


Analysis of Repeated Measures and Time Series: An Introduction with Forestry Examples

Biometrics Information Handbook No.6

15/1995



Province of British Columbia
Ministry of Forests Research Program



Analysis of Repeated Measures and Time Series: An Introduction with Forestry Examples

Biometrics Information Handbook No.6

Amanda F. Linnell Nemec



Province of British Columbia
Ministry of Forests Research Program

The use of trade, firm, or corporation names in this publication is for the information and convenience of the reader. Such use does not constitute an official endorsement or approval by the Government of British Columbia of any product or service to the exclusion of any others that may also be suitable. Contents of this report are presented for discussion purposes only.

Citation:

Nemec, Amanda F. Linnell. 1996. Analysis of repeated measures and time series: an introduction with forestry examples. Biom. Inf. Handb. 6. Res. Br., B.C. Min. For., Victoria, B.C. Work. Pap. 15/1996.

Prepared by

Amanda F. Linnell Nemec
International Statistics and Research Corporation
P.O. Box 496
Brentwood Bay, BC V8M 1R3
for
B.C. Ministry of Forests
Research Branch
31 Bastion Square
Victoria, BC V8W 3E7

Copies of this report may be obtained, depending upon supply, from:
B.C. Ministry of Forests
Forestry Division Services Branch
1205 Broad Street
Victoria, BC V8W 3E7

© 1996 Province of British Columbia

The contents of this report may not be cited in whole or in part without the approval of the Director of Research, B.C. Ministry of Forests, Victoria, B.C.

Repeated measures and time-series data are common in forestry. Because such data tend to be serially correlated—that is, current measurements are correlated with past measurements—they require special methods of analysis. This handbook is an introduction to two broad classes of methods developed for this purpose: repeated-measures analysis of variance and time-series analysis. Both types of analyses are described briefly and are illustrated with forestry examples. Several procedures for the analysis of repeated measures and time series are available in the SAS/STAT and SAS/ETS libraries. Application of the REPEATED statement in PROC GLM (and PROC ANOVA) and the time-series procedures PROC AUTOREG, PROC ARIMA, and PROC FORECAST are discussed.

ACKNOWLEDGEMENTS

The author thanks all individuals who responded to the request for repeated-measures and time-series data. Their contributions were essential in the development of this handbook. Many constructive criticisms, suggestions, and references were received from the 12 reviewers of the first and second drafts. Discussions with Vera Sit, Wendy Bergerud, Ian Cameron, and Dave Spittlehouse of the Research Branch were particularly helpful and contributed much to the final content of the handbook. Financial support was provided by the B.C. Ministry of Forests and International Statistics and Research Corporation.

Abstract	iii
Acknowledgements	iv
1 Introduction	1
1.1 Examples	1
1.1.1 Repeated measurement of seedling height	2
1.1.2 Missing tree rings	3
1.1.3 Correlation between ring index and rainfall	3
1.2 Definitions	4
1.2.1 Trend, cyclic variation, and irregular variation	6
1.2.2 Stationarity	7
1.2.3 Autocorrelation and cross-correlation	8
2 Repeated-measures Analysis	9
2.1 Objectives	9
2.2 Univariate Analysis of Repeated Measures	11
2.3 Multivariate Analysis of Repeated Measures	13
3 Time-series Analysis	14
3.1 Objectives	14
3.2 Descriptive Methods	14
3.2.1 Time plot	15
3.2.2 Correlogram and cross-correlogram	15
3.2.3 Tests of randomness	20
3.3 Trend	21
3.4 Seasonal and Cyclic Components	22
3.5 Time-Series Models	23
3.5.1 Autoregressions and moving averages	24
3.5.2 Advanced topics	25
3.6 Forecasting	27
4 Repeated-measures and Time-series Analysis with SAS	28
4.1 Repeated-measures Analysis	28
4.1.1 Repeated-measures data sets	28
4.1.2 Univariate analysis	30
4.1.3 Multivariate analysis	38
4.2 Time-series Analysis	40
4.2.1 Time-series data sets	40
4.2.2 PROC ARIMA	43
4.2.3 PROC AUTOREG	57
4.2.4 PROC FORECAST	58
5 SAS Examples	62
5.1 Repeated-measures Analysis of Seedling Height Growth	62
5.2 Cross-correlation Analysis of Missing Tree Rings	70
6 Conclusions	78

APPENDIX 1	Average height of seedlings	79
APPENDIX 2	Ring widths	80
APPENDIX 3	Ring index and rainfall	81
References	82

TABLES

1	Split-plot ANOVA model for seedling experiment	12
2	Analysis of annual height increments: summary of <i>p</i> -values	70

FIGURES

1	Average height of seedlings	2
2	Missing tree rings	4
3	Comparison of ring index with annual spring rainfall	5
4	Temporal variation	6
5	Daily photosynthetically active radiation	8
6	Null hypotheses for repeated-measures analysis	10
7	Time plots of annual snowfall for Victoria, B.C.	16
8	White noise	17
9	Soil temperatures	18
10	Correlograms for ring-index and rainfall series	19
11	Cross-correlogram for prewhitened ring-index and rainfall series	20
12	Smoothed daily soil temperatures	23
13	Time series generated by AR, MA, ARMA, and ARIMA models ...	26
14	Univariate repeated-measures analysis of seedling data: univariate data set	31
15	Univariate repeated-measures analysis of seedling data: multivariate data set	34
16	Multivariate repeated-measures analysis of seedling data	39
17	Time plot of weekly soil temperatures created with PROC TIMEPLOT	43
18	Time-series analysis of ring-index series: model identification	45
19	Cross-correlation of prewhitened ring-index and rainfall series ...	48
20	Time-series analysis of ring-index series: model estimation	51
21	Time-series analysis of ring-index series: PROC AUTOREG	59
22	Ring-index forecasts generated with PROC FORECAST	62
23	Repeated-measures analysis of the growth of Douglas-fir and lodgepole pine seedlings	64
24	Cross-correlation analysis of missing tree rings	73

Forestry data are often collected over time or space.¹ In trials to compare several treatments, tree height and diameter are typically measured before treatments are applied and on one or more occasions after application. Sometimes data are collected more frequently or over extended periods. Microclimatic conditions are generally monitored on a daily or hourly basis, or at even shorter intervals, for periods of several weeks, months, or years. Tree rings, growth and yield, timber prices, reforestation costs, forest fire occurrence, insect infestations, animal populations, and water quality are also observed at regular intervals so that trends or cyclic patterns can be studied. These diverse examples have one common feature: the same unit or process is measured on more than one occasion. Such data tend to be serially correlated, or autocorrelated, which means that the most recent measurements are dependent on, or to some extent predictable from, past observations. Because this violates the independence assumption on which many standard statistical methods are based, alternative methods are required for their analysis. Two broad classes of methods have been developed for this purpose: repeated-measures analysis and time-series analysis.

This handbook is a brief introduction to repeated-measures and time-series analysis, with an emphasis on methods that are most likely to be applicable to forestry data. The objective of the handbook is to help the reader recognize when repeated-measures or time-series methods are applicable, and to provide general guidance in their selection and use. Most mathematical details have been omitted, but some familiarity with analysis of variance and regression analysis, and an understanding of such basic statistical concepts as the mean and variance of a random variable and the correlation between two variables are required. Readers are also assumed to have a working knowledge of SAS.² The discussion begins with three examples (Section 1.1), which are used to illustrate the ideas and methods that are covered in subsequent sections. The examples are followed by some definitions (Section 1.2). Repeated-measures analysis of variance is discussed in Section 2 and general time-series methods are described in Section 3. Elementary SAS programs for carrying out repeated-measures analyses and some simple time-series analyses are included in Section 4. Additional examples are given in Section 5. For more information about a particular topic, the reader should consult the list of references at the end of the handbook.

1.1 Examples

Before proceeding with the definitions and a discussion of methods, it will be helpful to describe some situations in which repeated-measures or time-series data arise. The first example (Section 1.1.1) is a typical repeated-

1 The methods discussed in this handbook can be generalized to data collected over space (e.g., Rossi et al. 1992), or any other index by which measurements can be arranged in a logical sequence or array.

2 SAS is a registered trademark of SAS Institute Inc., Cary, N.C.

measures experiment. Repeated-measures designs are often used to assess treatment effects on trees or vegetation, and to monitor growth and yield in permanent sample plots. The second and third examples involve tree rings, which is an important area of application of time-series methods in forestry. Section 1.1.2 illustrates how several tree-ring series from a single tree can be used to reconstruct the growth history of a tree. In Section 1.1.3, the correspondence between ring width and rainfall is examined.

1.1.1 Repeated measurement of seedling height To assess the effects of three site-preparation treatments, four blocks comprising 12 rows of 25 seedlings were established at a single trial site in the Sub-Boreal Spruce (SBS) dry warm subzone in the Cariboo Forest Region. Three site-preparation treatments (V = v-plow, S = 30 × 30 cm hand screef, and U = an untreated control), two seedling species (FD = Douglas-fir and PL = lodgepole pine), and two types of stock (B = bareroot and P = plug) were randomly assigned to the rows, with one row for each of the 12 combinations. Seedling height, diameter, condition, and survival were measured at the time of planting (1983) and annually for the next six years (1984–1989). Figure 1 shows the average height of the seedlings that

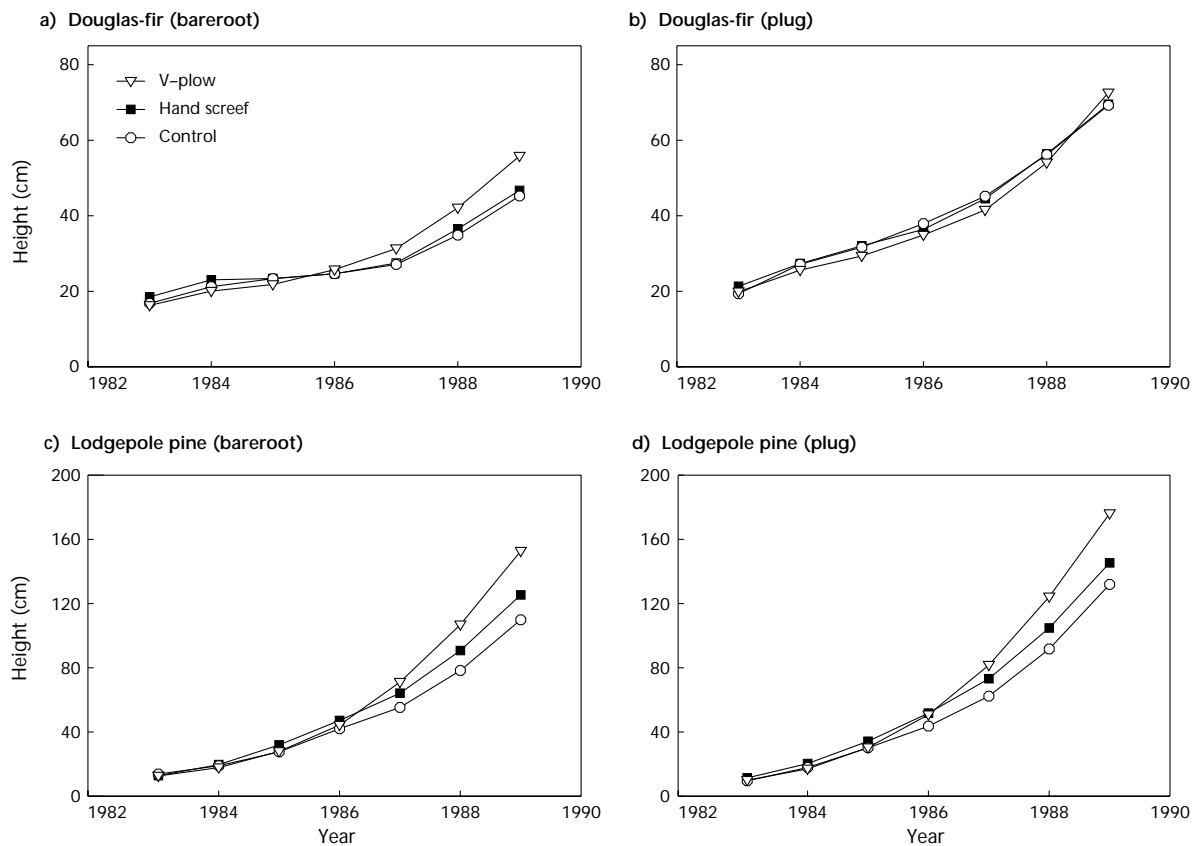


FIGURE 1 Average height of seedlings: (a) Douglas-fir grown from bareroot stock, (b) Douglas-fir grown from plugs, (c) lodgepole pine grown from bareroot stock, and (d) lodgepole pine grown from plugs.

survived to 1989 (i.e., the average over seedlings in all rows and blocks) plotted against year, for each of the three site-preparation treatments (the data are in Appendix 1). The objective of the experiment is to determine whether treatment or stock type affects the growth of either species of seedling.

