u+1}^n e^{-2\pi i(\omega_k - \omega_\ell)v} e^{-2\pi i\omega_k u} \right| \\
&\leq \sum_{u=-(n-1)}^0 (-u)|\gamma(u)| + \sum_{u=1}^{n-1} u|\gamma(u)| \\
&= \sum_{u=-(n-1)}^{(n-1)} |u| |\gamma(u)|.
\end{aligned}$$

Hence, we have

$$S_n(\omega_k, \omega_\ell) \leq \frac{\theta}{n},$$

and the asserted relations of the lemma follow. $\square$

Because the DFTs are approximately uncorrelated, say, of order $1/n$, when the frequencies are of the form $\omega_k = k/n$, we shall compute at those frequencies. The behavior of $f(\omega)$ at neighboring frequencies, however, will often be of interest and we shall use [Lemma C.3](#) below to handle such cases.

Lemma C.3 For $|\omega_k - \omega| \leq L/2n$ and θ in (C.9), we have

$$|f(\omega_k) - f(\omega)| \leq \frac{\pi\theta L}{n} \quad (\text{C.14})$$

or

$$f(\omega_k) - f(\omega) = O(L/n). \quad (\text{C.15})$$

Proof: Write the difference

$$\begin{aligned}
|f(\omega_k) - f(\omega)| &= \left| \sum_{h=-\infty}^{\infty} \gamma(h) (e^{-2\pi i\omega_k h} - e^{-2\pi i\omega h}) \right| \\
&\leq \sum_{h=-\infty}^{\infty} |\gamma(h)| |e^{-\pi i(\omega_k - \omega)h} - e^{\pi i(\omega_k - \omega)h}| \\
&= 2 \sum_{h=-\infty}^{\infty} |\gamma(h)| |\sin[\pi(\omega_k - \omega)h]| \\
&\leq 2\pi|\omega_k - \omega| \sum_{h=-\infty}^{\infty} |h| |\gamma(h)| \\
&\leq \frac{\pi\theta L}{n}
\end{aligned}$$

because $|\sin x| \leq |x|$. $\square$

The main use of the properties described by [Lemma C.1](#) and [Lemma C.2](#) is in identifying the covariance structure of the DFT, say,

$$\begin{aligned} d(\omega_k) &= n^{-1/2} \sum_{t=1}^n x_t e^{-2\pi i \omega_k t} \\ &= d_c(\omega_k) - i d_s(\omega_k), \end{aligned}$$

where

$$d_c(\omega_k) = n^{-1/2} \sum_{t=1}^n x_t \cos(2\pi \omega_k t)$$

and

$$d_s(\omega_k) = n^{-1/2} \sum_{t=1}^n x_t \sin(2\pi \omega_k t)$$

are the cosine and sine transforms, respectively, of the observed series, defined previously in (4.26) and (4.27). For example, assuming zero means for convenience, we will have

$$\begin{aligned} E[d_c(\omega_k) d_c(\omega_\ell)] &= \frac{1}{4} n^{-1} \sum_{s=1}^n \sum_{t=1}^n \gamma(s-t) (e^{2\pi i \omega_k s} + e^{-2\pi i \omega_k s}) (e^{2\pi i \omega_\ell t} + e^{-2\pi i \omega_\ell t}) \\ &= \frac{1}{4} [S_n(-\omega_k, \omega_\ell) + S_n(\omega_k, \omega_\ell) + S_n(\omega_\ell, \omega_k) + S_n(\omega_k, -\omega_\ell)]. \end{aligned}$$

[Lemma C.1](#) and [Lemma C.2](#) imply, for $k = \ell$,

$$\begin{aligned} E[d_c(\omega_k) d_c(\omega_\ell)] &= \frac{1}{4} [O(n^{-1}) + f(\omega_k) + O(n^{-1}) \\ &\quad + f(\omega_k) + O(n^{-1}) + O(n^{-1})] \\ &= \frac{1}{2} f(\omega_k) + O(n^{-1}). \end{aligned} \tag{C.16}$$

For $k \neq \ell$, all terms are $O(n^{-1})$. Hence, we have

$$E[d_c(\omega_k) d_c(\omega_\ell)] = \begin{cases} \frac{1}{2} f(\omega_k) + O(n^{-1}), & k = \ell \\ O(n^{-1}), & k \neq \ell. \end{cases} \tag{C.17}$$

A similar argument gives

$$E[d_s(\omega_k) d_s(\omega_\ell)] = \begin{cases} \frac{1}{2} f(\omega_k) + O(n^{-1}), & k = \ell \\ O(n^{-1}), & k \neq \ell \end{cases} \tag{C.18}$$

and we also have $E[d_s(\omega_k) d_c(\omega_\ell)] = O(n^{-1})$ for all k, ℓ . We may summarize the results of [Lemma C.1](#)–[Lemma C.3](#) as follows.

Theorem C.4 *For a stationary mean zero process with autocovariance function satisfying (C.9) and frequencies $\omega_{k:n}$, such that $|\omega_{k:n} - \omega| < 1/n$, are close to some target frequency ω , the cosine and sine transforms (4.26) and (4.27) are approximately uncorrelated with variances equal to $(1/2)f(\omega)$, and the error in the approximation can be uniformly bounded by $\pi\theta L/n$.*

Now, consider estimating the spectrum in a neighborhood of some target frequency ω , using the periodogram estimator

$$I(\omega_{k:n}) = |d(\omega_{k:n})|^2 = d_c^2(\omega_{k:n}) + d_s^2(\omega_{k:n}),$$

where we take $|\omega_{k:n} - \omega| \leq n^{-1}$ for each n . In case the series x_t is Gaussian with zero mean,

$$\begin{pmatrix} d_c(\omega_{k:n}) \\ d_s(\omega_{k:n}) \end{pmatrix} \xrightarrow{d} N\left\{\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \frac{1}{2} \begin{pmatrix} f(\omega) & 0 \\ 0 & f(\omega) \end{pmatrix}\right\},$$

and we have that

$$\frac{2 I(\omega_{k:n})}{f(\omega)} \xrightarrow{d} \chi_2^2,$$

where χ_ν^2 denotes a chi-squared random variable with ν degrees of freedom, as usual. Unfortunately, the distribution does not become more concentrated as $n \rightarrow \infty$, because the variance of the periodogram estimator does not go to zero.

We develop a fix for the deficiencies mentioned above by considering the average of the periodogram over a set of frequencies in the neighborhood of ω . For example, we can always find a set of $L = 2m + 1$ frequencies of the form $\{\omega_{j:n} + k/n; k = 0, \pm 1, \pm 2, \dots, m\}$, for which

$$f(\omega_{j:n} + k/n) = f(\omega) + O(Ln^{-1})$$

by **Lemma C.3**. As n increases, the values of the separate frequencies change.

Now, we can consider the smoothed periodogram estimator, $\hat{f}(\omega)$, given in (4.59); this case includes the averaged periodogram, $\bar{f}(\omega)$. First, we note that (C.9), $\theta = \sum_{h=-\infty}^{\infty} |h||\gamma(h)| < \infty$, is a crucial condition in the estimation of spectra. In investigating local averages of the periodogram, we will require a condition on the rate of (C.9), namely

$$\sum_{h=-n}^n |h||\gamma(h)| = O(n^{-1/2}). \quad (\text{C.19})$$

One can show that a sufficient condition for (C.19) is that the time series is the linear process given by,

$$x_t = \sum_{j=-\infty}^{\infty} \psi_j w_{t-j}, \quad \sum_{j=0}^{\infty} \sqrt{j} |\psi_j| < \infty \quad (\text{C.20})$$

where $w_t \sim \text{iid}(0, \sigma_w^2)$ and w_t has finite fourth moment,

$$E(w_t^4) = \eta \sigma_w^4 < \infty.$$

We leave it to the reader (see [Problem 4.37](#) for more details) to show (C.20) implies (C.19). If $w_t \sim wn(0, \sigma_w^2)$, then (C.20) implies (C.19), but we will require the noise to be iid in the following lemma.

Lemma C.4 *Suppose x_t is the linear process given by (C.20), and let $I(\omega_j)$ be the periodogram of the data $\{x_1, \dots, x_n\}$. Then*

$$\text{cov}\left(I(\omega_j), I(\omega_k)\right) = \begin{cases} 2f^2(\omega_j) + o(1) & \omega_j = \omega_k = 0, 1/2 \\ f^2(\omega_j) + o(1) & \omega_j = \omega_k \neq 0, 1/2 \\ O(n^{-1}) & \omega_j \neq \omega_k. \end{cases}$$

The proof of [Lemma C.4](#) is straightforward but tedious, and details may be found in Fuller (1976, Theorem 7.2.1) or in Brockwell and Davis (1991, Theorem 10.3.2). For demonstration purposes, we present the proof of the lemma for the pure white noise case; i.e., $x_t = w_t$, in which case $f(\omega) \equiv \sigma_w^2$. By definition, the periodogram in this case is

$$I(\omega_j) = n^{-1} \sum_{s=1}^n \sum_{t=1}^n w_s w_t e^{2\pi i \omega_j (t-s)},$$

where $\omega_j = j/n$, and hence

$$E\{I(\omega_j)I(\omega_k)\} = n^{-2} \sum_{s=1}^n \sum_{t=1}^n \sum_{u=1}^n \sum_{v=1}^n E(w_s w_t w_u w_v) e^{2\pi i \omega_j (t-s)} e^{2\pi i \omega_k (u-v)}.$$

Now when all the subscripts match, $E(w_s w_t w_u w_v) = \eta \sigma_w^4$, when the subscripts match in pairs (e.g., $s = t \neq u = v$), $E(w_s w_t w_u w_v) = \sigma_w^4$, otherwise, $E(w_s w_t w_u w_v) = 0$. Thus,

$$E\{I(\omega_j)I(\omega_k)\} = n^{-1}(\eta - 3)\sigma_w^4 + \sigma_w^4 \left(1 + n^{-2}[A(\omega_j + \omega_k) + A(\omega_k - \omega_j)]\right),$$

where

$$A(\lambda) = \left| \sum_{t=1}^n e^{2\pi i \lambda t} \right|^2.$$

Noting that $E I(\omega_j) = n^{-1} \sum_{t=1}^n E(w_t^2) = \sigma_w^2$, we have

$$\begin{aligned} \text{cov}\{I(\omega_j), I(\omega_k)\} &= E\{I(\omega_j)I(\omega_k)\} - \sigma_w^4 \\ &= n^{-1}(\eta - 3)\sigma_w^4 + n^{-2}\sigma_w^4[A(\omega_j + \omega_k) + A(\omega_k - \omega_j)]. \end{aligned}$$

Thus we conclude that

$$\begin{aligned} \text{var}\{I(\omega_j)\} &= n^{-1}(\eta - 3)\sigma_w^4 + \sigma_w^4 && \text{for } \omega_j \neq 0, 1/2 \\ \text{var}\{I(\omega_j)\} &= n^{-1}(\eta - 3)\sigma_w^4 + 2\sigma_w^4 && \text{for } \omega_j = 0, 1/2 \\ \text{cov}\{I(\omega_j), I(\omega_k)\} &= n^{-1}(\eta - 3)\sigma_w^4 && \text{for } \omega_j \neq \omega_k, \end{aligned}$$

which establishes the result in this case. We also note that if w_t is Gaussian, then $\eta = 3$ and the periodogram ordinates are independent. Using [Lemma C.4](#), we may establish the following fundamental result.

Theorem C.5 Suppose x_t is the linear process given by (C.20). Then, with $\widehat{f}(\omega)$ defined in (4.59) and corresponding conditions on the weights h_k , we have, as $n \rightarrow \infty$,

$$(i) E \left(\widehat{f}(\omega) \right) \rightarrow f(\omega)$$

$$(ii) \left(\sum_{k=-m}^m h_k^2 \right)^{-1} \text{cov} \left(\widehat{f}(\omega), \widehat{f}(\lambda) \right) \rightarrow f^2(\omega) \quad \text{for } \omega = \lambda \neq 0, 1/2.$$

In (ii), replace $f^2(\omega)$ by 0 if $\omega \neq \lambda$ and by $2f^2(\omega)$ if $\omega = \lambda = 0$ or $1/2$.

Proof: (i): First, recall (4.31)

$$E \left[I(\omega_{j:n}) \right] = \sum_{h=-(n-1)}^{n-1} \left(\frac{n-|h|}{n} \right) \gamma(h) e^{-2\pi i \omega_{j:n} h} \stackrel{\text{def}}{=} f_n(\omega_{j:n}).$$

But since $f_n(\omega_{j:n}) \rightarrow f(\omega)$ uniformly, and $|f(\omega_{j:n}) - f(\omega_{j:n} + k/n)| \rightarrow 0$ by the continuity of f , we have

$$\begin{aligned} E \widehat{f}(\omega) &= \sum_{k=-m}^m h_k E I(\omega_{j:n} + k/n) = \sum_{k=-m}^m h_k f_n(\omega_{j:n} + k/n) \\ &= \sum_{k=-m}^m h_k [f(\omega) + o(1)] \rightarrow f(\omega), \end{aligned}$$

because $\sum_{k=-m}^m h_k = 1$.

(ii): First, suppose we have $\omega_{j:n} \rightarrow \omega_1$ and $\omega_{\ell:n} \rightarrow \omega_2$, and $\omega_1 \neq \omega_2$. Then, for n large enough to separate the bands, using Lemma C.4, we have

$$\begin{aligned} \left| \text{cov} \left(\widehat{f}(\omega_1), \widehat{f}(\omega_2) \right) \right| &= \left| \sum_{|k| \leq m} \sum_{|r| \leq m} h_k h_r \text{cov} [I(\omega_{j:n} + k/n), I(\omega_{\ell:n} + r/n)] \right| \\ &= \left| \sum_{|k| \leq m} \sum_{|r| \leq m} h_k h_r O(n^{-1}) \right| \\ &\leq \frac{c}{n} \left(\sum_{|k| \leq m} h_k \right)^2 \quad (\text{where } c \text{ is a constant}) \\ &\leq \frac{cL}{n} \left(\sum_{|k| \leq m} h_k^2 \right), \end{aligned}$$

which establishes (ii) for the case of different frequencies. The case of the same frequencies, i.e., $\omega = \lambda$, is established in a similar manner to the above arguments. $\square$

Theorem C.5 justifies the distributional properties used throughout §4.5 and Chapter 7. We may extend the results of this section to vector series of the form $\mathbf{x}_t = (x_{t1}, \dots, x_{tp})'$, when the cross-spectrum is given by

$$f_{ij}(\omega) = \sum_{h=-\infty}^{\infty} \gamma_{ij}(h) e^{-2\pi i \omega h} = c_{ij}(\omega) - i q_{ij}(\omega), \quad (\text{C.21})$$

where

$$c_{ij}(\omega) = \sum_{h=-\infty}^{\infty} \gamma_{ij}(h) \cos(2\pi \omega h) \quad (\text{C.22})$$

and

$$q_{ij}(\omega) = \sum_{h=-\infty}^{\infty} \gamma_{ij}(h) \sin(2\pi \omega h) \quad (\text{C.23})$$

denote the cospectrum and quadpectrum, respectively. We denote the DFT of the series x_{tj} by

$$\begin{aligned} d_j(\omega_k) &= n^{-1/2} \sum_{t=1}^n x_{tj} e^{-2\pi i \omega_k t} \\ &= d_{cj}(\omega_k) - i d_{sj}(\omega_k), \end{aligned}$$

where d_{cj} and d_{sj} are the cosine and sine transforms of x_{tj} , for $j = 1, 2, \dots, p$. We bound the covariance structure as before and summarize the results as follows.