1.1.2 Missing tree rings Tree rings are a valuable source of information. When cross-sectional disks are cut at several heights, the growth history of a tree can be reconstructed by determining the year that the tree first reached the height of each disk (i.e., the year when the innermost ring of the disk was formed). For disks that have a complete complement of rings, this is a simple matter of counting backwards from the outermost ring (which is assumed to correspond to the year in which the tree was cut) to the year of the innermost ring. Dating rings is more complicated if, during the course of its growth, a tree experiences adverse growing conditions and in response fails to produce a uniform sheath of xylem each year. If this happens, one or more rings will be missing in at least some disks (e.g., the sheath might not fully encircle a disk or it might not extend down as far as the disk).

Figure 2 shows two tree-ring series from a paper birch tree (the data are in Appendix 2). Figure 2a is for a disk cut at a height of 1.3 m; Figure 2b shows the corresponding series for a disk taken at 2.0 m. Eleven additional disks were sampled at heights ranging from 0.3 m to 20 m. In Figure 2c, the height of each disk is plotted against the year of the innermost ring, with no adjustment for missing rings. Until it was felled in 1993, the tree in Figure 2 was growing in a mixed birch and conifer stand. In the early stages of development of the stand, the birch trees were taller than the conifers, but during the forty years before cutting they were overtopped by the conifers. Because paper birch is a shade-intolerant species, the trees were subject to increasing stress and therefore some of the outermost rings are expected to be missing, especially in disks cut near the base of the tree.

One method of adjusting for missing rings (Cameron 1993) is to align the tree-ring series by comparing patterns of growth. If there are no missing rings, then the best match should be achieved by aligning the outermost ring of each series. Otherwise, each series is shifted by an amount equal to the estimated number of missing rings and the growth curve is adjusted accordingly (Figure 2d). The same approach is used to date trees except that an undated ring series from one tree is aligned with a dated series from a second tree, or with a standard chronology. For more information about the time-series analysis of tree rings, refer to Monserud (1986).

1.1.3 Correlation between ring index and rainfall The width of a tree ring depends on the age of the tree. Typically, ring width increases rapidly when the tree is young, decreases as the tree matures, and eventually levels out. Ring width is also affected by climate and environmental conditions. To reveal the less obvious effects of rainfall, air temperature, or pollution, the dominant growth trend is removed from the ring-width series by a

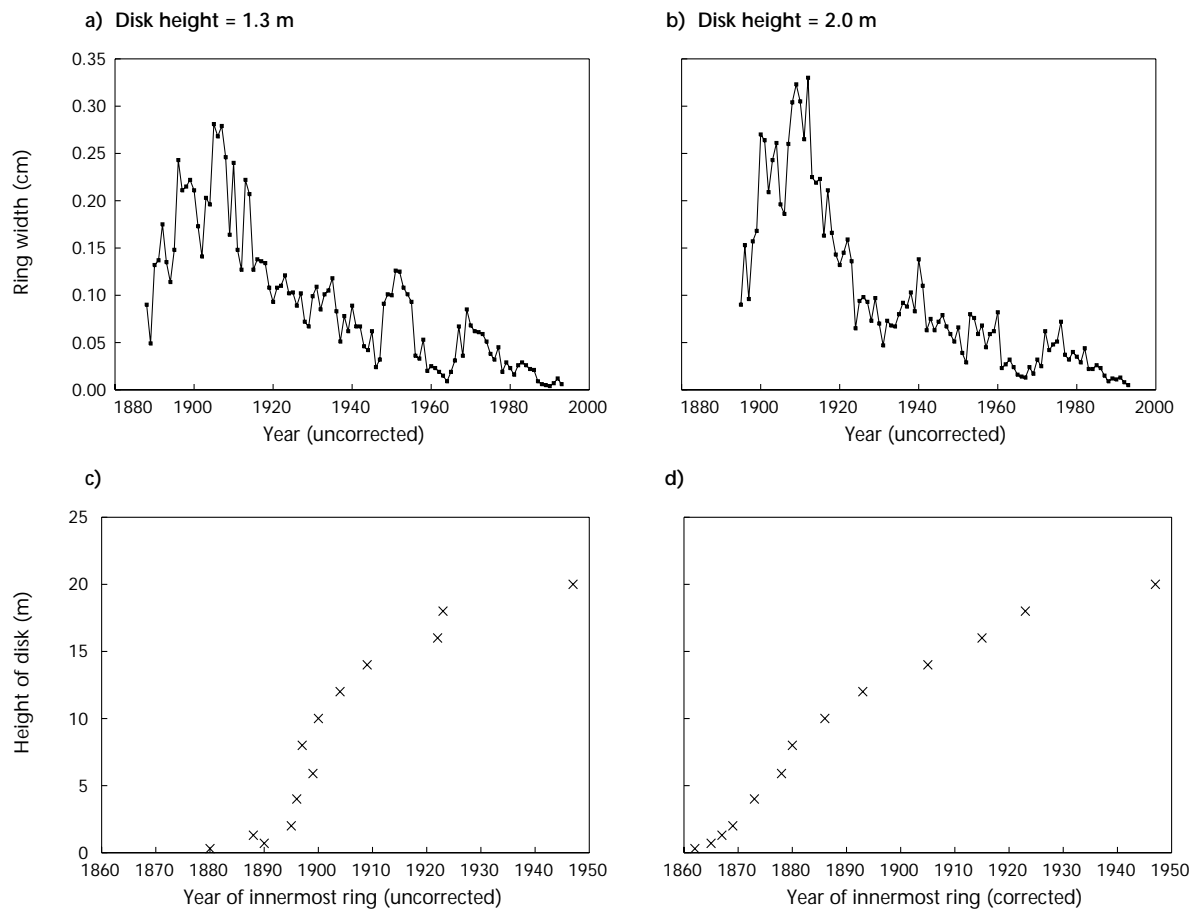


FIGURE 2 Missing tree rings: (a) ring widths for disk at 1.3 m, (b) ring width for disk at 2.0 m, (c) disk height versus year (no correction for missing rings), and (d) disk height versus year (corrected for missing rings).

process called *detrending* or “standardization”³ (refer to Section 3.3). The detrended ring width is called a *ring index*. Figure 3a shows a ring-index series for a Douglas-fir tree on the Saanich Peninsula, while Figure 3b gives the total rainfall during March, April, and May of the same years, as recorded at the Dominion Astrophysical Observatory (corrected, adjusted, and extended by comparison with stations at Gonzales Observatory and Victoria Airport). Data for the two series are given in Appendix 3. In this example, the investigator wants to determine whether annual spring rainfall has any effect on ring width.

1.2 Definitions

Let y_1, y_2, \dots, y_n be a sequence of measurements (average height of a row of seedlings, ring width, annual rainfall, etc.) made at n distinct times. Such data are called *repeated measures*, if the measurements are

³ A set of computer programs for standardizing tree-ring chronologies is available from the International Tree-Ring Data Bank (1993).

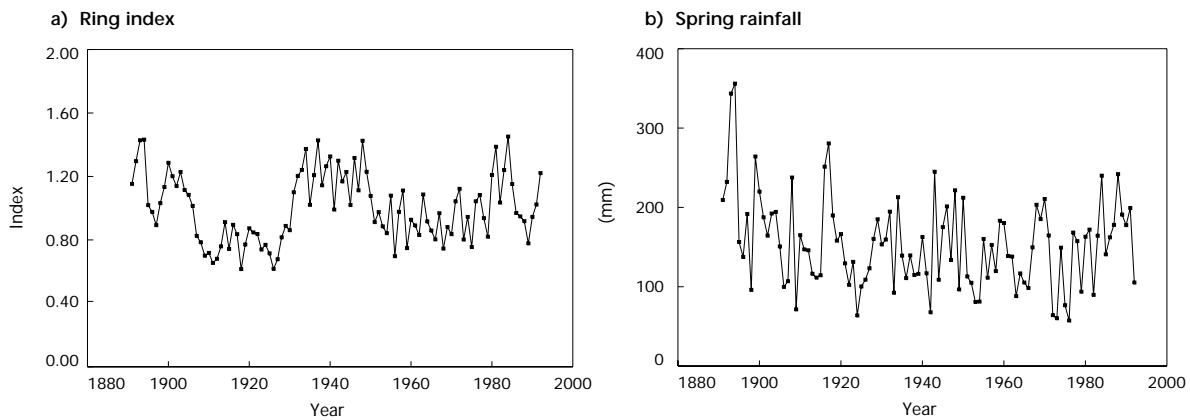


FIGURE 3 Comparison of ring index with annual spring rainfall: (a) ring index versus year and (b) total rainfall during March, April, and May versus year.

made on relatively few occasions (e.g., $n \leq 10$), or a *time series*, if the number of observations is large (e.g., $n \geq 25$). Thus the seedling measurements (Figure 1) would normally be considered repeated measures, while the tree-ring data and rainfall measurements (Figures 2a, 2b, 3a, and 3b) are time series. Figures 1–3 illustrate another common distinction between repeated measures and time series. Repeated-measures designs generally include experimental units (e.g., rows of trees) from two or more study groups (e.g., site-preparation, species, and stock-type combinations)—notice that each curve in Figure 1 represents a separate group of seedlings. In contrast, time series often originate from a single population or experimental unit (e.g., a single tree or weather station). This division, which is based on the number of observation times and presence or absence of experimental groups, is more or less arbitrary, but should help the reader recognize when a repeated-measures analysis is warranted and when the methods that are generally referred to as time-series analysis are applicable.

Repeated-measures analysis is a type of analysis of variance (ANOVA), in which variation between experimental units (often called “between-subjects” variation) and variation within units (called “within-subjects” variation) are examined. Between-units variation can be attributed to the factors that differ across the study groups (e.g., treatment, species, and stock type). Within-units variation is any change, such as an increase in height, that is observed in an individual experimental unit. In Figure 1, the between-units variation accounts for the separation of the curves, while the within-units variation determines the shape of the curves (if there were no within-units variation then the curves would be flat). The objectives of a repeated-measures analysis are twofold: (1) to determine how the experimental units change over time and (2) to compare the changes across study groups.

Time-series analysis encompasses a much wider collection of methods than repeated-measures analysis. It includes descriptive methods, model-fitting techniques, forecasting and regression-type methods, and spectral

analysis. Time-series analysis is concerned with short- and long-term changes, and the correlation or dependence between past and present measurements.

1.2.1 Trend, cyclic variation, and irregular variation Forestry researchers are frequently interested in temporal variation. If repeated-measures or time-series data are plotted against time, one or more of three distinct types of variation will be evident (Figure 4). The simplest type of variation is a *trend* (Figure 4a), which is a relatively slow shift in the level of the data. Trends can be linear (Figure 4a) or nonlinear (Figures 2a and 2b), and can correspond to an increase or decrease in the mean, or both. The growth in height of a tree or row of seedlings (e.g., Figure 1) is a familiar example of a trend.

Some data oscillate at more or less regular intervals as illustrated in Figure 4b. This type of variation is called *cyclic variation*. Insect and animal populations sometimes display cyclic variation in their numbers. *Seasonal variation* is cyclic variation that is controlled by seasonal factors




FIGURE 4 Temporal variation: (a) linear trend (and irregular variation), (b) seasonal (and irregular) variation, (c) irregular variation, and (d) trend, seasonal, and irregular variation.

and therefore completes exactly one cycle per year. Air temperatures typically exhibit a seasonal increase in the spring and summer, and a corresponding decrease in the fall and winter. The distinction between trend and cyclic variation can depend on the length of the observation period and on the frequency of the measurements. If only part of a cycle is completed during the period of observation, then cyclic variation becomes indistinguishable from trend. Identification of a cyclic component is also impossible if sampling is too infrequent to cover the full range of variability (e.g., if the observation times happen to coincide with the maximum of each cycle, the data will show no apparent periodicity).

The third type of temporal variation is called *residual* or *irregular variation*. It includes any noncyclic change that cannot be classified as a trend. Figure 4c shows a typical example. Notice that there is no trend—the data fluctuate irregularly about a constant mean (horizontal line)—and there are no obvious cycles. Irregular variation is the result of isolated or random events. Measurement error and sampling error are probably the most common and best-known sources of irregular variation. There are, however, many other factors that produce irregular variation. The rainfall series in Figure 3b is an example. It shows irregular variation resulting from random changes in the meteorological conditions that produce rain. The ring-index series (Figure 3a) also exhibits irregular variation, which probably reflects changes in environmental conditions.

Trend, cyclic variation, and irregular variation can occur simultaneously (as illustrated in Figure 4d) or in various combinations. One of the first steps in an analysis is to identify the components of interest. Because repeated-measures data comprise relatively few observation times, repeated-measures analysis is concerned mainly with trends. Time series are often sufficiently long and detailed that both trend and cyclic variation are potentially of interest. An irregular component is invariably present in both repeated measures and time series. In many applications, irregular variation is attributed entirely to error. Although this variation must be considered in the selection of a suitable probability model, it is not the focus of the study. In other studies, such as the ring-index and rainfall example, irregular variation is the main component under investigation.

1.2.2 Stationarity A time series is *stationary* if its statistical properties are invariant over time. This implies that the mean and variance are the same for all epochs (e.g., the mean and variance for the first 20 years are the same as those for the last 20 years). The series shown in Figure 4c is stationary. Notice that the data fluctuate about a fixed value and the amplitude of the fluctuations remains constant. Data that exhibit a trend (Figure 4a and 4d) or cyclic variation (Figures 4b and 4d) are *nonstationary* because the mean changes with time. A time-dependent variance is another common form of nonstationarity. In some cases, both the mean and variance vary. The daily photosynthetically active radiation (PAR) measurements plotted in Figure 5 display the latter behaviour. Notice that as the average light level falls off, the variability of the measurements also tends to decrease.




FIGURE 5 Daily photosynthetically active radiation (PAR).

Stationarity is a simplifying assumption that underlies many time-series methods. If a series is nonstationary, then the nonstationary components (e.g., trend) must be removed or the series transformed (e.g., to stabilize a nonstationary variance) before the methods can be applied. Nonstationarity must also be considered when computing summary statistics. For instance, the sample mean is not particularly informative if the data are seasonal, and should be replaced with a more descriptive set of statistics, such as monthly averages.

1.2.3 Autocorrelation and cross-correlation Repeated measures and time series usually exhibit some degree of autocorrelation. *Autocorrelation*, also known as *serial correlation*, is the correlation between any two measurements y_s and y_t in a sequence of measurements y_1, y_2, \dots, y_n (i.e., correlation between a series and itself, hence the prefix “auto”). Seedling heights and tree-ring widths are expected to be serially correlated because unusually vigorous or poor growth in one year tends to carry over to the next year. Serially correlated data violate the assumption of independence on which many ANOVA and regression methods are based. Therefore, the underlying models must be revised before they can be applied to such data.

The autocorrelation between y_s and y_t can be positive or negative, and the magnitude of the correlation can be constant, or decrease more or less quickly, as the time interval between the observations increases. The *autocorrelation function* (ACF) is a convenient way of summarizing the dependence between observations in a stationary time series. If the observations y_1, y_2, \dots, y_n are made at n equally spaced times and y_t is the observation at time t , let y_{t+1} be the next observation (i.e., the measurement made one step ahead), let y_{t+2} be the measurement made two steps

ahead and, in general, let y_{t+k} be the observation made k steps ahead. The time interval, or delay, between y_t and y_{t+k} is called the *lag* (i.e., y_t lags y_{t+k} by k time steps) and the autocorrelation function evaluated at lag k is

$$ACF(k) = \frac{Cov(y_t, y_{t+k})}{Var(y_t)}$$

The numerator of the function is the covariance between y_t and y_{t+k} and the denominator is the variance of y_t , which is the same as the variance of y_{t+k} , since the series is assumed to be stationary. Notice that, by definition, $ACF(0) = 1$ and $ACF(k) = ACF(-k)$. The latter symmetry property implies that the autocorrelation function need only be evaluated for $k \geq 0$ (or $k \leq 0$).

The ACF can be extended to two stationary series x_1, x_2, \dots, x_n and y_1, y_2, \dots, y_n (e.g., the ring index and rainfall series of Section 1.1.3) by defining the *cross-correlation function* (CCF). At lag k , this function is the correlation between x_t and y_{t+k} :

$$CCF(k) = \frac{Cov(x_t, y_{t+k})}{\sqrt{Var(x_t) Var(y_t)}}$$

Notice that, unlike the autocorrelation function, the cross-correlation function is not necessarily one at lag 0 (because the correlation between x_t and y_t is not necessarily one) and $CCF(k)$ is not necessarily equal to $CCF(-k)$ (i.e., the CCF is not necessarily symmetric). Therefore, the cross-correlation function must be evaluated at $k = 0, \pm 1, \pm 2$, etc.

The auto- and cross-correlation functions play key roles in time-series analysis. They are used extensively for data summary, model identification, and verification.

2 REPEATED-MEASURES ANALYSIS

For simplicity, the discussion of repeated measures is restricted to a single repeated factor, as illustrated by the seedling example in Section 1.1.1. In this and in many other forestry applications, year, or more generally time, is the only repeated factor. If the heights of the seedlings were measured in the spring and fall of each year, or for several years before and after a fertilizer is applied, the design would include two repeated factors—season and year, or fertilizer and year. Designs with two or more repeated factors lead to more complicated analyses than the one-factor case considered here, but the overall approach (i.e., the breakdown of the analysis into a within- and between-units analysis) is the same.

2.1 Objectives

There are three types of hypotheses to be tested in a repeated-measures analysis:

- H_{01} : the growth curves or trends are parallel for all groups (i.e., there are no interactions involving time),
- H_{02} : there are no trends (i.e., there are no time effects), and
- H_{03} : there are no overall differences between groups (i.e., the between-units factors have no effect).

The three hypotheses are illustrated in Figure 6 with the simple case of one between-units factor. This figure shows the expected average height of a row of lodgepole pine seedlings (Section 1.1.1), grown from plugs, plotted against year. Each curve corresponds to a different site-preparation treatment, which is the between-units factor. Hypotheses H_{01} and H_{02} concern changes over time, which are examined as part of the within-units analysis. Hypothesis H_{01} (Figure 6a) implies that site-preparation treatment has no effect on the rate of growth of the seedlings. If this hypothesis is retained, it is often appropriate to test whether the growth curves are flat (H_{02} , Figure 6b). Acceptance of H_{02} implies that there is no change over time. In this example, H_{02} is not very interesting because the seedlings are expected to show some growth over the seven-year period of the study. The last hypothesis concerns the separation of the three growth curves and is tested as part of the between-units analysis. If the groups show parallel trends (i.e., H_{01} is true) then, H_{03} implies that growth patterns are identical for the three groups (Figure 6c). Otherwise, H_{03} implies

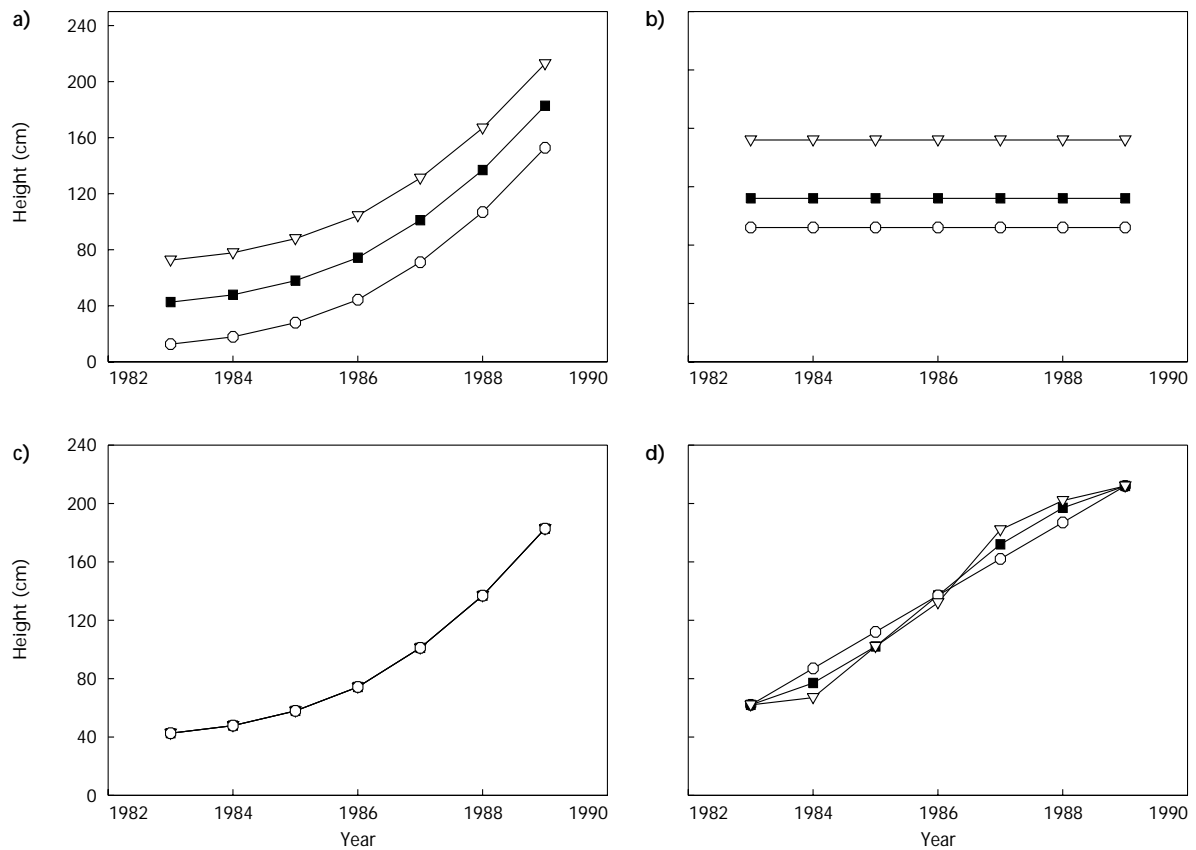


FIGURE 6 Null hypotheses for repeated-measures analysis: (a) parallel trends, (b) no trends, (c) no difference between groups, and (d) differences between groups cancel over time.

that there is no difference between the groups when the effects of the site-preparation treatments are averaged over time. Relative gains and losses are cancelled over time, as illustrated in Figure 6d.

Rejection of H_{01} implies that the trends are not parallel for at least two groups. When this occurs, additional hypotheses can be tested to determine the nature of the divergence (just as multiple comparisons are used to pinpoint group differences in a factorial ANOVA). One approach is to compare the groups at each time, as suggested by Looney and Stanley (1989). Alternatively, one of the following can be tested:

- H_{04} : the expected difference between two consecutive values, $y_t - y_{t-1}$, is the same for all groups,
- H_{05} : the expected difference between an observation at time t and its initial value, $y_t - y_1$, is the same for all groups, or
- H_{06} : the expected value of the k th-order polynomial contrast $a_{1k}y_1 + a_{2k}y_2 + \dots + a_{nk}y_n$ is the same for all groups.

Each hypothesis comprises a series of $n - 1$ hypotheses about the within-row effects. In the first two cases (H_{04} and H_{05}), the $n - 1$ increments $y_2 - y_1, y_3 - y_2, \dots, y_n - y_{n-1}$, or cumulative increments $y_2 - y_1, y_3 - y_1, \dots, y_n - y_1$, are tested for significant group differences by carrying $n - 1$ separate analyses of variance. If the trends are expected to be parallel for at least part of the observation period, then H_{04} is often of interest. Alternatively, the trends might be expected to diverge initially and then converge, in which case H_{05} might be more relevant. The last hypothesis (H_{06}) is of interest when the trends are adequately described by a polynomial of order k (i.e., $\beta_0 + \beta_1 t + \dots + \beta_k t^k$). In this case, a set of coefficients $a_{1k}, a_{2k}, \dots, a_{nk}$ (see Bergerud 1988 for details) can be chosen so that the expected value of linear combination $a_{1k}y_1 + a_{2k}y_2 + \dots + a_{nk}y_n$ depends only on β_k . Thus an ANOVA of the transformed values $a_{1k}y_1 + a_{2k}y_2 + \dots + a_{nk}y_n$ is equivalent to assessing the effects of the between-units factors on β_k . If the order of the polynomial is unknown, ANOVA tests can be performed sequentially, starting with polynomials of order $n - 1$ (which is the highest-order polynomial that can be tested when there are n observation times) and ending with a comparison of linear components. Refer to Littell (1989), Bergerud (1991), Meredith and Stehman (1991), Sit (1992a), and Gumpertz and Brownie (1993) for a discussion of the comparison of polynomial and other nonlinear trends.

2.2 Univariate Analysis of Repeated Measures

Univariate repeated-measures analysis is based on a split-plot ANOVA model in which time is the split-plot factor (refer to Keppel 1973; Moser et al. 1990; or Milliken and Johnson 1992 for details). As an illustration, consider the average row heights for the seedling data (Section 1.1.1). The split-plot ANOVA (with an additive block effect) is summarized in Table 1. The top part of the table summarizes the between-rows analysis. It is equivalent to an ANOVA of the time-averaged responses $(y_1 + y_2 + \dots + y_n)/n$ and has the same sources of variation, degrees of freedom, sums of squares, expected mean squares, and F -tests as a factorial ($SPP \times STK \times TRT$) randomized block design with no repeated factors. The bottom part of the table summarizes the within-rows analysis.

TABLE 1 *Split-plot ANOVA model for seedling experiment (Section 1.1.1)*

Source of variation	Degrees of freedom	Error term for testing effect
Between rows		
Block, BLK	3	
Species, SPP	1	Error – row
Stock type, STK	1	Error – row
Site-preparation treatment, TRT	2	Error – row
SPP × STK	1	Error – row
SPP × TRT	2	Error – row
STK × TRT	2	Error – row
SPP × STK × TRT	2	Error – row
Error – row	33	
Within rows		
Time, YEAR	6	YEAR × BLK
YEAR × BLK	18	
YEAR × SPP	6	Error – row × year
YEAR × STK	6	Error – row × year
YEAR × TRT	12	Error – row × year
YEAR × SPP × STK	6	Error – row × year
YEAR × SPP × TRT	12	Error – row × year
YEAR × STK × TRT	12	Error – row × year
YEAR × SPP × STK × TRT	12	Error – row × year
Error – row × year	198	
Total	335	

It includes the main effect of time (YEAR) and all other time-related sources of variation (YEAR × SPP, YEAR × STK, YEAR × TRT, etc.), which are readily identified by forming time interactions with the factors listed in the top part of the table. If all the interactions involving time are significant, each of the 12 groups (2 species × 2 stock types × 3 site-preparation treatments) will have had a different pattern of growth. The absence of one or more interactions can simplify the comparison of growth curves. For example, if there are no interactions involving treatment and year (i.e., the terms YEAR × TRT, YEAR × SPP × TRT, YEAR × STK × TRT, and YEAR × SPP × STK × TRT are absent from the model), then the three growth curves corresponding to the three site-preparation treatments are parallel for each species and stock type.