Theorem C.6 *The covariance structure of the multivariate cosine and sine transforms, subject to*

$$\theta_{ij} = \sum_{h=-\infty}^{\infty} |h| |\gamma_{ij}(h)| < \infty, \quad (\text{C.24})$$

is given by

$$E[d_{ci}(\omega_k) d_{cj}(\omega_\ell)] = \begin{cases} \frac{1}{2} c_{ij}(\omega_k) + O(n^{-1}), & k = \ell \\ O(n^{-1}), & k \neq \ell. \end{cases} \quad (\text{C.25})$$

$$E[d_{ci}(\omega_k) d_{sj}(\omega_\ell)] = \begin{cases} -\frac{1}{2} q_{ij}(\omega_k) + O(n^{-1}), & k = \ell \\ O(n^{-1}), & k \neq \ell \end{cases} \quad (\text{C.26})$$

$$E[d_{si}(\omega_k) d_{cj}(\omega_\ell)] = \begin{cases} \frac{1}{2} q_{ij}(\omega_k) + O(n^{-1}), & k = \ell \\ O(n^{-1}), & k \neq \ell \end{cases} \quad (\text{C.27})$$

$$E[d_{si}(\omega_k) d_{sj}(\omega_\ell)] = \begin{cases} \frac{1}{2} c_{ij}(\omega_k) + O(n^{-1}), & k = \ell \\ O(n^{-1}), & k \neq \ell. \end{cases} \quad (\text{C.28})$$

Proof: We define

$$S_n^{ij}(\omega_k, \omega_\ell) = \sum_{s=1}^n \sum_{t=1}^n \gamma_{ij}(s-t) e^{-2\pi i \omega_k s} e^{2\pi i \omega_\ell t}. \quad (\text{C.29})$$

Then, we may verify the theorem with manipulations like

$$\begin{aligned}
E[d_{ci}(\omega_k)d_{sj}(\omega_k)] &= \frac{1}{4i} \sum_{s=1}^n \sum_{t=1}^n \gamma_{ij}(s-t)(e^{2\pi i \omega_k s} + e^{-2\pi i \omega_k s})(e^{2\pi i \omega_k t} - e^{-2\pi i \omega_k t}) \\
&= \frac{1}{4i} \left[S_n^{ij}(-\omega_k, \omega_k) + S_n^{ij}(\omega_k, \omega_k) - S_n^{ij}(\omega_k, \omega_k) - S_n^{ij}(\omega_k, -\omega_k) \right] \\
&= \frac{1}{4i} \left[c_{ij}(\omega_k) - iq_{ij}(\omega_k) - (c_{ij}(\omega_k) + iq_{ij}(\omega_k)) + O(n^{-1}) \right] \\
&= -\frac{1}{2} q_{ij}(\omega_k) + O(n^{-1}),
\end{aligned}$$

where we have used the fact that the properties given in [Lemma C.1](#)–[Lemma C.3](#) can be verified for the cross-spectral density functions $f_{ij}(\omega)$, $i, j = 1, \dots, p$. $\square$

Now, if the underlying multivariate time series $\mathbf{x}_t$ is a normal process, it is clear that the DFTs will be jointly normal and we may define the vector DFT, $\mathbf{d}(\omega_k) = (d_1(\omega_k), \dots, d_p(\omega_k))'$ as

$$\begin{aligned}
\mathbf{d}(\omega_k) &= n^{-1/2} \sum_{t=1}^n \mathbf{x}_t e^{-2\pi i \omega_k t} \\
&= \mathbf{d}_c(\omega_k) - i\mathbf{d}_s(\omega_k),
\end{aligned} \tag{C.30}$$

where

$$\mathbf{d}_c(\omega_k) = n^{-1/2} \sum_{t=1}^n \mathbf{x}_t \cos(2\pi \omega_k t) \tag{C.31}$$

and

$$\mathbf{d}_s(\omega_k) = n^{-1/2} \sum_{t=1}^n \mathbf{x}_t \sin(2\pi \omega_k t) \tag{C.32}$$

are the cosine and sine transforms, respectively, of the observed vector series $\mathbf{x}_t$. Then, constructing the vector of real and imaginary parts $(\mathbf{d}'_c(\omega_k), \mathbf{d}'_s(\omega_k))'$, we may note it has mean zero and $2p \times 2p$ covariance matrix

$$\Sigma(\omega_k) = \frac{1}{2} \begin{pmatrix} C(\omega_k) & -Q(\omega_k) \\ Q(\omega_k) & C(\omega_k) \end{pmatrix} \tag{C.33}$$

to order n^{-1} as long as $\omega_k - \omega = O(n^{-1})$. We have introduced the $p \times p$ matrices $C(\omega_k) = \{c_{ij}(\omega_k)\}$ and $Q = \{q_{ij}(\omega_k)\}$. The complex random variable $\mathbf{d}(\omega_k)$ has covariance

$$\begin{aligned}
S(\omega_k) &= E[\mathbf{d}(\omega_k)\mathbf{d}^*(\omega_k)] \\
&= E\left[(\mathbf{d}_c(\omega_k) - i\mathbf{d}_s(\omega_k))(\mathbf{d}_c(\omega_k) - i\mathbf{d}_s(\omega_k))^*\right] \\
&= E[\mathbf{d}_c(\omega_k)\mathbf{d}_c(\omega_k)'] + E[\mathbf{d}_s(\omega_k)\mathbf{d}_s(\omega_k)'] \\
&\quad - i(E[\mathbf{d}_s(\omega_k)\mathbf{d}_c(\omega_k)'] - E[\mathbf{d}_c(\omega_k)\mathbf{d}_s(\omega_k)']) \\
&= C(\omega_k) - iQ(\omega_k).
\end{aligned} \tag{C.34}$$

If the process $\mathbf{x}_t$ has a multivariate normal distribution, the complex vector $\mathbf{d}(\omega_k)$ has approximately the complex multivariate normal distribution with mean zero and covariance matrix $S(\omega_k) = C(\omega_k) - iQ(\omega_k)$ if the real and imaginary parts have the covariance structure as specified above. In the next section, we work further with this distribution and show how it adapts to the real case. If we wish to estimate the spectral matrix $S(\omega)$, it is natural to take a band of frequencies of the form $\omega_{k:n} + \ell/n$, for $\ell = -m, \dots, m$ as before, so that the estimator becomes (4.94) of §4.6. A discussion of further properties of the multivariate complex normal distribution is deferred.

It is also of interest to develop a large sample theory for cases in which the underlying distribution is not necessarily normal. If x_t is not necessarily a normal process, some additional conditions are needed to get asymptotic normality. In particular, introduce the notion of a generalized linear process

$$\mathbf{y}_t = \sum_{r=-\infty}^{\infty} A_r \mathbf{w}_{t-r}, \quad (\text{C.35})$$

where $\mathbf{w}_t$ is a $p \times 1$ vector white noise process with $p \times p$ covariance $E[\mathbf{w}_t \mathbf{w}_t'] = G$ and the $p \times p$ matrices of filter coefficients A_t satisfy

$$\sum_{t=-\infty}^{\infty} \text{tr}\{A_t A_t'\} = \sum_{t=-\infty}^{\infty} \|A_t\|^2 < \infty. \quad (\text{C.36})$$

In particular, stable vector ARMA processes satisfy these conditions. For generalized linear processes, we state the following general result from Hannan (1970, p.224).

Theorem C.7 *If $\mathbf{x}_t$ is generated by a generalized linear process with a continuous spectrum that is not zero at ω and $\omega_{k:n} + \ell/n$ are a set of frequencies within L/n of ω , the joint density of the cosine and sine transforms (C.31) and (C.32) converges to that of L independent $2p \times 1$ normal vectors with covariance matrix $\Sigma(\omega)$ with structure given by (C.33). At $\omega = 0$ or $\omega = 1/2$, the distribution is real with covariance matrix $2\Sigma(\omega)$.*

The above result provides the basis for inference involving the Fourier transforms of stationary series because it justifies approximations to the likelihood function based on multivariate normal theory. We make extensive use of this result in Chapter 7, but will still need a simple form to justify the distributional result for the sample coherence given in (4.100). The next section gives an elementary introduction to the complex normal distribution.

C.3 The Complex Multivariate Normal Distribution

The multivariate normal distribution will be the fundamental tool for expressing the likelihood function and determining approximate maximum likelihood estimators and their large sample probability distributions. A detailed treatment of the multivariate

normal distribution can be found in standard texts such as Anderson (1984). We will use the multivariate normal distribution of the $p \times 1$ vector $\mathbf{x} = (x_1, x_2, \dots, x_p)'$, as defined by its density function

$$p(\mathbf{x}) = (2\pi)^{-p/2} |\Sigma|^{-1/2} \exp\left\{-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})' \Sigma^{-1} (\mathbf{x} - \boldsymbol{\mu})\right\}, \quad (\text{C.37})$$

which can be shown to have mean vector $E[\mathbf{x}] = \boldsymbol{\mu} = (\mu_1, \dots, \mu_p)'$ and covariance matrix

$$\Sigma = E[(\mathbf{x} - \boldsymbol{\mu})(\mathbf{x} - \boldsymbol{\mu})']. \quad (\text{C.38})$$

We use the notation $\mathbf{x} \sim N_p(\boldsymbol{\mu}, \Sigma)$ for densities of the form (C.37) and note that linearly transformed multivariate normal variables of the form $\mathbf{y} = A\mathbf{x}$, with A a $q \times p$ matrix $q \leq p$, will also be multivariate normal with distribution

$$\mathbf{y} \sim N_q(A\boldsymbol{\mu}, A\Sigma A'). \quad (\text{C.39})$$

Often, the partitioned multivariate normal, based on the vector $\mathbf{x} = (\mathbf{x}'_1, \mathbf{x}'_2)'$, split into two $p_1 \times 1$ and $p_2 \times 1$ components $\mathbf{x}_1$ and $\mathbf{x}_2$, respectively, will be used where $p = p_1 + p_2$. If the mean vector $\boldsymbol{\mu} = (\boldsymbol{\mu}'_1, \boldsymbol{\mu}'_2)'$ and covariance matrices

$$\Sigma = \begin{pmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{pmatrix} \quad (\text{C.40})$$

are also compatibly partitioned, the marginal distribution of any subset of components is multivariate normal, say,

$$\mathbf{x}_1 \sim N_{p_1}\{\boldsymbol{\mu}_1, \Sigma_{11}\},$$

and that the conditional distribution $\mathbf{x}_2$ given $\mathbf{x}_1$ is normal with mean

$$E[\mathbf{x}_2|\mathbf{x}_1] = \boldsymbol{\mu}_2 + \Sigma_{21}\Sigma_{11}^{-1}(\mathbf{x}_1 - \boldsymbol{\mu}_1) \quad (\text{C.41})$$

and conditional covariance

$$\text{cov}[\mathbf{x}_2|\mathbf{x}_1] = \Sigma_{22} - \Sigma_{21}\Sigma_{11}^{-1}\Sigma_{12}. \quad (\text{C.42})$$

In the previous section, the real and imaginary parts of the DFT had a partitioned covariance matrix as given in (C.33), and we use this result to say the complex $p \times 1$ vector

$$\mathbf{z} = \mathbf{x}_1 - i\mathbf{x}_2 \quad (\text{C.43})$$

has a complex multivariate normal distribution, with mean vector $\boldsymbol{\mu}_z = \boldsymbol{\mu}_1 - i\boldsymbol{\mu}_2$ and $p \times p$ covariance matrix

$$\Sigma_z = C - iQ \quad (\text{C.44})$$

if the real multivariate $2p \times 1$ normal vector $\mathbf{x} = (\mathbf{x}'_1, \mathbf{x}'_2)'$ has a real multivariate normal distribution with mean vector $\boldsymbol{\mu} = (\boldsymbol{\mu}'_1, \boldsymbol{\mu}'_2)'$ and covariance matrix

$$\Sigma = \frac{1}{2} \begin{pmatrix} C & -Q \\ Q & C \end{pmatrix}. \quad (\text{C.45})$$

The restrictions $C' = C$ and $Q' = -Q$ are necessary for the matrix Σ to be a covariance matrix, and these conditions then imply $\Sigma_z = \Sigma_z^*$ is Hermitian. The probability density function of the complex multivariate normal vector $\mathbf{z}$ can be expressed in the concise form

$$p_{\mathbf{z}}(\mathbf{z}) = \pi^{-p} |\Sigma_z|^{-1} \exp\{-(\mathbf{z} - \boldsymbol{\mu}_z)^* \Sigma_z^{-1} (\mathbf{z} - \boldsymbol{\mu}_z)\}, \quad (\text{C.46})$$

and this is the form that we will often use in the likelihood. The result follows from showing that $p_{\mathbf{x}}(\mathbf{x}_1, \mathbf{x}_2) = p_{\mathbf{z}}(\mathbf{z})$ exactly, using the fact that the quadratic and Hermitian forms in the exponent are equal and that $|\Sigma_x| = |\Sigma_z|^2$. The second assertion follows directly from the fact that the matrix Σ_x has repeated eigenvalues, $\lambda_1, \lambda_2, \dots, \lambda_p$ corresponding to eigenvectors $(\alpha'_1, \alpha'_2)'$ and the same set, $\lambda_1, \lambda_2, \dots, \lambda_p$ corresponding to $(\alpha'_2, -\alpha'_1)'$. Hence

$$|\Sigma_x| = \prod_{i=1}^p \lambda_i^2 = |\Sigma_z|^2.$$

For further material relating to the complex multivariate normal distribution, see Goodman (1963), Giri (1965), or Khatri (1965).