The univariate model allows for correlation between repeated measurements of the same experimental unit (e.g., successive height measurements of the same row of seedlings). This correlation is assumed to be the same for all times and all experimental units.⁴ The univariate model, like any randomized block design, also allows for within-block correlation—that is, correlation between measurements made on different experimental units in the same block (e.g., the heights of two rows of seedlings in the same

⁴ This condition can be replaced with the less restrictive “Huynh-Feldt condition,” which is described in Chapter 26 of Milliken and Johnson (1992).

block). In the repeated-measures model, inclusion of a year-by-block interaction ($\text{YEAR} \times \text{BLK}$) implies that the within-block correlation depends on whether measurements are made on different experimental units in the same block (e.g., the heights of two rows of seedlings in the same block). Measurements made in the same year (and in the same block) are assumed to be more strongly correlated than those made in different years. However, in both cases, the correlation is assumed to be the same for all blocks and experimental units. All other measurements (e.g., the heights of two rows in different blocks) are assumed to be independent. In addition, all measurements are assumed to have the same variance.

In a univariate analysis, repeated measurements are treated as separate observations and time is included as a factor in the ANOVA model. In the multivariate approach, repeated measurements are considered elements of a single multivariate observation and the univariate within-units ANOVA is replaced with a multivariate ANOVA, or MANOVA. The main advantage of the multivariate analysis is a less restrictive set of assumptions. Unlike the univariate ANOVA model, the MANOVA model does not require the variance of the repeated measures, or the correlation between pairs of repeated measures, to remain constant over time (e.g., the variance of the average height of a row of seedlings might increase with time, and the correlation between two measurements of the same row might decrease as the time interval between the measurements increases). Both models do, however, require the variances and correlations to be homogeneous across units (e.g., for any given year, the variance of the average height of a row of seedlings is the same for all rows, as are the inter-year correlations of row heights). The greater applicability of the multivariate model is not without cost. Because the model is more general than the univariate model, more parameters (i.e., more variances and correlations) need to be estimated and therefore there are fewer degrees of freedom for a fixed sample size. Thus, for reliable results, multivariate analyses typically require larger sample sizes than univariate analyses.

The multivariate analysis of the between-units variation is equivalent to the corresponding univariate analysis. However, differences in the underlying models lead to different within-units analyses. Several multivariate test statistics can be used to test H_{01} and H_{02} in a multivariate repeated-measures analysis: Wilks' Lambda, Pillai's trace, Hotelling-Lawley trace, and Roy's greatest root. To assess the statistical significance of an effect, each statistic is referred to an F -distribution with the appropriate degrees of freedom. If the effect has one degree of freedom, then the tests based on the four statistics are equivalent. Otherwise the tests differ, although in many cases they lead to similar conclusions. In some situations, the tests lead to substantially different conclusions so the analyst must consider other factors, such as the relative power of the tests (i.e., the probability that departures from the null hypothesis will be detected), before arriving at a conclusion. For a more detailed discussion of multivariate repeated-measures analysis, refer to Morrison (1976), Hand and Taylor (1987), and Tabachnick and Fidell (1989), who discuss the pros and cons of the four MANOVA test statistics; Moser et al. (1990), who compare the multi-

variate and univariate approaches to the analysis of repeated measures; and Gumpertz and Brownie (1993), who provide a clear and detailed exposition of the multivariate analysis of repeated measures in randomized block and split-plot experiments.

3 TIME-SERIES ANALYSIS

Time series can be considered from two perspectives: the *time domain* and the *frequency domain*. Analysis in the time domain relies on the auto-correlation and cross-correlation functions (defined in Section 1.2.3) to describe and explain the variability in a time series. In the frequency domain, temporal variation is represented as a sum of sinusoidal components, and the ACF and CCF are replaced by the corresponding Fourier transforms, which are known as the spectral and cross-spectral density functions. Analysis in the frequency domain, or *spectral analysis* as it is more commonly called, is useful for detecting hidden periodicities (e.g., cycles in animal populations), but is generally inappropriate for analyzing trends and other nonstationary behaviour. Because the results of a spectral analysis tend to be more difficult to interpret than those of a time-domain analysis, the following discussion is limited to the time domain. For a comprehensive introduction to time-series analysis in both the time and frequency domains, the reader should refer to Kendall and Ord (1990); Diggle (1991), who includes a discussion (Section 4.10, Chapter 4) of the strengths and weaknesses of spectral analysis; or Chatfield (1992). For more information about spectral analysis, the reader should consult Jenkins and Watts (1968) or Bloomfield (1976).

3.1 Objectives

The objectives of a time-series analysis range from simple description to model development. In some applications, the trend or cyclic components of a series are of special interest, and in others, the irregular component is more important. In either case, the objectives usually include one or more of the following:

- data summary and description
- detection, description, or removal of trend and cyclic components
- model development and parameter estimation
- prediction of a future value (i.e., forecasting)

Many time-series methods assume that the data are equally spaced in time. Therefore, the following discussion is limited to equally spaced series (i.e., the measurements y_1, y_2, \dots, y_n are made at times $t_0 + d, t_0 + 2d, \dots, t_0 + nd$ where d is the fixed interval between observations). This is usually not a serious restriction because in many applications, observations occur naturally at regular intervals (e.g., annual tree rings) or they can be made at equally spaced times by design.

3.2 Descriptive Methods

Describing a time series is similar to describing any other data set. Standard devices include graphs and, if the series is stationary, such familiar summary statistics as the sample mean and variance. The correlogram and cross-correlogram, which are plots of the sample auto- and cross-

correlation functions, are powerful tools. They are unique to time-series analysis and offer a simple way of displaying the correlation within or between time series.

3.2.1 Time plot All time series should be plotted before attempting an analysis. A time plot—that is, a plot of the response variable y_t versus time t —is the easiest and most obvious way to describe a time series. Trends (Figure 4a), cyclic behaviour (Figure 4b), nonstationarity (Figures 2a, 2b, 5), outliers, and other prominent features of the data are often most readily detected with a time plot. Because the appearance of a time plot is affected by the choice of symbols and scales, it is always advisable to experiment with different types of plots. Figure 7 illustrates how the look of a series (Figure 7a) changes when the connecting lines are omitted (Figure 7b) and when the data are plotted on a logarithmic scale (Figure 7c). Notice that the asymmetric (peaks in one direction) appearance of the series (Figures 7a and 7b) is eliminated by a log transformation (Figure 7c). If the number of points is very large, time plots are sometimes enhanced by decimating (i.e., retaining one out of every ten points) or aggregating the data (e.g., replacing the points in an interval with their sum or average).

3.2.2 Correlogram and cross-correlogram The *correlogram*, or sample autocorrelation function, is obtained by replacing $Cov(y_t, y_{t-k})$ and $Var(y_t)$ in the true autocorrelation function (Section 1.2.3) with the corresponding sample covariance and variance:

$$\widehat{ACF}(k) = \frac{\sum_{t=1}^{n-k} (y_t - \bar{y})(y_{t+k} - \bar{y})}{\sum_{t=1}^n (y_t - \bar{y})^2} = r_k$$

and plotting *autocorrelation coefficient* r_k against k . For reliable estimates, the sample size n should be large relative to k (e.g., $n > 4k$ and $n > 50$) and, because the autocorrelation coefficients are sensitive to extreme points, the data should be free of outliers.

The correlogram contains a lot of information. The sample ACF for a purely random or “white noise” series (i.e., a series of independent, identically distributed observations) is expected to be approximately zero for all non-zero lags (Figure 8). If a time series has a trend, then the ACF falls off slowly (e.g., linearly) with increasing lags. This behaviour is illustrated in Figure 9b, which shows the correlogram for the nonstationary series of average daily soil temperatures displayed in Figure 9a. If a time series contains a seasonal or cyclic component, the correlogram also exhibits oscillatory behaviour. The correlogram for a seasonal series with monthly intervals (e.g., total monthly rainfall) might, for example, have large negative values at lags 6, 18, etc. (because measurements made in the summer and winter are negatively correlated) and large positive values at lags 12, 24, etc. (because measurements made in the same season are positively correlated).

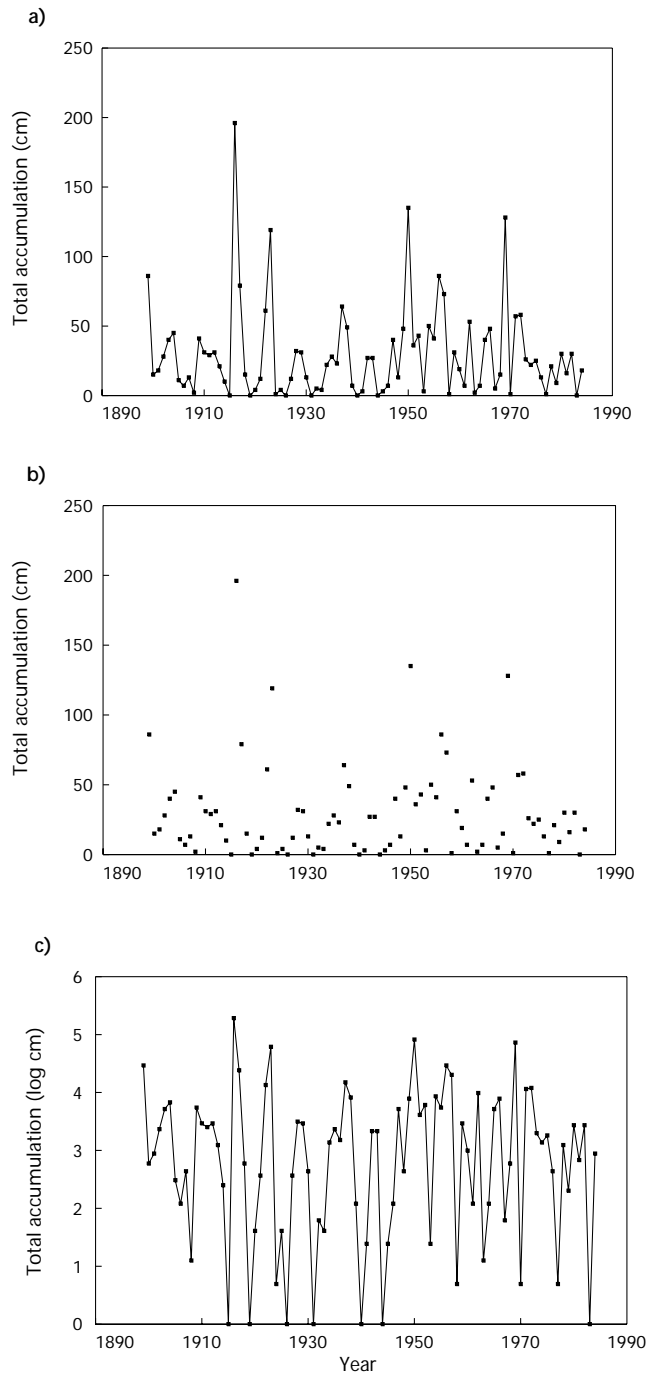


FIGURE 7 Time plots of annual snowfall for Victoria, B.C.: (a) with connecting lines, (b) without connecting lines, and (c) natural logarithmic scale.

Since the theoretical ACF is defined for stationary time series, further interpretation of the correlogram is possible only after trend and seasonal components are eliminated. Trend can often be removed by calculating the *first difference* between successive observations (i.e., $y_t - y_{t-1}$). Figure 9c shows the first difference of the soil temperature series and Figure 9d is

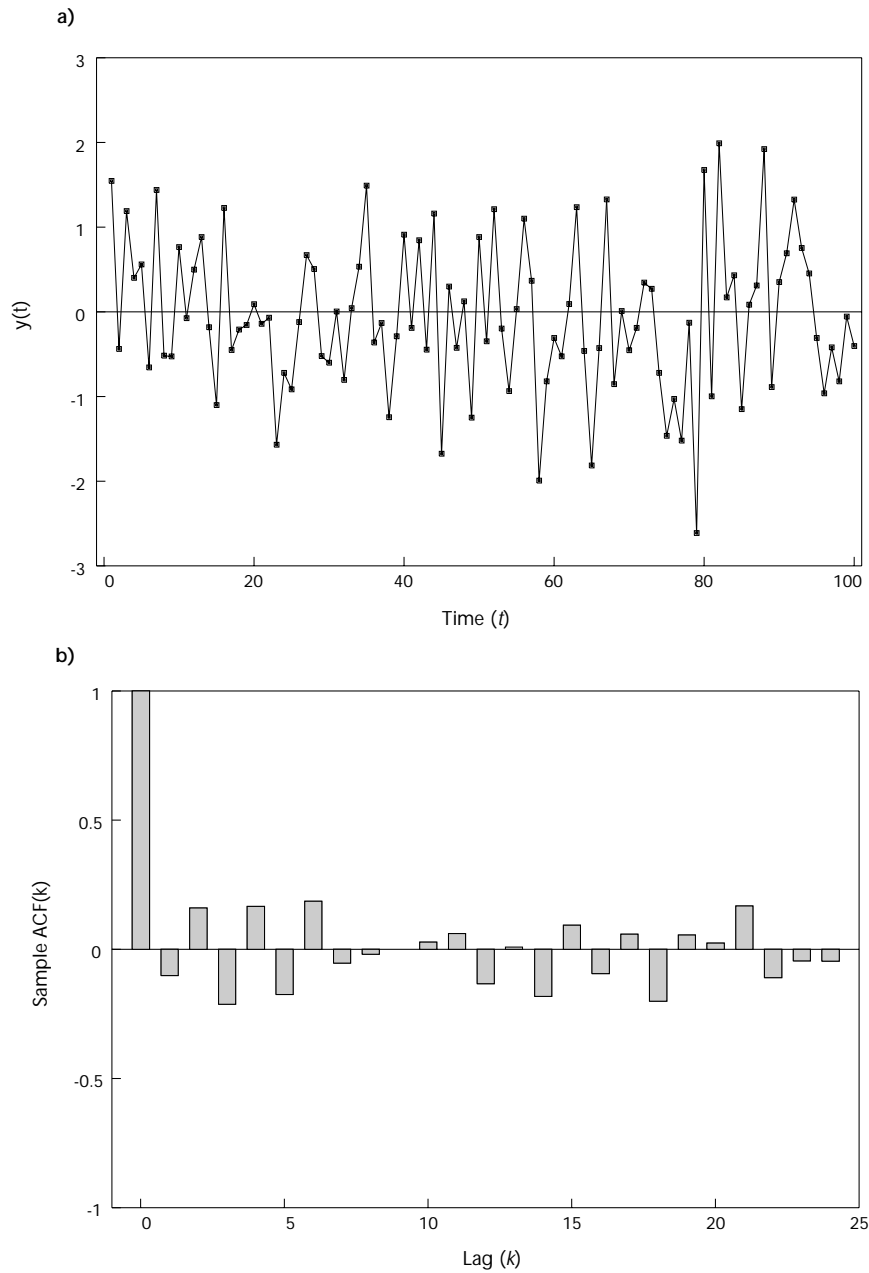


FIGURE 8 White noise: (a) time plot and (b) correlogram.

the corresponding correlogram. Notice that the trend that was evident in the original series (Figures 9a and 9b) is absent in the transformed series. The first difference is usually sufficient to remove simple (e.g., linear) trends. If more complicated (e.g., polynomial) trends are present, the difference operator can be applied repeatedly—that is, the *second difference* $[(y_t - y_{t-1}) - (y_{t-1} - y_{t-2}) = y_t - 2y_{t-1} + y_{t-2}]$, etc. can be applied to the series. Seasonal components can be eliminated by calculating an appropriate seasonal difference (e.g., $y_t - y_{t-12}$ for a monthly series or $y_t - y_{t-4}$ for a quarterly series).

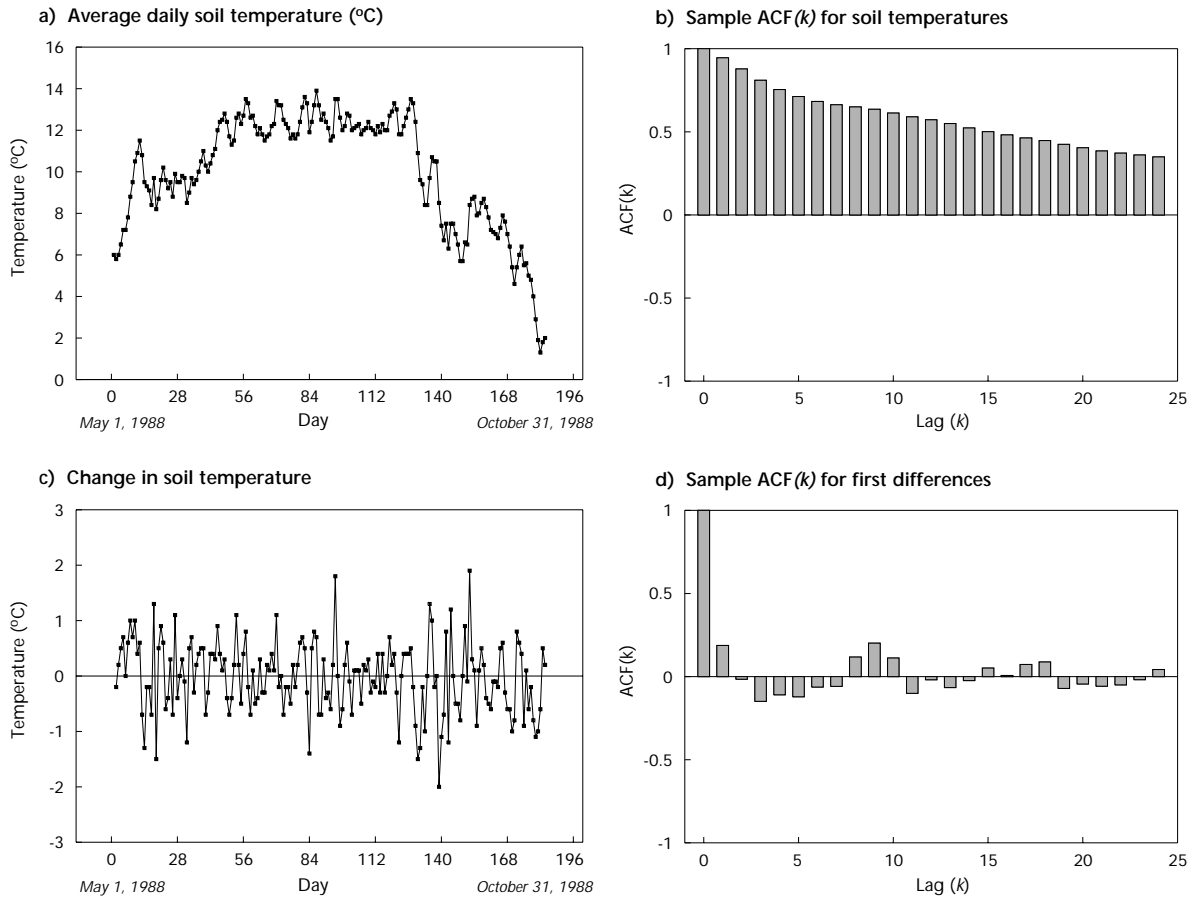


FIGURE 9 Soil temperatures: (a) daily soil temperatures, (b) correlogram for daily soil temperatures, (c) first difference of daily soil temperatures, and (d) correlogram for first differences.

Correlograms of stationary time series approach zero more quickly than processes with a nonstationary mean. The correlograms for the ring-index (Figure 3a) and rainfall (Figure 3b) series, both of which appear to be stationary, are shown in Figure 10. For some series, the ACF tails off (i.e., falls off exponentially, or consists of a mixture of damped exponentials and damped sine waves) and for others, it cuts off abruptly. The former behaviour is characteristic of autoregressions and mixed autoregressive-moving average processes, while the latter is typical of a moving average (refer to Section 3.5 for more information about these processes).

The *cross-correlogram*, or sample CCF of two series x_1, x_2, \dots, x_n and y_1, y_2, \dots, y_n , is:

$$\widehat{CCF}(k) = \frac{\sum_{t=1}^{n-k} (x_t - \bar{x})(y_{t+k} - \bar{y})}{\sqrt{\sum_{t=1}^n (x_t - \bar{x})^2 \sum_{t=1}^n (y_t - \bar{y})^2}}$$




FIGURE 10 Correlograms for ring-index and rainfall series: (a) correlogram for ring-index series shown in Figure 3a and (b) correlogram for rainfall series shown in Figure 3b.

A cross-correlogram can be more difficult to interpret than a correlogram because its statistical properties depend on the autocorrelation of the individual series, as well as on the cross-correlation between the series. A cross-correlogram might, for example, suggest that two series are cross-correlated when they are not, simply because one or both series is auto-correlated. Trends can also affect interpretation of the cross-correlogram because they dominate the cross-correlogram in much the same way as

they dominate the correlogram. To overcome these problems, time series are usually detrended and prewhitened prior to computing the sample CCF. Detrending is the removal of trend from a time series, which can be achieved by the methods described in Section 3.3. Prewhitening is the elimination of autocorrelation (and cyclic components). Time series are prewhitened by fitting a suitable model that describes the autocorrelation (e.g., one of the models described in Section 3.5.1) and then subtracting the fitted values. The resultant series of residuals is said to be *prewhitened* because it is relatively free of autocorrelation and therefore resembles white noise. More information about the cross-correlogram, and detrending and prewhitening can be obtained in Chapter 8 of Diggle (1991) or Chapter 8 of Chatfield (1992).

Figure 11 shows the cross-correlogram for the ring-index and rainfall series of Section 1.1.3 (Figure 3). Both series have been prewhitened (refer to Section 4.2.2 for details) to reduce the autocorrelation that is evident from the correlograms (Figure 10). Detrending is not required because the trend has already been removed from the ring index and the rainfall series shows no obvious trend. Notice that the cross-correlogram has an irregular pattern with one small but statistically significant spike at lag zero. There is no significant cross-correlation at any of the other lags. This suggests that ring width is weakly correlated with the amount of rainfall in the spring of the same year, but the effect does not carry over to subsequent years.




FIGURE 11 Cross-correlogram for prewhitened ring-index and rainfall series.

3.2.3 Tests of randomness Autocorrelation can often be detected by inspecting the time plot or correlogram of a series. However, some series are not easily distinguished from white noise. In such cases, it is useful to have a formal test of randomness.

There are several ways to test the null hypothesis that a stationary time series is random. If the mean of the series is zero, then the *Durbin-Watson test* can be used to test the null hypothesis that there is no first-order autocorrelation (i.e., $ACF(1) = 0$). The Durbin-Watson statistic is

$$D - W = \frac{\sum_{t=1}^{n-1} (y_{t+1} - y_t)^2}{\sum_{t=1}^n y_t^2}$$

which is expected to be approximately equal to two, if the series $\{y_t\}$ is random. Otherwise it will tend to be near zero if the series is positively correlated, or near four if the series is negatively correlated. To calculate a p -value, $D - W$ must be referred to a special table, such as that given in Durbin and Watson (1951).

Other tests of randomness use one or more of the autocorrelation coefficients r_k . If a time series is random, then r_k is expected to be zero for all $k \neq 0$. This assumption can be verified by determining the standard error of r_k under the null hypothesis (assuming the observations have the same normal distribution) and carrying out an appropriate test of significance. Alternatively, a “portmanteau” chi-squared test can be derived to test the hypothesis that autocorrelations for the first k lags are simultaneously zero (refer to Chapter 48 of Kendall et al. [1983], and Chapters 2 and 8 of Box and Jenkins [1976] for details). For information on additional tests of randomness, refer to Chapter 45 of Kendall et al. (1983).

3.3 Trend

There are two main ways to estimate trend: (1) by fitting a function of time (e.g., a polynomial or logistic growth curve) or (2) by smoothing the series to eliminate cyclic and irregular variation. The first method is applicable when the trend can be described by a fixed or deterministic function of time (i.e., a function that depends only on the initial conditions, such as the linear trend shown in Figure 4a). Once the function has been identified, the associated parameters (e.g., polynomial coefficients) can be estimated by standard regression methods, such as least squares or maximum likelihood estimation.

The second method requires no specific knowledge of the trend and is useful for describing stochastic trends (i.e., trends that vary randomly over time, such as the trend in soil temperatures shown in Figure 9a). Smoothing can be accomplished in a number of ways. An obvious method is to draw a curve by eye. A more objective estimate is obtained by calculating a weighted average of the observations surrounding each point, as well as the point itself; that is,

$$m_t = \sum_{k=-q}^p w_k y_{t+k}$$

where m_t is the smoothed value at time t and $\{w_k\}$ are weights. The simplest example is an ordinary *moving average*⁵

$$\frac{\sum_{k=-q}^p y_{t+k}}{1+p+q}$$

which is the arithmetic average of the points $y_{t-q}, \dots, y_{t-1}, y_t, y_{t+1}, \dots, y_{t+p}$. This smooths the data to a greater or lesser degree as the number of points included in the average is increased or decreased. This is illustrated in Figure 12, which compares the results of applying a 11-point (Figure 12a) and 21-point (Figure 12b) moving average to the daily soil temperatures in Figure 9a.

Moving averages attach equal weight to each of the $p + q + 1$ points $y_{t-q}, \dots, y_{t-1}, y_t, y_{t+1}, \dots, y_{t+p}$. Better estimates of trend are sometimes obtained when less weight is given to the observations that are farthest from y_t . Exponential smoothing uses weights that fall off exponentially. Many other methods of smoothing are available, including methods that are less sensitive to extreme points than moving averages (e.g., moving medians or trimmed means).

After the trend has been estimated, it can be removed by subtraction if the trend is additive, or by division if it is multiplicative. The process of removing a trend is called *detrending* and the resultant series is called a *detrended series*. Detrended series typically contain cyclic and irregular components, which more or less reflect the corresponding components of the original series. However, detrended series should be interpreted with caution because some methods of detrending can introduce spurious periodicity or otherwise alter the statistical properties of a time series (refer to Section 46.14 of Kendall et al. [1983] for details).