Example C.2 A Complex Normal Random Variable

To fix ideas, consider a very simple complex random variable

$$z = \Re(z) - i\Im(z) = z_1 - iz_2,$$

where $z_1 \sim N(0, \frac{1}{2}\sigma^2)$ independent of $z_2 \sim N(0, \frac{1}{2}\sigma^2)$. Then the joint density of (z_1, z_2) is

$$f(z_1, z_2) \propto \sigma^{-1} \exp\left(-\frac{z_1^2}{\sigma^2}\right) \times \sigma^{-1} \exp\left(-\frac{z_2^2}{\sigma^2}\right) = \sigma^{-2} \exp\left\{-\left(\frac{z_1^2 + z_2^2}{\sigma^2}\right)\right\}.$$

More succinctly, we write $z \sim N_c(0, \sigma^2)$, and

$$f(z) \propto \sigma^{-2} \exp\left(-\frac{z^* z}{\sigma^2}\right).$$

In Fourier analysis, z_1 would be the cosine transform of the data at a fundamental frequency (excluding the end points) and z_2 the corresponding sine transform. If the process is Gaussian, z_1 and z_2 are independent normals with zero means and variances that are half of the spectral density at the particular frequency. Consequently, the definition of the complex normal distribution is natural in the context of spectral analysis.

Example C.3 A Bivariate Complex Normal Distribution

Consider the joint distribution of the complex random variables $u_1 = x_1 - ix_2$ and $u_2 = y_1 - iy_2$, where the partitioned vector $(x_1, x_2, y_1, y_2)'$ has a real multivariate normal distribution with mean $(0, 0, 0, 0)'$ and covariance matrix

$$\Sigma = \frac{1}{2} \left(\begin{array}{cc|cc} c_{xx} & 0 & c_{xy} & -q_{xy} \\ 0 & c_{xx} & q_{xy} & c_{xy} \\ \hline c_{xy} & q_{xy} & c_{yy} & 0 \\ -q_{xy} & c_{yx} & 0 & c_{yy} \end{array} \right). \quad (\text{C.47})$$

Now, consider the conditional distribution of $\mathbf{y} = (y_1, y_2)'$, given $\mathbf{x} = (x_1, x_2)'$. Using (C.41), we obtain

$$E(\mathbf{y} | \mathbf{x}) = \begin{pmatrix} x_1 & -x_2 \\ x_2 & x_1 \end{pmatrix} \begin{pmatrix} b_1 \\ b_2 \end{pmatrix}, \quad (\text{C.48})$$

where

$$(b_1, b_2) = \left(\frac{c_{yx}}{c_{xx}}, \frac{q_{yx}}{c_{xx}} \right). \quad (\text{C.49})$$

It is natural to identify the cross-spectrum

$$f_{xy} = c_{xy} - iq_{xy}, \quad (\text{C.50})$$

so that the complex variable identified with the pair is just

$$b = b_1 - ib_2 = \frac{c_{yx} - iq_{yx}}{c_{xx}} = \frac{f_{yx}}{f_{xx}},$$

and we identify it as the complex regression coefficient. The conditional covariance follows from (C.42) and simplifies to

$$\text{cov}(\mathbf{y} | \mathbf{x}) = \frac{1}{2} f_{y \cdot x} I_2, \quad (\text{C.51})$$

where I_2 denotes the 2×2 identity matrix and

$$f_{y \cdot x} = c_{yy} - \frac{c_{xy}^2 + q_{xy}^2}{c_{xx}} = f_{yy} - \frac{|f_{xy}|^2}{f_{xx}} \quad (\text{C.52})$$

Example C.3 leads to an approach for justifying the distributional results for the function coherence given in (4.100). That equation suggests that the result can be derived using the regression results that lead to the F-statistics in §2.1. Suppose that we consider L values of the sine and cosine transforms of the input x_t and output y_t , which we will denote by $d_{x,c}(\omega_k + \ell/n)$, $d_{x,s}(\omega_k + \ell/n)$, $d_{y,c}(\omega_k + \ell/n)$, $d_{y,s}(\omega_k + \ell/n)$, sampled at $L = 2m+1$ frequencies, $\ell = -m, \dots, m$, in the neighborhood of some target frequency ω . Suppose these cosine and sine transforms are re-indexed and denoted by $d_{x,cj}$, $d_{x,sj}$, $d_{y,cj}$, $d_{y,sj}$, for $j = 1, 2, \dots, L$, producing $2L$ real random variables with a large sample normal distribution that have limiting covariance matrices of the form (C.47) for each j . Then, the conditional normal distribution of the 2×1 vector $d_{y,cj}$, $d_{y,sj}$ given $d_{x,cj}$, $d_{x,sj}$, given in **Example C.3**, shows that we may write, approximately, the regression model

$$\begin{pmatrix} d_{y,cj} \\ d_{y,sj} \end{pmatrix} = \begin{pmatrix} d_{x,cj} & -d_{x,sj} \\ d_{x,sj} & d_{x,cj} \end{pmatrix} \begin{pmatrix} b_1 \\ b_2 \end{pmatrix} + \begin{pmatrix} V_{cj} \\ V_{sj} \end{pmatrix},$$

where V_{cj}, V_{sj} are approximately uncorrelated with approximate variances

$$E[V_{cj}^2] = E[V_{sj}^2] = (1/2)f_{y \cdot x}.$$

Now, construct, by stacking, the $2L \times 1$ vectors $\mathbf{y}_c = (d_{y,c1}, \dots, d_{y,cL})'$, $\mathbf{y}_s = (d_{y,s1}, \dots, d_{y,sL})'$, $\mathbf{x}_c = (d_{x,c1}, \dots, d_{x,cL})'$ and $\mathbf{x}_s = (d_{x,s1}, \dots, d_{x,sL})'$, and rewrite the regression model as

$$\begin{pmatrix} \mathbf{y}_c \\ \mathbf{y}_s \end{pmatrix} = \begin{pmatrix} \mathbf{x}_c & -\mathbf{x}_s \\ \mathbf{x}_s & \mathbf{x}_c \end{pmatrix} \begin{pmatrix} b_1 \\ b_2 \end{pmatrix} + \begin{pmatrix} \mathbf{v}_c \\ \mathbf{v}_s \end{pmatrix}$$

where $\mathbf{v}_c$ and $\mathbf{v}_s$ are the error stacks. Finally, write the overall model as the regression model in Chapter 2, namely,

$$\mathbf{y} = \mathbf{Z}\mathbf{b} + \mathbf{v},$$

making the obvious identifications in the previous equation. Conditional on $\mathbf{Z}$, the model becomes exactly the regression model considered in Chapter 2 where there are $q = 2$ regression coefficients and $2L$ observations in the observation vector $\mathbf{y}$. To test the hypothesis of no regression for that model, we use an F-Statistic that depends on the difference between the residual sum of squares for the full model, say,

$$SSE = \mathbf{y}'\mathbf{y} - \mathbf{y}'\mathbf{Z}(\mathbf{Z}'\mathbf{Z})^{-1}\mathbf{Z}'\mathbf{y} \quad (\text{C.53})$$

and the residual sum of squares for the reduced model, $SSE_0 = \mathbf{y}'\mathbf{y}$. Then,

$$F_{2,2L-2} = (L-1) \frac{SSE_0 - SSE}{SSE} \quad (\text{C.54})$$

has the F-distribution with 2 and $2L - 2$ degrees of freedom. Also, it follows by substitution for $\mathbf{y}$ that

$$SSE_0 = \mathbf{y}'\mathbf{y} = \mathbf{y}'_c\mathbf{y}_c + \mathbf{y}'_s\mathbf{y}_s = \sum_{j=1}^L (d_{y,cj}^2 + d_{y,sj}^2) = L\widehat{f}_y(\omega),$$

which is just the sample spectrum of the output series. Similarly,

$$\mathbf{Z}'\mathbf{Z} = \begin{pmatrix} L\widehat{f}_x & 0 \\ 0 & L\widehat{f}_x \end{pmatrix}$$

and

$$\begin{aligned} \mathbf{Z}'\mathbf{y} &= \begin{pmatrix} (\mathbf{x}'_c\mathbf{y}_c + \mathbf{x}'_s\mathbf{y}_s) \\ (\mathbf{x}'_c\mathbf{y}_s - \mathbf{x}'_s\mathbf{y}_c) \end{pmatrix} \\ &= \begin{pmatrix} \sum_{j=1}^L (d_{x,cj}d_{y,cj} + d_{x,sj}d_{y,sj}) \\ \sum_{j=1}^L (d_{x,cj}d_{y,sj} - d_{x,sj}d_{y,cj}) \end{pmatrix} \\ &= \begin{pmatrix} L\widehat{C}_{yx} \\ L\widehat{Q}_{yx} \end{pmatrix}. \end{aligned}$$

together imply that

$$\mathbf{y}'\mathbf{Z}(\mathbf{Z}'\mathbf{Z})^{-1}\mathbf{Z}'\mathbf{y} = L |\widehat{f}_{xy}|^2 / \widehat{f}_x.$$

Substituting into (C.54) gives

$$F_{2,2L-2} = (L-1) \frac{|\widehat{f}_{xy}|^2 / \widehat{f}_x}{\left(\widehat{f}_y - |\widehat{f}_{xy}|^2 / \widehat{f}_x\right)},$$

which converts directly into the F-statistic (4.100), using the sample coherence defined in (4.99).

Appendix R

R Supplement

R.1 First Things First

The website for the text is <http://www.stat.pitt.edu/stoffer/tsa3/>. If you do not already have R, point your browser to the Comprehensive R Archive Network (CRAN), <http://cran.r-project.org/> and download and install it. The installation includes help files and some user manuals. You can find helpful tutorials by following CRAN's link to *Contributed Documentation*. If you are new to R/S-PLUS, then *R for Beginners* by Emmanuel Paradis is a great introduction. There is also a lot of advice out there in cyberspace, but some of it will be outdated because R goes through many revisions. Also, beginners will find RStudio (<https://www.rstudio.com/>) helpful for using R. [Advice: If you've never used a command based program, get R studio.]

R.2 `astsa`

There is an R package for the text called `astsa` (*Applied Statistical Time Series Analysis*), which was the name of the software distributed with the first and second editions of this text, and the original version, Shumway (1988). The package can be obtained from CRAN and its mirrors in the usual way. To download and install `astsa`, start R and type

```
install.packages("astsa")
```

You will be asked to choose the closest CRAN mirror to you. As with all packages, you have to load `astsa` before you use it by issuing the command

```
require(astsa)
```

All the data are loaded when the package is loaded. If you create a `.First` function as follows,

```
.First <- function(){require(astsa)}
```

and save the workspace when you quit, `astsa` will be loaded at every start until you change `.First`.

R is not consistent with help files across different operating systems. The best help system is the html help, which can be started issuing the command `help.start()` and then following the *Packages* link to `astsa`. In Windows, to get to html help quickly, type `help(package=astsa)`. Another option is to type `?astsa` and then use the link to the package *Index* at the bottom of the page. A pdf version of the `astsa` manual may be found at CRAN: <http://cran.r-project.org/web/packages/astsa/astsa.pdf>. Further details on R packages may be found in §R.3.2. A useful command to see all the data files available to you, including those loaded with `astsa`, is `data()`

R.3 Getting Started

The best way to use the rest of this appendix is to start up R and enter the example code as it is presented. Also, you can use the results and help files to get a better understanding of how R works (or doesn't work).

The convention throughout the text is that R code is in blue, output is purple and comments are # green. Get comfortable, then start her up and try some simple tasks.

```
2+2      # addition
[1] 5
5*5 + 2  # multiplication and addition
[1] 27
5/5 - 3  # division and subtraction
[1] -2
log(exp(pi)) # log, exponential, pi
[1] 3.141593
sin(pi/2) # sinusoids
[1] 1
exp(1)^(-2) # power
[1] 0.1353353
sqrt(8)   # square root
[1] 2.828427
1:5       # sequences
[1] 1 2 3 4 5
seq(1, 10, by=2) # sequences
[1] 1 3 5 7 9
rep(2,3)   # repeat 2 three times
[1] 2 2 2
```

Next, we'll use *assignment* to make some *objects*:

```
x <- 1 + 2 # put 1 + 2 in object x
x = 1 + 2  # same as above with fewer keystrokes
1 + 2 -> x # same
x          # view object x
[1] 3
(y = 9*3)  # put 9 times 3 in y and view the result
[1] 27
(z = rnorm(5)) # put 5 standard normals into z and print z
[1] 0.96607946 1.98135811 -0.06064527 0.31028473 0.02046853
```

It is worth pointing out R's *recycling rule* for doing arithmetic. In the code below, `c()` [concatenation] is used to create a vector. Note the use of the semicolon for multiple commands on one line.

```
x = c(1, 2, 3, 4); y = c(2, 4, 6, 8); z = c(10, 20); w = c(8, 3, 2)
x*y # 1*2, 2*4, 3*6, 4*8
[1] 2 8 18 32
x+z # 1+10, 2+20, 3+10, 4+20
[1] 11 22 13 24
y+w # what happened here?
[1] 10 7 8 16
Warning message:
In y + w : longer object length is not a multiple of
shorter object length
```

To list your objects, remove objects, get help, find out which directory is current (or to change it) or to quit, use the following commands:

```
ls() # list all objects
"dummy" "mydata" "x" "y" "z"
ls(pattern = "my") # list every object that contains "my"
"dummy" "mydata"
rm(dummy) # remove object "dummy"
rm(list=ls()) # remove almost everything (use with caution)
help.start() # html help and documentation (use it)
data() # list of available data sets
help(exp) # specific help (?exp is the same)
getwd() # get working directory
setwd() # change working directory
q() # end the session (keep reading)
```

When you quit, R will prompt you to save an image of your current workspace. Answering “yes” will save all the work you have done so far, and load it up when you next start R.

To create your own data set, you can make a data vector as follows:

```
mydata = c(1,2,3,2,1)
```

Now you have an object called `mydata` that contains five elements. R calls these objects *vectors* even though they have no dimensions (no rows, no columns); they do have order and length:

```
mydata # display the data
[1] 1 2 3 2 1
mydata[3] # the third element
[1] 3
mydata[3:5] # elements three through five
[1] 3 2 1
mydata[-(1:2)] # everything except the first two elements
[1] 3 2 1
length(mydata) # number of elements
[1] 5
dim(mydata) # no dimensions
NULL
mydata = as.matrix(mydata) # make it a matrix
dim(mydata) # now it has dimensions
[1] 5 1
```

If you have an external data set, you can use `scan` or `read.table` (or some variant) to input the data. For example, suppose you have an ASCII (text) data file called `dummy.txt` in your working directory, and the file looks like this:

1	2	3	2	1
9	0	2	1	0

```
(dummy = scan("dummy.txt"))           # scan and view it
Read 10 items
[1] 1 2 3 2 1 9 0 2 1 0
(dummy = read.table("dummy.txt"))      # read and view it
V1 V2 V3 V4 V5
 1  2  3  2  1
 9  0  2  1  0
```

There is a difference between `scan` and `read.table`. The former produced a data vector of 10 items while the latter produced a *data frame* with names `V1` to `V5` and two observations per variate. In this case, if you want to list (or use) the second variate, `V2`, you would use

```
dummy$V2
[1] 2 0
```

and so on. You might want to look at the help files `?scan` and `?read.table` now. Data frames (`?data.frame`) are “used as the fundamental data structure by most of R’s modeling software.” Notice that R gave the columns of `dummy` generic names, `V1`, ..., `V5`. You can provide your own names and then use the names to access the data without the use of `$` as above.

```
colnames(dummy) = c("Dog", "Cat", "Rat", "Pig", "Man")
attach(dummy)
Cat
[1] 2 0
Rat*(Pig - Man) # animal arithmetic
[1] 3 2
detach(dummy)   # clean up (if desired)
```

R is case sensitive, thus `cat` and `Cat` are different. Also, `cat` is a reserved name (`?cat`) in R, so using `"cat"` instead of `"Cat"` may cause problems later. You may also include a *header* in the data file to avoid using line 5 above. For example, if the file `dummy.txt` looks like this,

Dog	Cat	Rat	Pig	Man
1	2	3	2	1
9	0	2	1	0

then use the following command to read the data.

```
(dummy = read.table("dummy.txt", header=TRUE))
  Dog Cat Rat Pig Man
1   1   2   3   2   1
2   9   0   2   1   0
```

Another popular format for data files is `.csv` (comma separated values). In this case, use `read.csv` instead of `read.table`. The default for `.csv` files is `header=TRUE`; type `?read.table` for further information.

Two commands that are used frequently to manipulate data are `cbind` for *column binding* and `rbind` for *row binding*. The following is an example.

```

x = runif(4) # generate 4 values from a uniform(0,1) into object x
y = runif(4) # generate 4 more and put them into object y
(u = cbind(x,y)) # column bind the two vectors (4 by 2 matrix)
      x      y
[1,] 0.6547304 0.7503984
[2,] 0.8222048 0.1335557
[3,] 0.4555755 0.2151735
[4,] 0.9843289 0.8483795
(u = rbind(x,y)) # row bind the two vectors (2 by 4 matrix)
      [,1]      [,2]      [,3]      [,4]
x  0.6547304 0.8222048 0.4555755 0.9843289
y  0.7503984 0.1335557 0.2151735 0.8483795

```

R.3.1 Basic Statistics

Summary statistics are fairly easy to obtain. We will simulate 25 normals with $\mu = 10$ and $\sigma = 4$ and then perform some basic analyses.

```

set.seed(90210) # so you can reproduce these results
x = rnorm(25, 10, 4) # generate the data
mean(x)
[1] 9.473883
median(x) # compute the median
[1] 9.448511
var(x) # compute the variance
[1] 13.9267
sd(x) # compute the standard deviation
[1] 3.73185
max(x) # find the largest value
[1] 17.32609
which.max(x) # index of the max (x[25] in this case)
[1] 25
min(x) # find the smallest value
[1] 2.678173
summary(x) # a five number summary (with the mean added)
      Min. 1st Qu. Median      Mean      3rd Qu.      Max.
2.678   7.824   9.449   9.474   11.180   17.330
par(mfrow=c(1,2)) # multifigure setup (?par for info)
hist(x, col="lightblue", prob=TRUE) # see figure below
lines(density(x, bw=2)) # fit a density over it
boxplot(x, main="Boxplot of x", col="lightblue") # see figure below

```

We introduced some R graphics without saying much about it. Aside from the manual, there are many useful websites that expand on this topic; e.g., <https://www.stat.auckland.ac.nz/~paul/RGraphics/rgraphics.html>.

It can't hurt to learn a little about programming in R because you will see some of it along the way. Consider a simple program that we will call **crazy** to produce a graph of a sequence of sample means of increasing sample sizes from a Cauchy distribution with location parameter zero. The code is:

```

crazy <- function(num) {
  x <- rep(NA, num)
  for (n in 1:num) x[n] <- mean(rcauchy(n))
  plot(x, type="l", xlab="sample size", ylab="sample mean")
}

```




Fig. R.1. Histogram and boxplot of simulated normal data.




Fig. R.2. Crazy example.

The first line creates the function `crazy` and gives it one argument, `num`, that is the sample size that will end the sequence. Line 2 makes a vector, `x`, of `num` missing values `NA`, that will be used to store the sample means. Line 3 generates `n` random Cauchy variates [`rcauchy(n)`], finds the mean of those values, and puts the result into `x[n]`, the n -th value of `x`. The process is repeated in a “do loop” `num` times so that `x[1]` is the sample mean from a sample of size one, `x[2]` is the sample mean from a sample of size two, and so on, until finally, `x[num]` is the sample mean from a sample of size `num`. After the do loop is complete, the fourth line generates a graphic (see [Figure R.2](#)). The fifth line closes the function. To use `crazy` with a limit sample size of 100, for example, type

```
crazy(100)
```

and you will get a graphic that looks like [Figure R.2](#)

R.3.2 Packages

We have already discussed how to obtain and load the R package for this text, `astsa`. You may want to use one of the many other R packages; we have used a few throughout the text. In this case, you have to first download the package and then install it. For example,


```
install.packages(c("dynlm", "tseries"))
```

will download and install the packages `dynlm` that we use in Chapter 2 and `tseries` that we use in Chapter 5; you will be asked to choose the closest mirror to you. To use a package, you have to load it at each start up of R, for example:

```
require(dynlm) # load the package
```

A good way to get help for a package is to use html help

```
help.start()
```

and follow the *Packages* link.

To make sure you have the current version of your R packages, periodically update them by issuing the command

```
update.packages(ask=FALSE)
```

You will be asked to choose the nearest repository. Using "`ask=FALSE`" means you won't be asked about updating each package. If you want a list of your packages, run `installed.packages()[,1:3]`

giving the first 3 columns of info: the [1] package name, [2] library path, and [3] version number.

R.3.3 Word

Finally, a word of caution: `TRUE` and `FALSE` are reserved words, whereas `T` and `F` are initially set to these. Get in the habit of using the words rather than the letters `T` or `F` because you may get into trouble if you do something like

```
F = qf(p=.01, df1=3, df2=9)
```

so that `F` is no longer `FALSE`, but a quantile of the specified F -distribution. Better yet, just remember not to use `T` or `F` for anything.

R.4 Time Series Primer

In this section, we give a brief introduction on using R for time series. **We assume that `astsa` has been loaded.** To create a time series object, use the command `ts`. Related commands are `as.ts` to coerce an object to a time series and `is.ts` to test whether an object is a time series.

First, make a small data set:

```
(mydata = c(1,2,3,2,1)) # make it and view it
[1] 1 2 3 2 1
```

Now make it a time series:

```
(mydata = as.ts(mydata))
Time Series:
Start = 1
End = 5
Frequency = 1
[1] 1 2 3 2 1
```

Make it an annual time series that starts in 1950:

```
(mydata = ts(mydata, start=1950))
Time Series:
Start = 1950
End = 1954
Frequency = 1
[1] 1 2 3 2 1
```

Now make it a quarterly time series that starts in 1950-III:

```
(mydata = ts(mydata, start=c(1950,3), frequency=4))
Qtr1 Qtr2 Qtr3 Qtr4
1950          1    2
1951    3    2    1
time(mydata) # view the sampled times
Qtr1    Qtr2    Qtr3    Qtr4
1950          1950.50 1950.75
1951 1951.00 1951.25 1951.50
```

To use part of a time series object, use `window()`:

```
(x = window(mydata, start=c(1951,1), end=c(1951,3)))
Qtr1 Qtr2 Qtr3
1951    3    2    1
```

Next, we'll look at lagging and differencing. First make a simple series, x_t :

```
x = ts(1:5)
```

Now, column bind (`cbind`) lagged values of x_t and you will notice that `lag(x)` is *forward* lag, whereas `lag(x, -1)` is *backward* lag.

```
cbind(x, lag(x), lag(x,-1))
      x lag(x) lag(x, -1)
0    NA      1      NA
1     1      2      NA
2     2      3       1
3     3      4       2 <- in this row, for example, x is 3,
4     4      5       3   lag(x) is ahead at 4, and
5     5     NA       4   lag(x,-1) is behind at 2
6    NA     NA       5
```

Compare `cbind` and `ts.intersect`:

```
ts.intersect(x, lag(x,1), lag(x,-1))
Time Series: Start = 2 End = 4 Frequency = 1
      x lag(x, 1) lag(x, -1)
2     2          3          1
3     3          4          2
4     4          5          3
```

To difference a series, $\nabla x_t = x_t - x_{t-1}$, use

```
diff(x)
```

but note that

```
diff(x, 2)
```

is *not* second order differencing, it is $x_t - x_{t-2}$. For second order differencing, that is, $\nabla^2 x_t$, do one of these:

```
diff(diff(x))
diff(x, diff=2) # same thing
```

and so on for higher order differencing.




Fig. R.3. Demonstration of different R graphic tools for plotting time series.

For graphing time series, there are a few standard plotting mechanisms that we use repeatedly. If `x` is a time series, then `plot(x)` will produce a time plot. If `x` is not a time series object, then `plot.ts(x)` will coerce it into a time plot as will `ts.plot(x)`. There are differences, which we explore in the following. It would be a good idea to skim the graphical parameters help file (`?par`) while you are here. See [Figure R.3](#) for the resulting graphic.

```
x = -5:5      # x is NOT a time series object
y = 5*cos(x)  # neither is y
par(mfrow=c(3,2)) # multifigure setup: 3 rows, 2 cols
plot(x, main="plot(x)")
plot(x, y, main="plot(x,y)")
plot.ts(x, main="plot.ts(x)")
plot.ts(x, y, main="plot.ts(x,y)")
ts.plot(x, main="ts.plot(x)")
ts.plot(ts(x), ts(y), col=1:2, main="ts.plot(x,y)")
```

We will also make use of regression via `lm()`. First, suppose we want to fit a simple linear regression, $y = \alpha + \beta x + \epsilon$. In R, the formula is written as `y~x`:

```
set.seed(1999)      # so you can reproduce the result
x = rnorm(10,0,1)
y = x + rnorm(10,0,1)
summary(fit <- lm(y~x))
```

```

Residuals:
    Min       1Q   Median       3Q      Max
-0.8851 -0.3867  0.1325  0.3896  0.6561

Coefficients:
              Estimate Std. Error t value Pr(>t)
(Intercept)   0.2576     0.1892   1.362  0.2104
x              0.4577     0.2016   2.270  0.0529

Residual standard error: 0.58 on 8 degrees of freedom
Multiple R-squared: 0.3918,    Adjusted R-squared: 0.3157
F-statistic: 5.153 on 1 and 8 DF,  p-value: 0.05289
plot(x, y)    # draw a scatterplot of the data (not shown)
abline(fit)   # add the fitted line to the plot (not shown)

```

All sorts of information can be extracted from the `lm` object, which we called `fit`. For example,

```

resid(fit)     # will display the residuals (not shown)
fitted(fit)    # will display the fitted values (not shown)
lm(y ~ 0 + x)  # will exclude the intercept (not shown)

```

You have to be careful if you use `lm()` for lagged values of a time series. If you use `lm()`, then what you have to do is “tie” the series together using `ts.intersect`. If you do not tie the series together, they will not be aligned properly. Please read the warning *Using time series* in the `lm()` help file [`help(lm)`]. Here is an example regressing `astsa` data, weekly cardiovascular mortality (`cmort`) on particulate pollution (`part`) at the present value and lagged four weeks (`part4`). First, we create `ded`, which consists of the intersection of the three series:

```
ded = ts.intersect(cmort, part, part4=lag(part,-4))
```

Now the series are all aligned and the regression will work.

```

fit = lm(cmort~part+part4, data=ded, na.action=NULL)
summary(fit)
Coefficients:
              Estimate Std. Error t value
(Intercept)  69.01020     1.37498  50.190
part          0.15140     0.02898   5.225
part4         0.26297     0.02899   9.071

```

```

Residual standard error: 8.323 on 501 degrees of freedom
Multiple R-Squared: 0.3091,    Adjusted R-squared: 0.3063
F-statistic: 112.1 on 2 and 501 DF,  p-value: < 2.2e-16

```

There was no need to rename `lag(part, -4)` to `part4`, it’s just an example of what you can do. There is a package called `dynlm` that makes it easy to fit lagged regressions. The basic advantage of `dynlm` is that it avoids having to make a data frame; that is, line 2 would be avoided.