3.4 Seasonal and Cyclic Components

After a series has been detrended, seasonal or cyclic components (with known periods) can be estimated by regression methods or by calculating a weighted average. The first approach is applicable if the component is adequately represented by a periodic (e.g., sinusoidal) function. The second approach is similar to the methods described in the previous section except that the averaging must take into account the periodic nature of the data. For a monthly series, a simple way to estimate a seasonal component is to average values in the same month; that is,

$$s_t = \frac{\sum_{k=-p}^q y_{t+12k}}{1+p+q}$$

where the $\{y_{t+12k}\}$ are detrended values and s_t is the estimated seasonal component.

5 This moving average should not be confused with the moving average model defined in Section 3.5. The former is a function that operates on a time series and the latter is a model that describes the statistical properties of a time series.

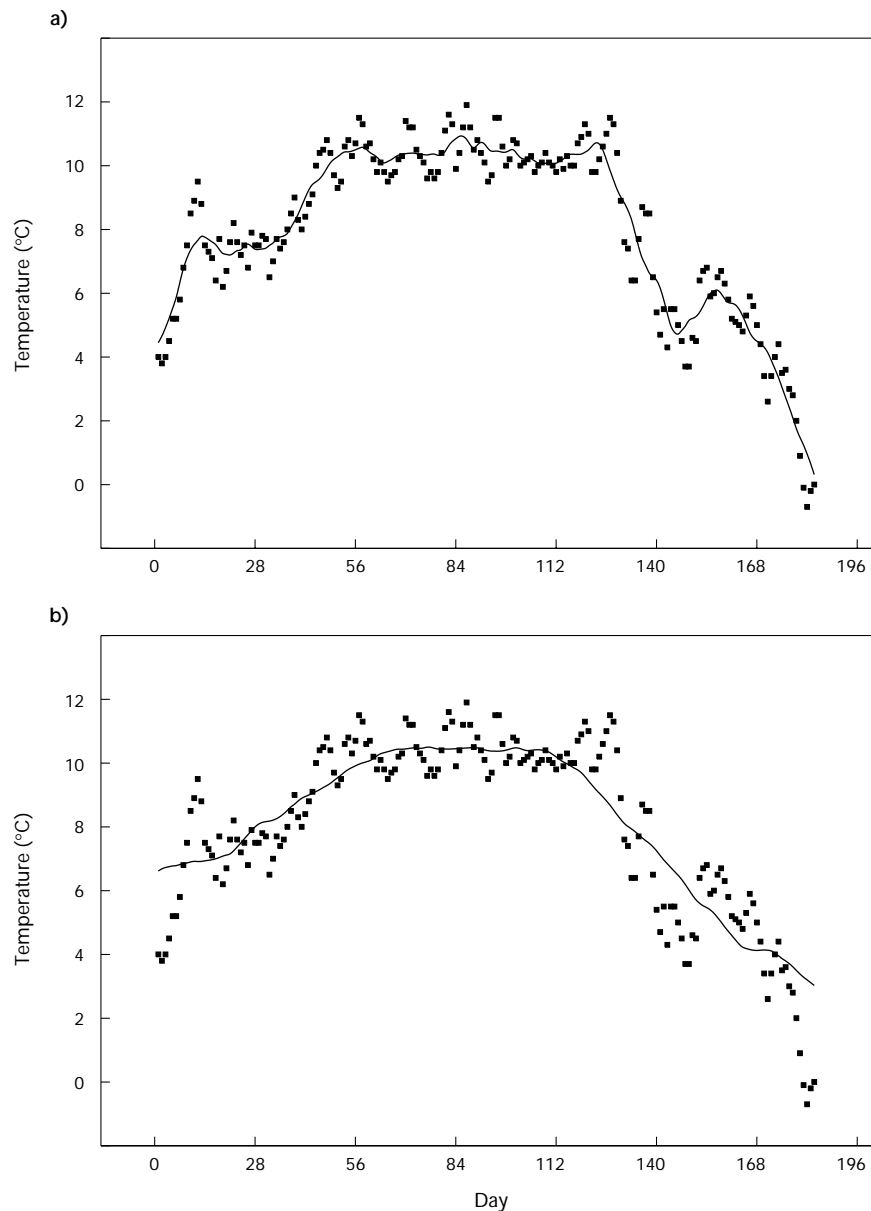


FIGURE 12 Smoothed daily soil temperatures: (a) 11-point moving average and (b) 21-point moving average.

A time series is said to be *seasonally adjusted* if a seasonal component has been removed. This can be accomplished by subtracting the estimated seasonal component s_t from the series, or by dividing by s_t , depending on whether the seasonal component is additive or multiplicative. For more information about the purposes and methods of seasonal adjustment, refer to Chapter 6 of Kendall et al. (1983) or Chapters 18 and 19 of Levenbach and Cleary (1981).

3.5 Time-series Models

Successful application of time-series methods requires a good understanding of the models on which they are based. This section provides an

overview of some models that are fundamental to the description and analysis of a single time series (Section 3.5.1). More complicated models for the analysis of two or more series are mentioned in Section 3.5.2.

3.5.1 Autoregressions and moving averages One of the simplest models for autocorrelated data is the autoregression. An *autoregression* has the same general form as a linear regression model:

$$y_t = \nu + \sum_{i=1}^p \phi_i y_{t-i} + \varepsilon_t$$

except in this case the response variables y_1, y_2, \dots, y_n are correlated because they appear on both sides of the equation (hence the name “auto” regression). The maximum lag (p) of the variables on the right side of the equation is called the *order* of the autoregression, ν is a constant, and the $\{\phi_i\}$ are unknown autoregressive parameters. Like other regression models, the errors $\{\varepsilon_t\}$, are assumed to be independent and identically (usually normally) distributed. Autoregressions are often denoted AR or AR(p).

Another simple model for autocorrelated data is the *moving average*:

$$y_t = \nu + \varepsilon_t - \sum_{i=1}^q \theta_i \varepsilon_{t-i}$$

Here the observed value y_t is a moving average of an unobserved series of independent and identically (normally) distributed random variables $\{\varepsilon_t\}$. The maximum lag q is the order of the moving average and the $\{\theta_i\}$ are unknown coefficients. Because there is overlap of the moving averages on the right side of the equation, the corresponding response variables are correlated. Moving average models are often denoted MA or MA(q).

The autoregressive and moving average models can be combined to produce a third type of model known as mixed *autoregressive-moving average*:

$$y_t = \nu + \sum_{i=1}^p \phi_i y_{t-i} + \varepsilon_t - \sum_{j=1}^q \theta_j \varepsilon_{t-j}$$

which is usually abbreviated as ARMA or ARMA(p,q). A related class of nonstationary models is obtained by substituting the first difference $\tilde{y}_t = y_t - y_{t-1}$ or the second difference $\tilde{y}_t = y_t - 2y_{t-1} + y_{t-2}$ etc., for y_t in the preceding ARMA model. The resultant model

$$\tilde{y}_t = \nu + \sum_{i=1}^p \phi_i \tilde{y}_{t-i} + \varepsilon_t - \sum_{k=1}^q \theta_k \varepsilon_{t-k}$$

is called an *autoregressive-integrated-moving average* and is abbreviated as ARIMA, or ARIMA(p,d,q), where d is the order of the difference (i.e.,

$d = 1$ if \tilde{y}_t is a first difference, $d = 2$ if \tilde{y}_t is a second difference, etc.). The ARMA model can be extended to include seasonal time series by substituting a seasonal difference for \tilde{y}_t .

The class of AR, MA, ARMA, and ARIMA models embodies a wide variety of stationary and nonstationary time series (Figure 13), which have many practical applications. All MA processes are stationary (Figures 13a, b). In contrast, all time series generated by ARIMA models are nonstationary (Figure 13f). Pure AR models and mixed ARMA models are either stationary or nonstationary (depending on the particular combination of autoregressive parameters $\{\phi_i\}$, although attention is generally restricted to the stationary case (Figures 13c, d, e). More information about AR, MA, ARMA, and ARIMA models is available in Chapter 3 of Chatfield (1992).

Box and Jenkins (1976) developed a general scheme for fitting AR, MA, ARMA, and ARIMA models, which has become known as *Box-Jenkins modelling*. The procedure has three main steps: (1) model identification (i.e., selection of p , d , and q), (2) model estimation (i.e., estimation of the parameters $\phi_1, \phi_2, \dots, \phi_p$ and $\theta_1, \theta_2, \dots, \theta_q$), and (3) model verification. Because the models have characteristic patterns of autocorrelation that depend on the values of p , d , and q , the correlogram is an important tool for model identification. The autocorrelation function is generally used in conjunction with the *partial autocorrelation function* (PACF), which measures the amount of autocorrelation that remains unaccounted for after fitting autoregressions of orders $k = 1, 2$, etc. (i.e., $\text{PACF}(k)$ is the amount of autocorrelation that cannot be explained by an autoregression of order k). These two functions provide complementary information about the underlying model: for an $\text{AR}(p)$ process, the ACF tails off and the PACF cuts off after lag p (i.e., the PACF is zero if the order of the fitted autoregression is greater than or equal to true value p); for an $\text{MA}(q)$ process, the ACF cuts off after lag q and the PACF tails off. Refer to Table 6.1 of Kendall and Ord (1990) or Figure 6.2 of Diggle (1991) for a handy guide to model identification using the ACF and PACF. Other tools for model identification include the *inverse autocorrelation function* (IACF) of an autoregressive moving average process, which is the ACF of the “inverse” process obtained by interchanging the parameters $\phi_1, \phi_2, \dots, \phi_p$ and $\theta_1, \theta_2, \dots, \theta_q$ (see Chatfield 1979 for details), and various automatic model-selection procedures, which are described in Section 11.4 of Chatfield (1992) and in Sections 7.26–31 of Kendall and Ord (1990).

After the model has been identified, the model parameters are estimated (e.g., by maximum likelihood estimation) and the adequacy of the fitted model is assessed by analyzing the residuals. For a detailed exposition of the Box-Jenkins procedure, the reader should consult Box and Jenkins (1976), which is the standard reference; McCleary and Hay (1980); or Chapters 3 and 4 of Chatfield (1992), for a less mathematical introduction to the subject.

3.5.2 Advanced topics The AR, MA, ARMA, and ARIMA models are useful for describing and analyzing individual time series. However, in the