In **Problem 2.1**, you are asked to fit a regression model

$$x_t = \beta t + \alpha_1 Q_1(t) + \alpha_2 Q_2(t) + \alpha_3 Q_3(t) + \alpha_4 Q_4(t) + w_t$$

where x_t is logged Johnson & Johnson quarterly earnings ($n = 84$), and $Q_i(t)$ is the indicator of quarter $i = 1, 2, 3, 4$. The indicators can be made using `factor`.

```
trend = time(jj) - 1970 # helps to 'center' time
Q = factor(cycle(jj))   # make (Q)uarter factors
reg = lm(log(jj)~0 + trend + Q, na.action=NULL) # no intercept
model.matrix(reg) # view the model matrix
```

	trend	Q1	Q2	Q3	Q4
1	-10.00	1	0	0	0
2	-9.75	0	1	0	0
3	-9.50	0	0	1	0
4	-9.25	0	0	0	1
.	.	.	.	.	.
.	.	.	.	.	.
83	10.50	0	0	1	0
84	10.75	0	0	0	1

```
summary(reg) # view the results (not shown)
```

The workhorse for ARIMA simulations is `arima.sim`. Here are some examples; no output is shown here so you're on your own.

```
x = arima.sim(list(order=c(1,0,0),ar=.9),n=100)+50 # AR(1) w/mean 50
x = arima.sim(list(order=c(2,0,0),ar=c(1,-.9)),n=100) # AR(2)
x = arima.sim(list(order=c(1,1,1),ar=.9,ma=-.5),n=200) # ARIMA(1,1,1)
```

Next, we'll discuss ARIMA estimation. This gets a bit tricky because R is not user friendly when it comes to fitting ARIMA models. Much of the story is spelled out in the "R Issues" page of the website for the text. In Chapter 3, we use the scripts `acf2`, `sarima`, and `sarima.for` that are included with `astsa`. But we will also show you how to use the scripts included with R.

First, we'll fit an ARMA(1,1) model to some simulated data (with diagnostics and forecasting):

```
set.seed(666)
x = 50 + arima.sim(list(order=c(1,0,1), ar=.9, ma=-.5), n=200)
acf(x); pacf(x) # display sample ACF and PACF ... or ...
acf2(x) # use our script (no output shown)
(x.fit = arima(x, order = c(1, 0, 1))) # fit the model
Call: arima(x = x, order = c(1, 0, 1))
Coefficients:
      ar1      mal  intercept
    0.8340  -0.432   49.8960
s.e.  0.0645   0.111   0.2452
sigma^2 estimated as 1.070: log likelihood = -290.79, aic = 589.58
```

Note that the reported `intercept` estimate is an estimate of the mean and *not* the constant. That is, the fitted model is

$$\hat{x}_t - 49.896 = .834(x_{t-1} - 49.896) + \hat{w}_t$$

where $\hat{\sigma}_w^2 = 1.070$. **Incorrect** diagnostics can be accomplished as follows:

```
tsdiag(x.fit, gof.lag=20) # !!!! don't use this !!!!
```

That's right, the Ljung-Box-Pierce test is **NOT** correct because it does not take into account the fact that the residuals are from a fitted model. If the analysis is repeated using the `sarima` script, a partial output would look like the following (`sarima` will also display the correct diagnostics as a graphic; e.g., see [Figure 3.19](#) and [Figure R.4](#)):

```
sarima(x, 1, 0, 1)
Coefficients:
      ar1      ma1      xmean
      0.8340  -0.432  49.8960
s.e.    0.0645   0.111   0.2452
sigma^2 estimated as 1.070: log likelihood = -290.79, aic = 589.58
$AIC [1] 1.097494  $AICc [1] 1.108519  $BIC [1] 0.1469684
```




Fig. R.4. Diagnostics from `sarima`

Forecasting using the script `sarima.for` can be accomplished in one line:

```
sarima.for(x, 10, 1, 0, 1)
```

If you don't use `astsa`, you would obtain and plot the forecasts using the following:

```
x.fore = predict(x.fit, n.ahead=10)
U = x.fore$pred + 2*x.fore$se # x.fore$pred holds predicted values
L = x.fore$pred - 2*x.fore$se # x.fore$se holds stdn errors
miny = min(x,L); maxy = max(x,U)
ts.plot(x, x.fore$pred, col=1:2, ylim=c(miny, maxy))
lines(U, col="blue", lty="dashed")
lines(L, col="blue", lty="dashed")
```

We close this appendix with a quick spectral analysis. This material is covered in detail in Chapter 4, so we will not discuss this example in much detail here. We will simulate an AR(2) and then estimate the spectrum via nonparametric and parametric methods. No graphics are shown, but we have confidence that you are proficient enough in R to display them yourself.

```

x = arima.sim(list(order=c(2,0,0), ar=c(1,-.9)), n=2^8) # some data
(u = polyroot(c(1,-1,.9))) # x is AR(2) w/complex roots
[1] 0.5555556+0.8958064i 0.5555556-0.8958064i
Arg(u[1])/(2*pi) # dominant frequency around .16
[1] 0.1616497
par(mfcol=c(2,2))
plot.ts(x)
mvspec(x, spans=c(5,5), taper=.1, log="no") # nonparametric estimate
spec.ar(x, log="no") # parametric spectral estimate
arma.spec(ar=c(1,-.9), log="no") # true spectral density

```

The script `arma.spec` is included in `astsa`. See `spectrum` and `spec.pgram` as alternatives to `mvspec`, which is part of `astsa`. Finally, note the easiest way to get a raw periodogram is:

```
per = abs(fft(x))^2/length(x) # abs() and Mod() are same here
```

References

- Akaike, H. (1969). Fitting autoregressive models for prediction. *Ann. Inst. Stat. Math.*, 21, 243-247.
- Akaike, H. (1973). Information theory and an extension of the maximum likelihood principal. In *2nd Int. Symp. Inform. Theory*, 267-281. B.N. Petrov and F. Csake, eds. Budapest: Akademia Kiado.
- Akaike, H. (1974). A new look at statistical model identification. *IEEE Trans. Automat. Contr.*, AC-19, 716-723.
- Alagón, J. (1989). Spectral discrimination for two groups of time series. *J. Time Series Anal.*, 10, 203-214.
- Alspach, D.L. and H.W. Sorensen (1972). Nonlinear Bayesian estimation using Gaussian sum approximations. *IEEE Trans. Automat. Contr.*, AC-17, 439-447.
- Anderson, B.D.O. and J.B. Moore (1979). *Optimal Filtering*. Englewood Cliffs, NJ: Prentice-Hall.
- Anderson, T.W. (1978). Estimation for autoregressive moving average models in the time and frequency domain. *Ann. Stat.*, 5, 842-865.
- Anderson, T.W. (1984). *An Introduction to Multivariate Statistical Analysis*, 2nd ed. New York: Wiley.
- Ansley, C.F. and P. Newbold (1980). Finite sample properties of estimators for autoregressive moving average processes. *J. Econ.*, 13, 159-183.
- Ansley, C.F. and R. Kohn (1982). A geometrical derivation of the fixed interval smoothing algorithm. *Biometrika*, 69, 486-487.
- Antognini, J.F., M.H. Buonocore, E.A. Disbrow, and E. Carstens (1997). Isoflurane anesthesia blunts cerebral responses to noxious and innocuous stimuli: a fMRI study. *Life Sci.*, 61, PL349-PL354.
- Bandettini, A., A. Jesmanowicz, E.C. Wong, and J.S. Hyde (1993). Processing strategies for time-course data sets in functional MRI of the human brain. *Magnetic Res. Med.*, 30, 161-173.
- Bar-Shalom, Y. (1978). Tracking methods in a multi-target environment. *IEEE Trans. Automat. Contr.*, AC-23, 618-626.

- Bar-Shalom, Y. and E. Tse (1975). Tracking in a cluttered environment with probabilistic data association. *Automatica*, 11, 4451-4460.
- Bazza, M., R.H. Shumway, and D.R. Nielsen (1988). Two-dimensional spectral analysis of soil surface temperatures. *Hilgardia*, 56, 1-28.
- Bedrick, E.J. and C.-L. Tsai (1994). Model selection for multivariate regression in small samples. *Biometrics*, 50, 226-231.
- Benjamini, Y. and Y. Hochberg (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J. Roy. Statist. Soc. Ser. B*, 289-300.
- Beran, J. (1994). *Statistics for Long Memory Processes*. New York: Chapman and Hall.
- Berk, K.N. (1974). Consistent autoregressive spectral estimates. *Ann. Stat.*, 2, 489-502.
- Besag, J. (1974). Spatial interaction and the statistical analysis of lattice systems (with discussion). *J. R. Stat. Soc. B*, 36, 192-236.
- Bhat, R.R. (1985). *Modern Probability Theory*, 2nd ed. New York: Wiley.
- Bhattacharya, A. (1943). On a measure of divergence between two statistical populations. *Bull. Calcutta Math. Soc.*, 35, 99-109.
- Billingsley, P. (1999). *Convergence of Probability Measures*, (2nd edition). New York: Wiley.
- Blackman, R.B. and J.W. Tukey (1959). *The Measurement of Power Spectra from the Point of View of Communications Engineering*. New York: Dover.
- Blight, B.J.N. (1974). Recursive solutions for the estimation of a stochastic parameter *J. Am. Stat. Assoc.*, 69, 477-481
- Bloomfield, P. (1976). *Fourier Analysis of Time Series: An Introduction*. New York: Wiley.
- Bloomfield, P. (2000). *Fourier Analysis of Time Series: An Introduction*, 2nd ed. New York: Wiley.
- Bloomfield, P. and J.M. Davis (1994). Orthogonal rotation of complex principal components. *Int. J. Climatol.*, 14, 759-775.
- Bogart, B. P., M. J. R. Healy, and J.W. Tukey (1962). The Quefrency Analysis of Time Series for Echoes: Cepstrum, Pseudo-Autocovariance, Cross-Cepstrum and Saphe Cracking. In *Proc. of the Symposium on Time Series Analysis*, pp. 209-243, Brown University, Providence, USA.
- Bollerslev, T. (1986). Generalized autoregressive conditional heteroscedasticity. *J. Econ.*, 31, 307- 327.
- Box, G.E.P. and D.A. Pierce (1970). Distributions of residual autocorrelations in autoregressive integrated moving average models. *J. Am. Stat. Assoc.*, 72, 397-402.
- Box, G.E.P. and G.M. Jenkins (1970). *Time Series Analysis, Forecasting, and Control*. Oakland, CA: Holden-Day.
- Box, G.E.P. and G.C. Tiao (1973). *Bayesian Inference in Statistical Analysis*. New York: Wiley.
- Box, G.E.P., G.M. Jenkins and G.C. Reinsel (1994). *Time Series Analysis, Forecasting, and Control*, 3rd ed. Englewood Cliffs, NJ: Prentice Hall.
- Breiman, L. and J. Friedman (1985). Estimating optimal transformations for multiple regression and correlation (with discussion). *J. Am. Stat. Assoc.*, 80, 580-619.
- Brillinger, D.R. (1973). The analysis of time series collected in an experimental design. In *Multivariate Analysis-III.*, pp. 241-256. P.R. Krishnaiah ed. New York: Academic Press.

- Brillinger, D.R. (1975). *Time Series: Data Analysis and Theory*. New York: Holt, Rinehart & Winston Inc.
- Brillinger, D.R. (1980). Analysis of variance and problems under time series models. In *Handbook of Statistics*, Vol I, pp. 237-278. P.R. Krishnaiah and D.R. Brillinger, eds. Amsterdam: North Holland.
- Brillinger, D.R. (1981, 2001). *Time Series: Data Analysis and Theory*, 2nd ed. San Francisco: Holden-Day. Republished in 2001 by the Society for Industrial and Applied Mathematics, Philadelphia.
- Brockwell, P.J. and R.A. Davis (1991). *Time Series: Theory and Methods*, 2nd ed. New York: Springer-Verlag.
- Bruce, A. and H-Y. Gao (1996). *Applied Wavelet Analysis with S-PLUS*. New York: Springer-Verlag.
- Caines, P.E. (1988). *Linear Stochastic Systems*. New York: Wiley.
- Carlin, B.P., N.G. Polson, and D.S. Stoffer (1992). A Monte Carlo approach to nonnormal and nonlinear state-space modeling. *J. Am. Stat. Assoc.*, 87, 493-500.
- Carter, C. K. and R. Kohn (1994). On Gibbs sampling for state space models. *Biometrika*, 81, 541-553.
- Chan, N.H. (2002). *Time Series: Applications to Finance*. New York: Wiley.
- Chernoff, H. (1952). A measure of asymptotic efficiency for tests of a hypothesis based on the sum of the observations. *Ann. Math. Stat.*, 25, 573-578.
- Cleveland, W.S. (1979). Robust locally weighted regression and smoothing scatterplots. *J. Am. Stat. Assoc.*, 74, 829-836.
- Cochrane, D. and G.H. Orcutt (1949). Applications of least squares regression to relationships containing autocorrelated errors. *J. Am. Stat. Assoc.*, 44, 32-61.
- Cooley, J.W. and J.W. Tukey (1965). An algorithm for the machine computation of complex Fourier series. *Math. Comput.*, 19, 297-301.
- Cressie, N.A.C. (1993). *Statistics for Spatial Data*. New York: Wiley.
- Dahlhaus, R. (1989). Efficient parameter estimation for self-similar processes. *Ann. Stat.*, 17, 1749-1766.
- Dargahi-Noubary, G.R. and P.J. Laycock (1981). Spectral ratio discriminants and information theory. *J. Time Series Anal.*, 16, 201-219.
- Danielson, J. (1994). Stochastic volatility in asset prices: Estimation with simulated maximum likelihood. *J. Econometrics*, 61, 375-400.
- Daubechies, I. (1992). *Ten Lectures on Wavelets*. Philadelphia: CBMS-NSF Regional Conference Series in Applied Mathematics.
- Davies, N., C.M. Triggs, and P. Newbold (1977). Significance levels of the Box-Pierce portmanteau statistic in finite samples. *Biometrika*, 64, 517-522.
- Dent, W. and A.-S. Min. (1978). A Monte Carlo study of autoregressive-integrated-moving average processes. *J. Econ.*, 7, 23-55.
- Dempster, A.P., N.M. Laird and D.B. Rubin (1977). Maximum likelihood from incomplete data via the EM algorithm. *J. R. Stat. Soc. B*, 39, 1-38.
- Ding, Z., C.W.J. Granger, and R.F. Engle (1993). A long memory property of stock market returns and a new model. *J. Empirical Finance*, 1, 83-106.

- Donoho, D.L. and I.M. Johnstone (1994). Ideal spatial adaptation by wavelet shrinkage. *Biometrika*, 81, 425-455.
- Donoho, D.L. and I.M. Johnstone (1995). Adapting to unknown smoothness via wavelet shrinkage. *J. of Am. Stat. Assoc.*, 90, 1200-1224.
- Douc, R., E. Moulines, and D.S. Stoffer (2014). *Nonlinear Time Series: Theory, Methods, and Applications with R Examples*. Boca Raton: CRC Press.
- Durbin, J. (1960). Estimation of parameters in time series regression models. *J. R. Stat. Soc. B*, 22, 139-153.
- Durbin, J. and S.J. Koopman (2001). *Time Series Analysis by State Space Methods* Oxford: Oxford University Press.
- Efron, B. and R. Tibshirani (1994). *An Introduction to the Bootstrap*. New York: Chapman and Hall.
- Engle, R.F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 50, 987-1007.
- Engle, R.F., D. Nelson, and T. Bollerslev (1994). ARCH Models. In *Handbook of Econometrics*, Vol IV, pp. 2959-3038. R. Engle and D. McFadden, eds. Amsterdam: North Holland.
- Fox, R. and M.S. Taquq (1986). Large sample properties of parameter estimates for strongly dependent stationary Gaussian time series. *Ann. Stat.*, 14, 517-532.
- Friedman, J.H. (1984). A Variable Span Smoother. Tech. Rep. No. 5, Lab. for Computational Statistics, Dept. Statistics, Stanford Univ., California.
- Friedman, J.H. and W. Stuetzle. (1981). Projection pursuit regression. *J. Am. Stat. Assoc.*, 76, 817-823.
- Frühwirth-Schnatter, S. (1994). Data Augmentation and Dynamic Linear Models. *J. Time Series Anal.*, 15, 183-202.
- Fuller, W.A. (1976). *Introduction to Statistical Time Series*. New York: Wiley.
- Fuller, W.A. (1996). *Introduction to Statistical Time Series*, 2nd ed. New York: Wiley.
- Gelfand, A.E. and A.F.M. Smith (1990). Sampling-based approaches to calculating marginal densities. *J. Am. Stat. Assoc.*, 85, 398-409.
- Gelman, A., J. Carlin, H. Stern, and D. Rubin (1995). *Bayesian Data Analysis*. London: Chapman and Hall.
- Geman, S. and D. Geman (1984). Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Trans. Pattern Anal. Machine Intell.*, 6, 721-741.
- Geweke, J.F. (1977). The dynamic factor analysis of economic time series models. In *Latent Variables in Socio-Economic Models*, pp 365-383. D. Aigner and A. Goldberger, eds. Amsterdam: North Holland.
- Geweke, J.F. and K.J. Singleton (1981). Latent variable models for time series: A frequency domain approach with an application to the Permanent Income Hypothesis. *J. Econ.*, 17, 287-304.
- Geweke, J.F. and S. Porter-Hudak (1983). The estimation and application of long-memory time series models. *J. Time Series Anal.*, 4, 221-238.
- Gilks, W.R., S. Richardson, and D.J. Spiegelhalter (eds.) (1996). *Markov Chain Monte Carlo in Practice*. London: Chapman and Hall.
- Giri, N. (1965). On complex analogues of T^2 and R^2 tests. *Ann. Math. Stat.*, 36, 664-670.

- Goldfeld, S.M. and R.E. Quandt (1973). A Markov model for switching regressions. *J. Econ.*, 1, 3-16.
- Goodman, N.R. (1963). Statistical analysis based on a certain multivariate complex Gaussian distribution. *Ann. Math. Stat.*, 34, 152-177.
- Gordon, K. and A.F.M. Smith (1988). Modeling and monitoring discontinuous changes in time series. In *Bayesian Analysis of Time Series and Dynamic Models*, 359-392.
- Gordon, K. and A.F.M. Smith (1990). Modeling and monitoring biomedical time series. *J. Am. Stat. Assoc.*, 85, 328-337.
- Gouriéroux, C. (1997). *ARCH Models and Financial Applications*. New York: Springer-Verlag.
- Granger, C.W. and R. Joyeux (1980). An introduction to long-memory time series models and fractional differencing. *J. Time Series Anal.*, 1, 15-29.
- Grenander, U. (1951). On empirical spectral analysis of stochastic processes. *Arkiv for Matematik*, 1, 503-531.
- Grenander, U. and M. Rosenblatt (1957). *Statistical Analysis of Stationary Time Series*. New York: Wiley.
- Grether, D.M. and M. Nerlove (1970). Some properties of optimal seasonal adjustment. *Econometrica*, 38, 682-703.
- Gupta, N.K. and R.K. Mehra (1974). Computational aspects of maximum likelihood estimation and reduction in sensitivity function calculations. *IEEE Trans. Automat. Contr.*, AC-19, 774-783.
- Hamilton, J.D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57, 357-384.
- Hannan, E.J. (1970). *Multiple Time Series*. New York: Wiley.
- Hannan, E. J. and B. G. Quinn (1979). The determination of the order of an autoregression. *J. Royal Statistical Society, B*, 41, 190-195.
- Hannan, E.J. and M. Deistler (1988). *The Statistical Theory of Linear Systems*. New York: Wiley.
- Hansen, J., M. Sato, R. Ruedy, K. Lo, D.W. Lea, and M. Medina-Elizade (2006). Global temperature change. *Proc. Natl. Acad. Sci.*, 103, 14288-14293.
- Harrison, P.J. and C.F. Stevens (1976). Bayesian forecasting (with discussion). *J. R. Stat. Soc. B*, 38, 205-247.
- Harvey, A.C. and P.H.J. Todd (1983). Forecasting economic time series with structural and Box-Jenkins models: A case study. *J. Bus. Econ. Stat.*, 1, 299-307.
- Harvey, A.C. and R.G. Pierse (1984). Estimating missing observations in economic time series. *J. Am. Stat. Assoc.*, 79, 125-131.
- Harvey, A.C. (1991). *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge: Cambridge University Press.
- Harvey, A.C. (1993). *Time Series Models*. Cambridge, MA: MIT Press.
- Harvey A.C., E. Ruiz and N. Shephard (1994). Multivariate stochastic volatility models. *Rev. Economic Studies*, 61, 247-264.
- Haslett, J. and A.E. Raftery (1989) Space-time modelling with long-memory dependence: Assessing Ireland's wind power resource (C/R: 89V38 p21-50) *Applied Statistics*, 38, 1-21

- Hastings, W.K. (1970). Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57, 97-109.
- Hosking, J.R.M. (1981). Fractional differencing. *Biometrika*, 68, 165-176.
- Hurst, H. (1951). Long term storage capacity of reservoirs. *Trans. Am. Soc. Civil Eng.*, 116, 778-808.
- Hurvich, C.M. and S. Zeger (1987). Frequency domain bootstrap methods for time series. *Tech. Report 87-115*, Department of Statistics and Operations Research, Stern School of Business, New York University.
- Hurvich, C.M and C.-L. Tsai (1989). Regression and time series model selection in small samples. *Biometrika*, 76, 297-307.
- Hurvich, C.M. and K.I. Beltrao (1993). Asymptotics for the low-frequency ordinates of the periodogram for a long-memory time series. *J. Time Series Anal.*, 14, 455-472.
- Hurvich, C.M., R.S. Deo and J. Brodsky (1998). The mean squared error of Geweke and Porter-Hudak's estimator of the memory parameter of a long-memory time series. *J. Time Series Anal.*, 19, 19-46.
- Hurvich, C.M. and R.S. Deo (1999). Plug-in selection of the number of frequencies in regression estimates of the memory parameter of a long-memory time series. *J. Time Series Anal.*, 20, 331-341.
- Jacquier, E., N.G. Polson, and P.E. Rossi (1994). Bayesian analysis of stochastic volatility models. *J. Bus. Econ. Stat.*, 12, 371-417.
- Jazwinski, A.H. (1970). *Stochastic Processes and Filtering Theory*. New York: Academic Press.
- Jenkins, G.M. and D.G. Watts. (1968). *Spectral Analysis and Its Applications*. San Francisco: Holden-Day.
- Johnson, R.A. and D.W. Wichern (1992). *Applied Multivariate Statistical Analysis*, 3rd ed.. Englewood Cliffs, NJ: Prentice-Hall.
- Jones, P.D. (1994). Hemispheric surface air temperature variations: A reanalysis and an update to 1993. *J. Clim.*, 7, 1794-1802.
- Jones, R.H. (1980). Maximum likelihood fitting of ARMA models to time series with missing observations. *Technometrics*, 22, 389-395.
- Jones, R.H. (1984). Fitting multivariate models to unequally spaced data. In *Time Series Analysis of Irregularly Observed Data*, pp. 158-188. E. Parzen, ed. Lecture Notes in Statistics, 25, New York: Springer-Verlag.
- Jones, R.H. (1993). *Longitudinal Data With Serial Correlation : A State-Space Approach*. London: Chapman and Hall.
- Journel, A.G. and C.H. Huijbregts (1978). *Mining Geostatistics*. New York: Academic Press.
- Juang, B.H. and L.R. Rabiner (1985). Mixture autoregressive hidden Markov models for speech signals, *IEEE Trans. Acoust., Speech, Signal Process.*, ASSP-33, 1404-1413.
- Kakizawa, Y., R. H. Shumway, and M. Taniguchi (1998). Discrimination and clustering for multivariate time series. *J. Am. Stat. Assoc.*, 93, 328-340.
- Kalman, R.E. (1960). A new approach to linear filtering and prediction problems. *Trans ASME J. Basic Eng.*, 82, 35-45.
- Kalman, R.E. and R.S. Bucy (1961). New results in filtering and prediction theory. *Trans. ASME J. Basic Eng.*, 83, 95-108.

- Kaufman, L. and P.J. Rousseeuw (1990). *Finding Groups in Data: An Introduction to Cluster Analysis*. New York: Wiley.
- Kay, S.M. (1988). *Modern Spectral Analysis: Theory and Applications*. Englewood Cliffs, NJ: Prentice-Hall.
- Kazakos, D. and P. Papantoni-Kazakos (1980). Spectral distance measuring between Gaussian processes. *IEEE Trans. Automat. Contr.*, AC-25, 950-959.
- Khatri, C.G. (1965). Classical statistical analysis based on a certain multivariate complex Gaussian distribution. *Ann. Math. Stat.*, 36, 115-119.
- Kim S., N. Shephard and S. Chib (1998). Stochastic volatility: likelihood inference and comparison with ARCH models. *Rev. Economic Studies*, 65, p.361-393.
- Kitagawa, G. and W. Gersch (1984). A smoothness priors modeling of time series with trend and seasonality. *J. Am. Stat. Assoc.*, 79, 378-389.
- Kitagawa, G. (1987). Non-Gaussian state-space modeling of nonstationary time series (with discussion). *J. Am. Stat. Assoc.*, 82, 1032-1041, (C/R: p1041-1063; C/R: V83 p1231).
- Kitagawa, G. and W. Gersch (1996). *Smoothness Priors Analysis of Time Series*. New York: Springer-Verlag.
- Kolmogorov, A.N. (1941). Interpolation und extrapolation von stationären zufälligen Folgen. *Bull. Acad. Sci. U.R.S.S.*, 5, 3-14.
- Krishnaiah, P.R., J.C. Lee, and T.C. Chang (1976). The distribution of likelihood ratio statistics for tests of certain covariance structures of complex multivariate normal populations. *Biometrika*, 63, 543-549.
- Kullback, S. and R.A. Leibler (1951). On information and sufficiency. *Ann. Math. Stat.*, 22, 79-86.
- Kullback, S. (1958). *Information Theory and Statistics*. Gloucester, MA: Peter Smith.
- Lachenbruch, P.A. and M.R. Mickey (1968). Estimation of error rates in discriminant analysis. *Technometrics*, 10, 1-11.
- Lam, P.S. (1990). The Hamilton model with a general autoregressive component: Estimation and comparison with other models of economic time series. *J. Monetary Econ.*, 26, 409-432.
- Lay, T. (1997). Research required to support comprehensive nuclear test ban treaty monitoring. *National Research Council Report, National Academy Press*, 2101 Constitution Ave., Washington, DC 20055.
- Levinson, N. (1947). The Wiener (root mean square) error criterion in filter design and prediction. *J. Math. Phys.*, 25, 262-278.
- Lindgren, G. (1978). Markov regime models for mixed distributions and switching regressions. *Scand. J. Stat.*, 5, 81-91.
- Ljung, G.M. and G.E.P. Box (1978). On a measure of lack of fit in time series models. *Biometrika*, 65, 297-303.
- Lütkepohl, H. (1985). Comparison of criteria for estimating the order of a vector autoregressive process. *J. Time Series Anal.*, 6, 35-52.
- Lütkepohl, H. (1993). *Introduction to Multiple Time Series Analysis, 2nd ed.* Berlin: Springer-Verlag.
- MacQueen, J.B. (1967). Some methods for classification and analysis of multivariate observations. *Proceedings of 5-th Berkeley Symposium on Mathematical Statistics and Probability*. Berkeley: University of California Press, 1:281-297

- Mallows, C.L. (1973). Some comments on C_p . *Technometrics*, 15, 661-675.
- McBratney, A.B. and R. Webster (1981). Detection of ridge and furrow pattern by spectral analysis of crop yield. *Int. Stat. Rev.*, 49, 45-52.
- McCulloch, R.E. and R.S. Tsay (1993). Bayesian inference and prediction for mean and variance shifts in autoregressive time series. *J. Am. Stat. Assoc.*, 88, 968-978.
- McDougall, A. J., D.S. Stoffer and D.E. Tyler (1997). Optimal transformations and the spectral envelope for real-valued time series. *J. Stat. Plan. Infer.*, 57, 195-214.
- McLeod A.I. (1978). On the distribution of residual autocorrelations in Box-Jenkins models. *J. R. Stat. Soc. B*, 40, 296-302.
- McLeod, A.I. and K.W. Hipel (1978). Preservation of the rescaled adjusted range, I. A reassessment of the Hurst phenomenon. *Water Resour. Res.*, 14, 491-508.
- McQuarrie, A.D.R. and C-L. Tsai (1998). *Regression and Time Series Model Selection*, Singapore: World Scientific.
- Meinhold, R.J. and N.D. Singpurwalla (1983). Understanding the Kalman filter. *Am. Stat.*, 37, 123-127.
- Meinhold, R.J. and N.D. Singpurwalla (1989). Robustification of Kalman filter models. *J. Am. Stat. Assoc.*, 84, 479-486.
- Meng X.L. and Rubin, D.B. (1991). Using EM to obtain asymptotic variance-covariance matrices: The SEM algorithm. *J. Am. Stat. Assoc.*, 86, 899-909.
- Metropolis N., A.W. Rosenbluth, M.N. Rosenbluth, A. H. Teller, and E. Teller (1953). Equations of state calculations by fast computing machines. *J. Chem. Phys.*, 21, 1087-1091.
- Mickens, R.E. (1990). *Difference Equations: Theory and Applications (2nd ed)*. New York: Springer.
- Nason, G.P. (2008). *Wavelet Methods in Statistics with R*. New York: Springer.
- Newbold, P. and T. Bos (1985). *Stochastic Parameter Regression Models*. Beverly Hills: Sage.
- Ogawa, S., T.M. Lee, A. Nayak and P. Glynn (1990). Oxygenation-sensitive contrast in magnetic resonance image of rodent brain at high magnetic fields. *Magn. Reson. Med.*, 14, 68-78.
- Palma, W. (2007). *Long-Memory Time Series: Theory and Methods*. New York: Wiley.
- Palma, W. and N.H. Chan (1997). Estimation and forecasting of long-memory time series with missing values. *J. Forecast.*, 16, 395-410.
- Paparoditis, E. and Politis, D.N. (1999). The local bootstrap for periodogram statistics. *J. Time Series Anal.*, 20, 193-222.
- Parzen, E. (1962). On estimation of a probability density and mode. *Ann. Math. Stat.*, 35, 1065-1076.
- Parzen, E. (1983). Autoregressive spectral estimation. In *Time Series in the Frequency Domain, Handbook of Statistics*, Vol. 3, pp. 211-243. D.R. Brillinger and P.R. Krishnaiah eds. Amsterdam: North Holland.
- Pawitan, Y. and R.H. Shumway (1989). Spectral estimation and deconvolution for a linear time series model. *J. Time Series Anal.*, 10, 115-129.
- Peña, D. and I. Guttman (1988). A Bayesian approach to robustifying the Kalman filter. In *Bayesian Analysis of Time Series and Dynamic Linear Models*, pp. 227-254. J.C. Spall, ed. New York: Marcel Dekker.