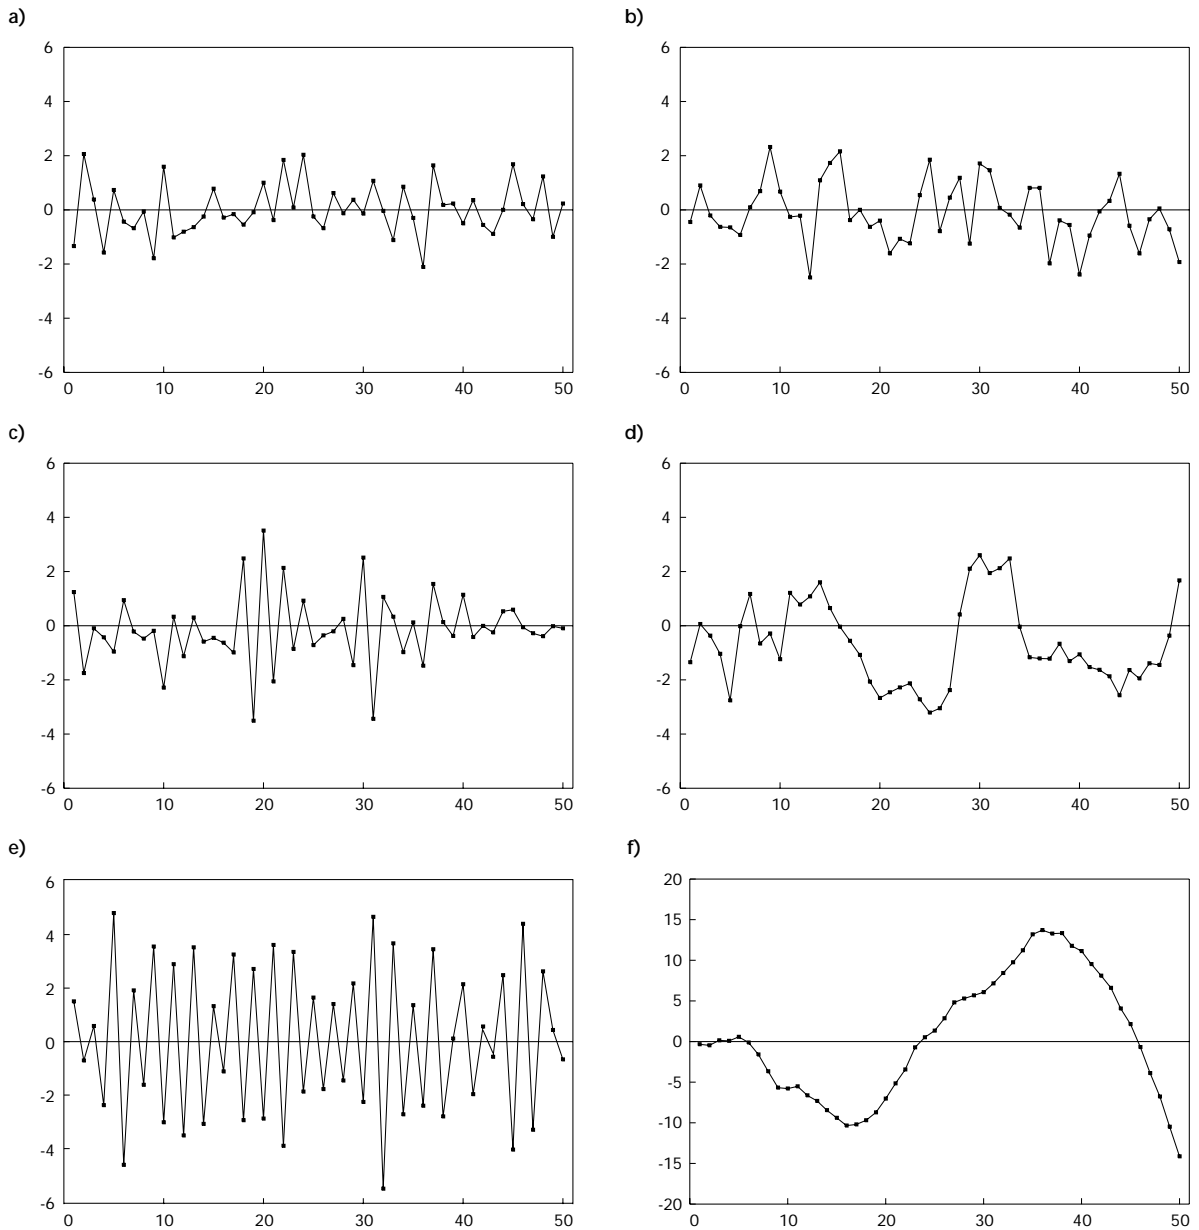


FIGURE 13 Time series generated by AR, MA, ARMA, and ARIMA models
 (a) $y_t = \varepsilon_t - 0.5 \varepsilon_{t-1}$, (b) $y_t = \varepsilon_t + 0.5 \varepsilon_{t-1}$, (c) $y_t = -0.8y_{t-1} + \varepsilon_t$, (d) $y_t = 0.8y_{t-1} + \varepsilon_t$,
 (e) $y_t = -0.8y_{t-1} - 0.5\varepsilon_{t-1} + \varepsilon_t$, and (f) $y_t - y_{t-1} = 0.8 (y_{t-1} - y_{t-2}) + \varepsilon_t$.

tree-ring examples described in Sections 1.1.2 and 1.1.3, and in many other situations, the goal is to relate one series to another. This requires a special type of time-series regression model known as a *transfer function model*

$$y_t = v + \sum_{i=0}^k \beta_i x_{t-i} + \varepsilon_t$$

in which the response variables $\{y_t\}$ and explanatory variables $\{x_t\}$ are both cross-correlated and autocorrelated. Readers who are interested in the identification, estimation, and interpretation of transfer function models should consult Chapters 11 and 12 of Kendall and Ord (1990) for more information.

Time series often arise as a collection of two or more time series. For instance, in the missing tree-ring problem (Section 1.1.2), each disk has an associated series of ring widths. In such situations, it seems natural to attempt to analyze the time series simultaneously by fitting a *multivariate model*. Multivariate AR, MA, ARMA, and ARIMA models have been developed for this purpose. They are, however, considerably more complex than their univariate counterparts. The reader should refer to Section 11.9 of Chatfield (1992) for an outline of difficulties and to Chapter 14 of Kendall and Ord (1990) for a description of the methods that can be used to identify and estimate multivariate models.

Other time-series models include *state-space models*, which are equivalent to multivariate ARIMA models and are useful for representing dependencies among one or more time series (see Chapter 9 of Kendall and Ord [1990] or Chapter 10 of Chatfield [1992]), and *intervention models*, useful for describing sudden changes in a time series, such as a disruption in growth caused by an unexpected drought, fire, or insect infestation (see Chapter 13 of Kendall and Ord [1990]). The reader should consult the appropriate reference for more information about these and other topics that are well beyond the introduction that this handbook is intended to provide.

3.6 Forecasting

Forecasting is the prediction of a future value y_{n+k} from a series of n previous values y_1, y_2, \dots, y_n . There are three general strategies for producing a forecast: (1) extrapolation of a deterministic trend, (2) exponential smoothing, and (3) the Box-Jenkins method. The preferred method depends on the properties of the time series (e.g., the presence or absence of a trend or seasonal component), the sample size n , the *lead time* k (i.e., the number of steps into the future for which the forecast is needed), and the required level of precision. If a time series is dominated by a deterministic trend, then the first method might be appropriate. On the other hand, this method sometimes produces unrealistic forecasts, in part, because it gives equal weight to current and past observations, even though the latter are generally less useful for predicting the future than the former. Exponential smoothing can be used to extrapolate short-term stochastic trends, as well as seasonal components. It is simple to use and automatically discounts remote observations. The Box-Jenkins method generates forecasts by fitting an ARIMA or seasonal ARIMA model to the data. It has considerable versatility, but is more difficult to apply than the other two methods because a suitable model must be identified and fitted.

All three methods have a subjective element, either in the selection of a model or in determination of the appropriate amount of smoothing. Various automatic forecasting methods have been developed in an attempt to eliminate this subjectivity (e.g., stepwise autoregression is a type of automatic Box-Jenkins procedure). More information about forecasting

can be found in McCleary and Hay (1980), Levenbach and Cleary (1981), Kendall et al. (1983), or Chatfield (1992).

4 REPEATED-MEASURES AND TIME-SERIES ANALYSIS WITH SAS

The SAS package is equipped to carry out both repeated-measures and time-series analyses. Repeated-measures analysis is available in the statistics module SAS/STAT (SAS Institute 1989). Procedures for time-series analysis are collected together in the econometric and time-series module SAS/ETS (SAS Institute 1991a).

4.1 Repeated-measures Analysis

A repeated-measures analysis is a special type of ANOVA, which is requested with a `REPEATED` statement in the general linear model procedure, `PROC GLM`, of SAS/STAT (SAS Institute 1989). The `REPEATED` statement performs a univariate or multivariate analysis, or both. If the design is balanced (i.e., the sample sizes are equal for all groups) and the residuals are not required, the same analyses can also be performed with `PROC ANOVA`.

4.1.1 Repeated-measures data sets Repeated-measures data sets have a univariate or multivariate structure. For the seedling experiment (Section 1.1.1), consider only the lodgepole pine seedlings grown from plugs. In this case, the experiment reduces to a simple randomized block design with measurements repeated on seven occasions. The objective is to compare the growth curves for the three site-preparation treatments (Figure 1d) by carrying out a repeated-measures analysis of the row means. If measurements made in successive years are treated as separate observations, then the data set for the seedling example has the following univariate structure (i.e., there is one response variable):

TRT	BLK	YEAR	HT
S	1	1983	12.43
S	1	1984	23.19
S	1	1985	36.71
S	1	1986	55.29
S	1	1987	75.71
S	1	1988	109.48
S	1	1989	155.76
.	.	.	.
.	.	.	.
.	.	.	.
V	4	1983	7.83
V	4	1984	13.58
V	4	1985	30.00
V	4	1986	53.42
V	4	1987	84.71
V	4	1988	130.38
V	4	1989	186.21

This data set has a total of 84 observations (3 treatments \times 4 blocks \times 7 years) with one response variable (HT = average height of seedlings) for each observation (row and year).

Alternatively, the same data can be converted to a multivariate data set with 12 observations (rows of seedlings) and seven response variables (average planting height, [PHT], and the average heights for 1984–89, [HT84, HT85 . . . , HT89]) for each row:

TRT	BLK	PHT	HT84	HT85	HT86	HT87	HT88	HT89
S	1	12.43	23.19	36.71	55.29	75.71	109.48	155.76
S	2	10.23	18.59	33.91	53.59	74.09	108.27	150.64
S	3	9.59	17.82	32.05	49.86	69.50	97.59	133.55
S	4	13.48	21.70	34.26	48.22	73.39	103.83	141.48
U	1	12.00	22.86	34.38	49.00	71.10	105.05	148.71
U	2	9.43	17.14	30.10	43.33	60.95	87.24	125.67
U	3	8.15	15.95	28.60	39.65	58.75	89.00	129.40
U	4	8.75	15.70	27.45	42.55	58.45	85.55	123.85
V	1	12.28	19.52	33.12	55.12	89.24	136.16	193.56
V	2	9.57	17.13	28.74	46.65	74.00	114.22	163.13
V	3	10.25	17.83	29.38	48.00	78.88	116.29	161.50
V	4	7.83	13.58	30.00	53.42	84.71	130.38	186.21

With SAS, the univariate data set (UVDATA) can be readily transformed to the multivariate form (MVDATA), as illustrated below:

```
PROC SORT DATA=UVDATA;
  BY TRT BLK;
DATA MVDATA(KEEP=TRT BLK PHT HT84-HT89);
  ARRAY H(7) PHT HT84-HT89;
  DO YEAR=1983 TO 1989;
    SET UVDATA;
    BY TRT BLK;
    H(YEAR-1982)=HT;
    IF LAST.BLK THEN RETURN;
  END;
RUN;
```

The reverse transformation can be achieved with the following statements:

```
DATA UVDATA(KEEP=TRT BLK YEAR HT);
  ARRAY H(7) PHT HT84-HT89;
  SET MVDATA;
  DO YEAR=1983 TO 1989;
    HT=H(YEAR-1982);
    OUTPUT;
  END;
RUN;
```

4.1.2 Univariate analysis A univariate repeated-measures ANOVA is requested in PROC GLM (or PROC ANOVA) by supplying the necessary MODEL and TEST statements, if the input data set has a univariate structure, or by replacing the TEST statement with a REPEATED statement, if the data set is multivariate. The two methods are illustrated below for the seedling data.

For the univariate data set, the SAS statements are:

```
PROC GLM DATA=UVDATA;
  TITLE1 'Univariate Repeated-Measures Analysis';
  TITLE2 'Method 1: univariate data set analyzed with TEST statement';
  CLASS BLK TRT YEAR;
  MODEL HT=BLK TRT BLK*TRT YEAR YEAR*TRT YEAR*BLK;
  TEST H=TRT E=BLK*TRT;
  TEST H=YEAR E=BLK*YEAR;
RUN;
```

Here the repeated-measures analysis is performed by specifying the between- and within-units factors with the MODEL statement. The correct error terms for testing the between-row effect of treatment and the within-row effect of year are selected with the TEST statement. Alternatively, the two TEST statements can be replaced with a single RANDOM statement with the TEST option:

```
RANDOM BLK BLK*TRT YEAR*BLK/TEST;
```

This calculates expected mean squares and automatically selects the correct error term for each effect, which is particularly useful if the investigator is uncertain about which error term to use. Application of the RANDOM statement is described in more detail in the SAS/STAT manual (SAS Institute 1989) and in Sit (1992b).

The results of the analysis of the univariate data set are displayed in Figure 14. Notice that the TRT*YEAR interaction is highly significant (**A**), so it is reasonable to conclude that the growth curves for the three treatments (Figure 1d) are not parallel. The tests of H_{02} —no growth over the seven-year study period—and H_{03} —no overall difference between treatments (after averaging over the seven years)—are labelled **B** and **C**, respectively. Both are highly significant, although the first is not particularly interesting in this example because the seedlings have obviously grown in height.

Results from a more informative univariate analysis, including an investigation of polynomial trends, are readily obtained from an analysis of the multivariate data set. The required SAS statements are:

```
PROC GLM DATA=MVDATA;
  TITLE1 'Univariate Repeated-Measures Analysis';
  TITLE2 'Method 2: multivariate data set analyzed with REPEATED statement';
  CLASS BLK TRT;
  MODEL PHT HT84-HT89=BLK TRT;
  REPEATED YEAR 7 (1983 1984 1985 1986 1987 1988 1989) POLYNOMIAL/SUMMARY
  PRINTM NOM;
RUN;
```

Univariate Repeated-Measures Analysis
Method 1: univariate data set analyzed with TEST statement

General Linear Models Procedure

Dependent Variable: HT

Source	DF	Sum of Squares	Mean Square	F Value	Pr > F
Model	47	198236.399	4217.796	290.70	0.0001
Error	36	522.326	14.509		
Corrected Total	83	198758.725			
	R-Square	C.V.	Root MSE		HT Mean
	0.997372	6.067216	3.80908		62.7813

Source	DF	Type III SS	Mean Square	F Value	Pr > F
BLK	3	1206.093	402.031	27.71	0.0001
TRT	2	3032.285	1516.143	104.50	0.0001
BLK*TRT	6	340.388	56.731	3.91	0.0041
YEAR	6	188835.114	31472.519	2169.16	0.0001
TRT*YEAR	12	4206.553	350.546	24.16	<u>0.0001</u> (A)
BLK*YEAR	18	615.966	34.220	2.36	0.0140

Tests of Hypotheses using the Type III MS for BLK*TRT as an error term

Source	DF	Type III SS	Mean Square	F Value	Pr > F
TRT	2	3032.28507	1516.14254	26.72	<u>0.0010</u> (B)

Tests of Hypotheses using the Type III MS for BLK*YEAR as an error term

Source	DF	Type III SS	Mean Square	F Value	Pr > F
YEAR	6	188835.114	31472.519	919.70	<u>0.0001</u> (C)

FIGURE 14 Univariate repeated-measures analysis of seedling data: univariate data set.