- Percival, D.B. and A.T. Walden (1993). *Spectral Analysis for Physical Applications: Multitaper and Conventional Univariate Techniques* Cambridge: Cambridge University Press.
- Percival, D.B. and A.T. Walden (2000). *Wavelet Methods for Time Series Analysis*. Cambridge: Cambridge University Press.
- Petrìs G, Petrone S, Campagnoli P (2009). *Dynamic Linear Models with R*. New York: Springer.
- Phillips, P.C.B. (1987). Time series regression with a unit root. *Econometrica*, 55, 227-301.
- Phillips, P.C.B. and P. Perron (1988). Testing for unit roots in time series regression. *Biometrika*, 75, 335-346.
- Pinsker, M.S. (1964). *Information and Information Stability of Random Variables and Processes*, San Francisco: Holden Day.
- Pole, P.J. and M. West (1988). Nonnormal and nonlinear dynamic Bayesian modeling. In *Bayesian Analysis of Time Series and Dynamic Linear Models*, pp. 167-198. J.C. Spall, ed. New York: Marcel Dekker.
- Press, W.H., S.A. Teukolsky, W. T. Vetterling, and B.P. Flannery (1993). *Numerical Recipes in C: The Art of Scientific Computing, 2nd ed.* Cambridge: Cambridge University Press.
- Priestley, M.B., T. Subba-Rao and H. Tong (1974). Applications of principal components analysis and factor analysis in the identification of multi-variable systems. *IEEE Trans. Automat. Contr.*, AC-19, 730-734.
- Priestley, M.B. and T. Subba-Rao (1975). The estimation of factor scores and Kalman filtering for discrete parameter stationary processes. *Int. J. Contr.*, 21, 971-975.
- Priestley, M.B. (1981). *Spectral Analysis and Time Series*. Vol. 1: Univariate Series; Vol 2: Multivariate Series, Prediction and Control. New York: Academic Press.
- Priestley, M.B. (1988). *Nonlinear and Nonstationary Time Series Analysis*. London: Academic Press.
- Quandt, R.E. (1972). A new approach to estimating switching regressions. *J. Am. Stat. Assoc.*, 67, 306-310.
- Rabiner, L.R. and B.H. Juang (1986). An introduction to hidden Markov models, *IEEE Acoust., Speech, Signal Process.*, ASSP-34, 4-16.
- Rao, C.R. (1973). *Linear Statistical Inference and Its Applications*. New York: Wiley.
- Rauch, H.E., F. Tung, and C.T. Striebel (1965). Maximum likelihood estimation of linear dynamic systems. *J. AIAA*, 3, 1445-1450.
- Reinsel, G.C. (1997). *Elements of Multivariate Time Series Analysis, 2nd ed.* New York: Springer-Verlag.
- Rényi, A. (1961). On measures of entropy and information. In *Proceedings of 4th Berkeley Symp. Math. Stat. and Probability*, pp. 547-561, Berkeley: Univ. of California Press.
- Rissanen, J. (1978). Modeling by shortest data description. *Automatica*, 14, 465-471.
- Robinson, P.M. (1995). Gaussian semiparametric estimation of long range dependence. *Ann. Stat.*, 23, 1630-1661.
- Robinson, P.M. (2003). *Time Series With Long Memory*. Oxford: Oxford University Press.
- Rosenblatt, M. (1956a). A central limit theorem and a strong mixing condition. *Proc. Nat. Acad. Sci.*, 42, 43-47.
- Rosenblatt, M. (1956b). Remarks on some nonparametric estimates of a density functions. *Ann. Math. Stat.*, 27 642-669.

- Royston, P. (1982). An extension of Shapiro and Wilk's W test for normality to large samples. *Applied Statistics*, 31, 115-124.
- Said E. and D.A. Dickey (1984). Testing for unit roots in autoregressive moving average models of unknown order. *Biometrika*, 71, 599-607.
- Sandmann, G. and S.J. Koopman (1998). Estimation of stochastic volatility models via Monte Carlo maximum likelihood. *J. Econometrics*, 87, 271-301.
- Sargan, J.D. (1964). Wages and prices in the United Kingdom: A study in econometric methodology. In *Econometric Analysis for National Economic Planning*, eds. P. E. Hart, G. Mills and J. K. Whitaker. London: Butterworths, reprinted in *Quantitative Economics and Econometric Analysis*, pp. 275-314, eds. K. F. Wallis and D. F. Hendry (1984). Oxford: Basil Blackwell.
- Scheffé, H. (1959). *The Analysis of Variance*. New York: Wiley.
- Schuster, A. (1898). On the investigation of hidden periodicities with application to a supposed 26 day period of meteorological phenomena. *Terrestrial Magnetism*, III, 11-41.
- Schuster, A. (1906). On the periodicities of sunspots. *Phil. Trans. R. Soc., Ser. A*, 206, 69-100.
- Schwarz, F. (1978). Estimating the dimension of a model. *Ann. Stat.*, 6, 461-464.
- Schweppe, F.C. (1965). Evaluation of likelihood functions for Gaussian signals. *IEEE Trans. Inform. Theory*, IT-4, 294-305.
- Shephard, N. (1996). Statistical aspects of ARCH and stochastic volatility. In *Time Series Models in Econometrics, Finance and Other Fields*, pp 1-100. D.R. Cox, D.V. Hinkley, and O.E. Barndorff-Nielsen eds. London: Chapman and Hall.
- Shumway, R.H. and W.C. Dean (1968). Best linear unbiased estimation for multivariate stationary processes. *Technometrics*, 10, 523-534.
- Shumway, R.H. (1970). Applied regression and analysis of variance for stationary time series. *J. Am. Stat. Assoc.*, 65, 1527-1546.
- Shumway, R.H. (1971). On detecting a signal in N stationarily correlated noise series. *Technometrics*, 10, 523-534.
- Shumway, R.H. and A.N. Unger (1974). Linear discriminant functions for stationary time series. *J. Am. Stat. Assoc.*, 69, 948-956.
- Shumway, R.H. (1982). Discriminant analysis for time series. In *Classification, Pattern Recognition and Reduction of Dimensionality, Handbook of Statistics Vol. 2*, pp. 1-46. P.R. Krishnaiah and L.N. Kanal, eds. Amsterdam: North Holland.
- Shumway, R.H. and D.S. Stoffer (1982). An approach to time series smoothing and forecasting using the EM algorithm. *J. Time Series Anal.*, 3, 253-264.
- Shumway, R.H. (1983). Replicated time series regression: An approach to signal estimation and detection. In *Time Series in the Frequency Domain, Handbook of Statistics Vol. 3*, pp. 383-408. D.R. Brillinger and P.R. Krishnaiah, eds. Amsterdam: North Holland.
- Shumway, R.H. (1988). *Applied Statistical Time Series Analysis*. Englewood Cliffs, NJ: Prentice-Hall.
- Shumway, R.H., R.S. Azari, and Y. Pawitan (1988). Modeling mortality fluctuations in Los Angeles as functions of pollution and weather effects. *Environ. Res.*, 45, 224-241.
- Shumway, R.H. and D.S. Stoffer (1991). Dynamic linear models with switching. *J. Am. Stat. Assoc.*, 86, 763-769, (Correction: V87 p. 913).

- Shumway, R.H. and K.L. Verosub (1992). State space modeling of paleoclimatic time series. In *Pro. 5th Int. Meeting Stat. Climatol.*, Toronto, pp. 22-26, June, 1992.
- Shumway, R.H., S.E. Kim and R.R. Blandford (1999). Nonlinear estimation for time series observed on arrays. Chapter 7, S. Ghosh, ed. *Asymptotics, Nonparametrics and Time Series*, pp. 227-258. New York: Marcel Dekker.
- Small, C.G. and D.L. McLeish (1994). *Hilbert Space Methods in Probability and Statistical Inference*. New York: Wiley.
- Smith, A.F.M. and M. West (1983). Monitoring renal transplants: An application of the multiprocess Kalman filter. *Biometrics*, 39, 867-878.
- Spliid, H. (1983). A fast estimation method for the vector autoregressive moving average model with exogenous variables. *J. Am. Stat. Assoc.*, 78, 843-849.
- Stoffer, D.S. (1982). Estimation of Parameters in a Linear Dynamic System with Missing Observations. Ph.D. Dissertation. Univ. California, Davis.
- Stoffer, D.S., M. Scher, G. Richardson, N. Day, and P. Coble (1988). A Walsh-Fourier analysis of the effects of moderate maternal alcohol consumption on neonatal sleep-state cycling. *J. Am. Stat. Assoc.*, 83, 954-963.
- Stoffer, D.S. and K.D. Wall (1991). Bootstrapping state space models: Gaussian maximum likelihood estimation and the Kalman filter. *J. Am. Stat. Assoc.*, 86, 1024-1033.
- Stoffer, D.S., D.E. Tyler, and A.J. McDougall (1993). Spectral analysis for categorical time series: Scaling and the spectral envelope. *Biometrika*, 80, 611-622.
- Stoffer, D.S. (1999). Detecting common signals in multiple time series using the spectral envelope. *J. Am. Stat. Assoc.*, 94, 1341-1356.
- Stoffer, D.S. and K.D. Wall (2004). Resampling in State Space Models. In *State Space and Unobserved Component Models Theory and Applications*, Chapter 9, pp. 227-258. Andrew Harvey, Siem Jan Koopman, and Neil Shephard, eds. Cambridge: Cambridge University Press.
- Sugiura, N. (1978). Further analysis of the data by Akaike's information criterion and the finite corrections, *Commun. Statist. A, Theory Methods*, 7, 13-26.
- Taniguchi, M., M.L. Puri, and M. Kondo (1994). Nonparametric approach for non-Gaussian vector stationary processes. *J. Mult. Anal.*, 56, 259-283.
- Tanner, M. and W.H. Wong (1987). The calculation of posterior distributions by data augmentation (with discussion). *J. Am. Stat. Assoc.*, 82, 528-554.
- Taylor, S. J. (1982). Financial returns modelled by the product of two stochastic processes – A study of daily sugar prices, 1961-79. In Anderson, O. D., editor, *Time Series Analysis: Theory and Practice*, Volume 1, pages 203–226. New York: Elsevier/North-Holland.
- Tiao, G.C. and R.S. Tsay (1989). Model specification in multivariate time series (with discussion). *J. Roy. Statist. Soc. B*, 51, 157-213.
- Tiao, G. C. and R.S. Tsay (1994). Some advances in nonlinear and adaptive modeling in time series analysis. *J. Forecast.*, 13, 109-131.
- Tiao, G.C., R.S. Tsay and T. Wang (1993). Usefulness of linear transformations in multivariate time series analysis. *Empir. Econ.*, 18, 567-593.
- Tierney, L. (1994). Markov chains for exploring posterior distributions (with discussion). *Ann. Stat.*, 22, 1701-1728.

- Tong, H. (1983). *Threshold Models in Nonlinear Time Series Analysis*. Springer Lecture Notes in Statistics, 21. New York: Springer-Verlag.
- Tong, H. (1990). *Nonlinear Time Series: A Dynamical System Approach*. Oxford: Oxford Univ. Press.
- Tsay, Ruey S. (2002). *Analysis of Financial Time Series*. New York: Wiley.
- Venables, W.N. and B.D. Ripley (1994). *Modern Applied Statistics with S-Plus*. New York: Springer-Verlag.
- Watson, G.S. (1966). Smooth regression analysis. *Sankhya*, 26, 359-378.
- Weiss, A.A. (1984). ARMA models with ARCH errors. *J. Time Series Anal.*, 5, 129-143.
- West, M. and J. Harrison (1997). *Bayesian Forecasting and Dynamic Models 2nd ed.* New York: Springer-Verlag.
- Whittle, P. (1961). Gaussian estimation in stationary time series. *Bull. Int. Stat. Inst.*, 33, 1-26.
- Wiener, N. (1949). *The Extrapolation, Interpolation and Smoothing of Stationary Time Series with Engineering Applications*. New York: Wiley.
- Wu, C.F. (1983). On the convergence properties of the EM algorithm. *Ann. Stat.*, 11, 95-103.
- Young, P.C. and D.J. Pedregal (1998). Macro-economic relativity: Government spending, private investment and unemployment in the USA. Centre for Research on Environmental Systems and Statistics, Lancaster University, U.K.
- Yule, G.U. (1927). On a method of investigating periodicities in disturbed series with special reference to Wolfer's Sunspot Numbers. *Phil. Trans. R. Soc. Lond.*, A226, 267-298.

Index

- ACF, 19, 22
 - large sample distribution, 28, 486
 - multidimensional, 35
 - of an $AR(p)$, 99
 - of an $AR(1)$, 82
 - of an $AR(2)$, 95
 - of an $ARMA(1,1)$, 100
 - of an $MA(q)$, 98
 - sample, 28
- AIC, 51, 146, 209
 - multivariate case, 274
- AICc, 51, 146
 - multivariate case, 274
- Aliasing, 10, 172
- Amplitude, 171
 - of a filter, 220
- Analysis of Power, *see* ANOPOW
- ANOPOW, 384, 392–394
 - designed experiments, 401
- AR model, 12, 80
 - conditional sum of squares, 121
 - bootstrap, 132
 - conditional likelihood, 121
 - estimation
 - large sample distribution, 117, 496
 - likelihood, 120
 - maximum likelihood estimation, 120
 - missing data, 370
 - operator, 81
 - polynomial, 89
 - spectral density, 181
 - threshold, 264
 - unconditional sum of squares, 121
 - vector, *see* VAR
 - with observational noise, 293
- ARCH model
 - ARCH(m), 260
 - ARCH(1), 257
 - estimation, 258
 - GARCH, 261, 348
- ARFIMA model, 245, 249
- ARIMA model, 136
 - fractionally integrated, 249
 - multiplicative seasonal models, 154
 - multivariate, 273
- ARMA model, 88
 - ψ -weights, 97
 - conditional least squares, 123
 - pure seasonal models
 - behavior of ACF and PACF, 152
 - unconditional least squares, 123
 - backcasts, 116
 - behavior of ACF and PACF, 104
 - bootstrap, 331
 - causality of, 90
 - conditional least squares, 124
 - forecasts, 111
 - mean square prediction error, 112
 - based on infinite past, 111
 - prediction intervals, 114
 - truncated prediction, 113
- Gauss–Newton, 125
- in state-space form, 327
- invertibility of, 90
- large sample distribution of estimators, 128

- likelihood, 122
- MLE, 123
- multiplicative seasonal model, 152
- pure seasonal model, 150
- unconditional least squares, 125
- vector, *see* VARMA model
- ARMAX model, 282, 326
 - bootstrap, 331
 - in state-space form, 326
- ARX model, 275
- Autocorrelation function, *see* ACF
- Autocovariance function, 17, 22, 81
 - multidimensional, 35
 - random sum of sines and cosines, 172
 - sample, 27
- Autocovariance matrix, 33
 - sample, 33
- Autoregressive Integrated Moving Average Model, *see* ARIMA model
- Autoregressive models, *see* AR model
- Autoregressive Moving Average Models, *see* ARMA model

- Backcasting, 115
- Backshift operator, 59
- Bandwidth, 193
 - equivalent, 207
- Bartlett kernel, 203
- Beam, 389
- Best linear predictor, *see* BLP
- BIC, 52, 146, 209
 - multivariate case, 274, 277
- BLP, 105
 - m -step-ahead prediction, 109
 - mean square prediction error, 110
 - one-step-ahead prediction, 106
 - definition, 105
 - one-step-ahead prediction
 - mean square prediction error, 107
 - stationary processes, 106
- Bone marrow transplant series, 290, 316
- Bonferroni inequality, 198
- Bootstrap, 131, 194, 208, 331
 - stochastic volatility, 352
- Bounded in probability O_p , 472
- Brownian motion, 254

- Cauchy sequence, 489
- Cauchy–Schwarz inequality, 469, 489

- Causal, 83, 90, 493
 - conditions for an AR(2), 92
 - vector model, 282
- CCF, 20, 24
 - large sample distribution, 30
 - sample, 30
- Central Limit Theorem, 477
 - M-dependent, 478
- Cepstral analysis, 237
- Characteristic function, 474
- Chernoff information, 424
- Cluster analysis, 429
- Coherence, 212
 - estimation, 214
 - hypothesis test, 214, 519
 - multiple, 381
- Completeness of L^2 , 470
- Complex normal distribution, 514
- Complex roots, 96
- Conditional least squares, 123
- Convergence in distribution, 474
 - Basic Approximation Theorem, 475
- Convergence in probability, 471
- Convolution, 180
- Cosine transform
 - large sample distribution, 504
 - of a vector process, 377
 - properties, 188
- Cospectrum, 211
 - of a vector process, 377
- Cramér–Wold device, 475
- Cross-correlation function, *see* CCF
- Cross-covariance function, 19
 - sample, 30
- Cross-spectrum, 211
- Cycle, 171

- Daniell kernel, 200, 201
 - modified, 200, 201
- Deconvolution, 397
- Density function, 16
- Designed experiments, *see* ANOPOW
- Deterministic process, 499
- Detrending, 47
- DFT, 68
 - inverse, 184
 - large sample distribution, 504
 - multidimensional, 232
 - of a vector process, 377

- likelihood, 377
- Differencing, 57, 58, 60
- Discriminant analysis, 416
- DLM, 289, 324
 - Bayesian approach, 356
 - bootstrap, 331
 - innovations form, 330
 - maximum likelihood estimation
 - large sample distribution, 313
 - via EM algorithm, 310, 316
 - via Newton-Raphson, 305
 - MCMC methods, 362
 - observation equation, 289
 - state equation, 289
 - steady-state, 312
 - with switching, 336
 - EM algorithm, 342
 - maximum likelihood estimation, 342
- DNA series, 450, 454
- Durbin–Levinson algorithm, 108
- Earthquake series, 9, 375, 413, 419, 425, 430
- EEG sleep data, 448
- EM algorithm, 309
 - complete data likelihood, 309
 - DLM with missing observations, 316
 - expectation step, 309
 - maximization step, 310
- Explosion series, 9, 375, 413, 419, 425, 430
- Exponentially Weighted Moving Averages, 137
- Factor analysis, 438
 - EM algorithm, 440
- Fejér kernel, 203, 207
- FFT, 68
- Filter, 60
 - amplitude, 220, 221
 - band-pass, 230
 - design, 230
 - high-pass, 217, 230
 - linear, 216
 - low-pass, 217, 230
 - matrix, 221, 222
 - optimum, 228
 - phase, 220, 221
 - recursive, 230
 - seasonal adjustment, 230
 - spatial, 232
 - time-invariant, 470
- fMRI, *see* Functional magnetic resonance imaging series
- Folding frequency, 172, 174
- Fourier frequency, 68, 184
- Fractional difference, 60, 245
 - fractional noise, 245
- Frequency bands, 178, 192
- Frequency response function, 180
 - of a first difference filter, 216
 - of a moving average filter, 216
- Functional magnetic resonance imaging series, 8, 374, 402, 405, 408, 436, 441
- Fundamental frequency, 68, 174, 184
- Glacial varve series, 61, 126, 145, 247, 256
- Global temperature series, 4, 56, 60, 292
- Gradient vector, 306, 369
- Growth rate, 139, 256
- Harmonics, 197
- Hessian matrix, 306, 369
- Hidden Markov model, 337, 340
 - estimation, 342
- Hilbert space, 489
 - closed span, 490
 - conditional expectation, 492
 - projection mapping, 490
 - regression, 491
- Homogeneous difference equation
 - first order, 93
 - general solution, 97
 - second order, 93
 - solution, 94
- Impulse response function, 180
- Influenza series, 264, 343
- Infrasound series, 388, 391, 394, 397, 399
- Inner product space, 489
- Innovations, 143, 304
 - standardized, 143
 - steady-state, 312
- Innovations algorithm, 110
- Integrated models, 135, 137, 154
 - forecasting, 137
- Interest rate and inflation rate series, 332
- Invertible, 87
 - vector model, 283
- J-divergence measure, 429

- Johnson & Johnson quarterly earnings series, 3, 321
- Joint distribution function, 16
- Kalman filter, 295
 correlated noise, 325
 innovations form, 330
 Riccati equation, 312
 stability, 312
 with missing observations, 314
 with switching, 339
 with time-varying parameters, 296
- Kalman smoother, 299, 367
 as a smoothing spline, 318
 for the lag-one covariance, 303
 spline smoothing, 320
 with missing observations, 314
- Kullback-Leibler information, 75, 423
- LA Pollution – Mortality Study, 52, 73, 149, 275, 277, 328
- Lag, 19, 25
- Lagged regression model, 268
- Lake Shasta series, 373, 379, 384
- Lead, 25
- Leakage, 203
 sidelobe, 203
- Least squares estimation, *see* LSE
- Likelihood
 AR(1) model, 120
 conditional, 121
 innovations form, 122, 305
- Linear filter, *see* Filter
- Linear process, 25, 90
- Ljung–Box–Pierce statistic, 144
 multivariate, 280
- Local level model, 298, 301, 359
- Long memory, 60, 245
 estimation, 246
 estimation of d , 251
 spectral density, 250
- LSE, 48
 conditional sum of squares, 121
 Gauss–Newton, 124
 unconditional, 121
- MA model, 12, 85
 autocovariance function, 18, 98
 Gauss–Newton, 126
 mean function, 17
 operator, 85
 polynomial, 89
 spectral density, 181
- Maximum likelihood estimation, *see* MLE
- Mean function, 16
- Mean square convergence, 469
- Method of moments estimators, *see* Yule–Walker
- Minimum mean square error predictor, 105
- Missing data, 316
- MLE
 ARMA model, 123
 conditional likelihood, 121
 DLM, 305
 state-space model, 305
 via EM algorithm, 309
 via Newton–Raphson, 123, 305
 via scoring, 123
- Moving average model, *see* MA model
- New York Stock Exchange, 6, 350
- Newton–Raphson, 123
- Normal distribution
 marginal density, 16
 multivariate, 26, 514
- NYSE, *see* New York Stock Exchange
- Order in probability o_p , 472
- Orthogonality property, 490
- PACF, 102
 of an MA(1), 103
 iterative solution, 109
 large sample results, 118
 of an AR(p), 102
 of an AR(1), 102
 of an MA(q), 103
- Parameter redundancy, 88
- Partial autocorrelation function, *see* PACF
- Parzen window, 206
- Period, 171
- Periodogram, 69, 185
 distribution, 189
 matrix, 423
 scaled, 67
- Phase, 171
 of a filter, 220
- Pitch period, 5

- Prediction equations, 106
- Prewhiten, 32, 269
- Principal components, 432
- Projection Theorem, 490
- Q-test, *see* Ljung–Box–Pierce statistic
- Quadspectrum, 211
 - of a vector process, 377
- Random sum of sines and cosines, 172, 500, 502
- Random walk, 13, 17, 21, 137
 - autocovariance function, 19
- Recruitment series, 6, 31, 62, 104, 114, 190, 194, 201, 215, 224, 269
- Regression
 - ANOVA table, 50
 - autocorrelated errors, 147, 327
 - Cochrane-Orcutt procedure, 148
 - for jointly stationary series, 378
 - ANOPOW table, 384
 - Hilbert space, 491
 - lagged, 222
 - model, 47
 - multiple correlation, 50
 - multivariate, 273, 327
 - normal equations, 48
 - random coefficients, 396
 - spectral domain, 378
 - stochastic, 332, 396
 - ridge correction, 397
 - with deterministic inputs, 387
- Return, 6, 139, 256
- Riesz–Fisher Theorem, 470
- Scaling, 448
- Scatterplot matrix, 54, 62
- Scatterplot smoothers
 - kernel, 70
 - lowess, 71, 73
 - nearest neighbors, 71
 - splines, 72
- Score vector, 306
- SIC, 52
- Signal plus noise, 14, 15, 227, 389
 - mean function, 17
- Signal-to-noise ratio, 14, 228
- Sinc kernel, 207
- Sine transform
 - large sample distribution, 504
 - of a vector process, 377
 - properties, 188
- Smoothing splines, 71, 318
- Soil surface temperature series, 35, 36, 232
- Southern Oscillation Index, 6, 31, 62, 190, 194, 201, 203, 209, 215, 217, 224, 228, 269
- Spectral density, 178
 - autoregression, 208
 - estimation, 193
 - adjusted degrees of freedom, 194
 - bandwidth stability, 199
 - confidence interval, 193
 - degrees of freedom, 193
 - large sample distribution, 193
 - nonparametric, 207
 - parametric, 207
 - resolution, 199
 - matrix, 213
 - linear filter, 222
 - of a filtered series, 180
 - of a moving average, 181
 - of an AR(2), 181
 - of white noise, 179
 - wavenumber, 231
- Spectral distribution function, 178
- Spectral envelope, 447
 - categorical time series, 451
 - real-valued time series, 457
- Spectral Representation Theorem, 177, 178, 500, 502
 - vector process, 213, 503
- Speech series, 5, 29
- Spline smoothing, 320
- State-space model
 - Bayesian approach, 356
 - linear, *see* DLM
- Stationary
 - Gaussian series, 26
 - jointly, 24
 - strictly, 20
 - weakly, 21
- Stochastic process, 10
 - realization, 10
- Stochastic regression, 332
- Stochastic trend, 136
- Stochastic volatility model, 348
 - bootstrap, 352

- estimation, 350
- Structural component model, 73, 321, 343
- Taper, 202, 203
 - cosine bell, 203
- Taylor series expansion in probability, 473
- Tchebycheff inequality, 469
- Threshold autoregressive model, 264
- Time series, 10
 - categorical, 451
 - complex-valued, 432
 - multidimensional, 34, 231
 - multivariate, 20, 33
 - two-dimensional, 231
- Transfer function model, 268
- Transformation
 - Box-Cox, 60
 - via spectral envelope, 460
- Trend stationarity, 22
- Triangle inequality, 489
- Tukey-Hanning window, 206
- U.S. GNP series, 140, 144, 147, 260, 458
- U.S. macroeconomic series, 444
- U.S. population series, 146
- Unconditional least squares, 123
- Unit root tests, 253
 - Augmented Dickey-Fuller test, 255
 - Dickey-Fuller test, 255
 - Phillips-Perron test, 255
- VAR model, 274, 276
 - estimation
 - large sample distribution, 281
 - operator, 282
 - Variogram, 37, 44
 - VARMA model, 282
 - autocovariance function, 283
 - estimation
 - Spliid algorithm, 285
 - identifiability of, 285
 - Varve series, 251
 - VMA model, 282
 - operator, 282
 - Volatility, 6, 256
- Wavenumber spectrum, 231
 - estimation, 232
- Weak law of large numbers, 472
- White noise, 11
 - autocovariance function, 18
 - Gaussian, 11
 - vector, 274
- Whittle likelihood, 210, 422
- Wold Decomposition, 499
- Yule-Walker
 - equations, 117
 - vector model, 281
 - estimators, 117
 - AR(2), 118
 - MA(1), 119
 - large sample results, 117