When the multivariate data set is used, the repeated factor (YEAR) and its levels (1983, 1984, . . . , 1989) are specified in the REPEATED statement and the respective response variables (PHT HT84-HT89) are identified in the MODEL statement. The REPEATED statement carries out both univariate and multivariate analyses, unless one or the other is suppressed with the NOU (no univariate analysis) or NOM (no multivariate analysis) option. In this example, NOM is chosen because only univariate analyses are required (the multivariate analysis is discussed in the next section). The POLYNOMIAL transformation and SUMMARY option are also

selected. The POLYNOMIAL transformation generates all possible (in this case, six) polynomial combinations of the repeated measures and the SUMMARY option produces an ANOVA of each (which are used to test hypothesis H_{06} of Section 2.1). The same results can be obtained, without the REPEATED statement, by first creating and then analyzing the transformed variables (POLY1, POLY2, . . . , POLY6) as follows:

```
DATA POLY;
  SET MVDATA;
  POLY1=-3*PHT-2*HT84 -HT85 +0*HT86 +HT87+2*HT88+3*HT89;
  POLY2= 5*PHT+0*HT84 -3*HT85 -4*HT86 -3*HT87+0*HT88+5*HT89;
  POLY3=-1*PHT+1*HT84 +1*HT85 +0*HT86 -1*HT87-1*HT88+1*HT89;
  POLY4= 3*PHT-7*HT84 +1*HT85 +6*HT86 +1*HT87-7*HT88+3*HT89;
  POLY5=-1*PHT+4*HT84 -5*HT85 +0*HT86 +5*HT87-4*HT88+1*HT89;
  POLY6= 1*PHT-6*HT84+15*HT85-20*HT86+15*HT87-6*HT88+1*HT89;
PROC GLM DATA=POLY;
  CLASS BLK TRT;
  MODEL POLY1-POLY6=BLK TRT/SS3;
RUN;
```

Notice that the coefficients (-3, -2, -1, 0, 1, 2, 3) are the same as those used to define polynomial contrasts in an ANOVA, and can be found in such textbooks as Keppel (1973). Other linear combinations can be selected by making the appropriate substitution for POLYNOMIAL in the REPEATED statement: PROFILE for successive differences (refer to H_{04} of Section 2.1) and CONTRAST⁶ for cumulative differences (refer to H_{05} of Section 2.1). The PRINTM option can be added to the REPEATED statement to verify that the correct transformation has been selected. This causes the contrast coefficients to be printed.

When the REPEATED statement is applied to a randomized block design, or any other design with random effects (e.g., designs with subsampling), a complication arises. Because the REPEATED statement assumes that all effects are fixed, the residual mean square is used as the error term for all within-units tests. Consequently, the F -ratio for testing the main effect of time (YEAR) is incorrect. In the seedling example, the correct F -ratio can be calculated by dividing the mean square for time (YEAR) by the proper error term, which is the mean square for the time by block interaction (YEAR*BLK). The same problem occurs in the ANOVA of the transformed data (i.e., POLY1, POLY2, etc.), where the F -ratio for testing the statistical significance of the overall mean of each variable (labelled MEAN in the ANOVA tables produced by the REPEATED statement) is incorrect. In the seedling example, the correct F -tests for the polynomial contrasts can be obtained by inserting the following MANOVA statement after the REPEATED statement:

6 The CONTRAST transformation of the REPEATED statement is not to be confused with the CONTRAST statement of PROC GLM.

```
MANOVA H=INTERCEPT E=BLK M=( -3 -2 -1 0 1 2 3, 5 0 -3 -4 -3 0 5,
-1 1 1 0 -1 -1 1, 3 -7 1 6 1 -7 3,
-1 4 -5 0 5 -4 1, 1 -6 15 -20 15 -6 1)
MNames=POLY1 POLY2 POLY3 POLY4 POLY5 POLY6/ORTH SUMMARY;
```

A MANOVA statement is a multivariate analog of the TEST statement. It is used whenever a test involves the elements of a multivariate array (in this case, the repeated measures). The H=INTERCEPT option requests a test of the intercept or overall mean; the E=BLK option identifies BLK as the correct error matrix (this is a matrix of sums of squares and cross-products and is the multivariate equivalent of the univariate YEAR*BLK sum of squares); M=(-3 -2 -1 0 1 2 3, etc.) defines the transformed variables to be analyzed (in this case polynomial contrasts); and MNames=POLY1, etc., names the transformed variables. Proper use of the MANOVA statement requires a good understanding of multivariate models and multivariate hypothesis testing. Refer to Gumpertz and Brownie (1993) for a complete description of the analysis of randomized block designs with repeated measures and the use of MANOVA statements to obtain the correct analysis in SAS.

An edited version of the output from the second univariate analysis of the seedling data is shown in Figure 15. It has five main parts:

- an ANOVA for each year;
- a repeated-measures ANOVA of the between-rows (“Between Subjects”) variation;
- a repeated-measures ANOVA of the within-rows (“Within Subjects”) variation;
- an ANOVA of each of the six polynomial contrasts among years; and
- the correct *F*-tests for the overall means of the polynomial contrasts, which are labelled **A**, **B**, **C**, **D**, and **E**, respectively.

The first part of the output (**A**) gives the ANOVA results for each year. For brevity, only the results for 1989 are shown in Figure 15. Notice that there was a significant difference (**A.1**) between the three treatment groups in that year. The results for the other years (not shown) suggest that there were no significant differences between groups from the time of planting up to the end of the first year, but by 1985 and in subsequent years, there were significant differences between the groups.

The test for an overall treatment effect in the between-rows analysis (**B**) is the same as the TRT test (**C**) in Figure 14. Likewise, the test for a treatment effect in the within-rows analysis (**C.3**) is the same as the TRT*YEAR test (**A**) in Figure 14. However, notice that the within-rows test of year (**C.1**), which uses the incorrect error term, is not the same as the corresponding test (**B**) in Figure 14. The correct *F*-statistic for the year effect ($F = 919.70$) can be calculated by dividing the mean square for YEAR, which is 31472.5190 (**C.1**), by the mean square for YEAR*BLK, which is 34.2203 (**C.2**).

At the end of the within-units analysis there are two numbers labelled “Greenhouse-Geisser Epsilon” and “Huynh-Feldt Epsilon” (**C.4**). If the variances and correlations of the repeated measures are constant over time (as required by the univariate repeated-measures analysis), then these two

Univariate Repeated-Measures Analysis
Method 2: multivariate data set analyzed with REPEATED statement

General Linear Models Procedure

Dependent Variable: HT89 (A)

Source	DF	Sum of Squares	Mean Square	F Value	Pr > F
Model	5	5115.06413	1023.01283	13.29	0.0034
Error	6	461.95544	76.99257		
Corrected Total	11	5577.01957			
	R-Square	C.V.	Root MSE	HT89 Mean	
	0.917168	5.806303	8.77454	151.121	

Source	DF	Type III SS	Mean Square	F Value	Pr > F
BLK	3	1009.76203	336.58734	4.37	0.0591
TRT	2	4105.30210	2052.65105	26.66	<u>0.0010</u> (A.1)

Tests of Hypotheses for Between Subjects Effects

Source	DF	Type III SS	Mean Square	F Value	Pr > F
BLK	3	1206.0928	402.0309	7.09	0.0213
TRT	2	3032.2851	1516.1425	26.72	<u>0.0010</u> (B)
Error	6	340.3883	56.7314		

Univariate Tests of Hypotheses for Within Subject Effects (C)

Source: YEAR

DF	Type III SS	Mean Square	F Value	Pr > F	Adj G - G	Pr > F H - F
6	188835.113904	<u>31472.518984</u>	2169.16	0.0001	0.0001	0.0001 (C.1)

Source: YEAR*BLK

DF	Type III SS	Mean Square	F Value	Pr > F	Adj G - G	Pr > F H - F
18	615.966173	<u>34.220343</u>	2.36	0.0140	0.1482	0.0673 (C.2)

Source: YEAR*TRT

DF	Type III SS	Mean Square	F Value	Pr > F	Adj G - G	Pr > F H - F
12	4206.552681	350.546057	24.16	0.0001	0.0004	0.0001 (C.3)

FIGURE 15 Univariate repeated-measures analysis of seedling data: multivariate data set.

Source: Error (YEAR)

DF	Type III SS	Mean Square
36	522.326420	14.509067

Greenhouse-Geisser Epsilon = 0.2065 (C.4)
 Huynh-Feldt Epsilon = 0.4505

Analysis of Variance of Contrast Variables (D)

YEAR.N represents the nth degree polynomial contrast for YEAR

Contrast Variable: YEAR.1

Source	DF	Type III SS	Mean Square	F Value	Pr > F
MEAN	1	175659.5789	175659.5789	<u>2455.33</u>	0.0001 (D.1)
BLK	3	517.2866	172.4289	2.41	0.1654
TRT	2	3578.8219	1789.4110	25.01	0.0012
Error	6	429.2521	71.5420		

Contrast Variable: YEAR.2

Source	DF	Type III SS	Mean Square	F Value	Pr > F
MEAN	1	12900.59010	12900.59010	1532.32	0.0001
BLK	3	83.10595	27.70198	3.29	0.0999
TRT	2	587.24391	293.62195	34.88	0.0005
Error	6	50.51387	8.41898		

Contrast Variable: YEAR.3

Source	DF	Type III SS	Mean Square	F Value	Pr > F
MEAN	1	262.0476732	262.0476732	143.48	0.0001
BLK	3	10.5190901	3.5063634	1.92	0.2276
TRT	2	27.1102199	13.5551099	7.42	0.0239
Error	6	10.9585841	1.8264307		

Contrast Variable: YEAR.4

Source	DF	Type III SS	Mean Square	F Value	Pr > F
MEAN	1	0.84138232	0.84138232	0.28	0.6149
BLK	3	2.76045704	0.92015235	0.31	0.8195
TRT	2	11.52322226	5.76161113	1.93	0.2258
Error	6	17.94389412	2.99064902		

FIGURE 15 (continued)

Contrast Variable: YEAR.5

Source	DF	Type III SS	Mean Square	F Value	Pr > F
MEAN	1	11.63080680	11.63080680	16.80	0.0064
BLK	3	0.58185855	0.19395285	0.28	0.8381
TRT	2	0.88502506	0.44251253	0.64	0.5601 (D.2)
Error	6	4.15295236	0.69215873		

Contrast Variable: YEAR.6

Source	DF	Type III SS	Mean Square	F Value	Pr > F
MEAN	1	0.42502079	0.42502079	0.27	0.6230
BLK	3	1.71225628	0.57075209	0.36	0.7843
TRT	2	0.96838905	0.48419453	0.31	0.7475 (D.3)
Error	6	9.50504759	1.58417460		

Dependent Variable: POLY1

Tests of Hypotheses using the Type III MS for BLK as an error term

Source	DF	Type III SS	Mean Square	F Value	Pr > F (E)
INTERCEPT	1	175659.5789	175659.5789	<u>1018.74</u>	0.0001 (E.1)
Error	3	517.2866	172.4289		

Dependent Variable: POLY2

Tests of Hypotheses using the Type III MS for BLK as an error term

Source	DF	Type III SS	Mean Square	F Value	Pr > F
INTERCEPT	1	12900.59010	12900.59010	465.69	0.0002
Error	3	83.10595	27.70198		

Dependent Variable: POLY3

Tests of Hypotheses using the Type III MS for BLK as an error term

Source	DF	Type III SS	Mean Square	F Value	Pr > F
INTERCEPT	1	262.0476732	262.0476732	74.73	0.0033
Error	3	10.5190901	3.5063634		

FIGURE 15 (continued)

Dependent Variable: POLY4

Tests of Hypotheses using the Type III MS for BLK as an error term

Source	DF	Type III SS	Mean Square	F Value	Pr > F
INTERCEPT	1	0.84138232	0.84138232	0.91	0.4095
Error	3	2.76045704	0.92015235		

Dependent Variable: POLY5

Tests of Hypotheses using the Type III MS for BLK as an error term

Source	DF	Type III SS	Mean Square	F Value	Pr > F
INTERCEPT	1	11.63080680	11.63080680	59.97	0.0045 (E.2)
Error	3	0.58185855	0.19395285		

Dependent Variable: POLY6

Tests of Hypotheses using the Type III MS for BLK as an error term

Source	DF	Type III SS	Mean Square	F Value	Pr > F
INTERCEPT	1	0.42502079	0.42502079	0.74	0.4516 (E.3)
Error	3	1.71225628	0.57075209		

FIGURE 15 (concluded).

numbers should be approximately equal to one. When this assumption is violated, the numbers will be substantially less than one and the F -tests for the within-units analysis will be incorrect. To compensate for this problem, SAS uses the Greenhouse-Geisser epsilon and the Huynh-Feldt epsilon to make two separate adjustments to the p -value for each F -test. The adjusted p -values ("Adj Pr>F") are labelled "G - G" (Greenhouse-Geisser correction) and "H - F" (Huynh-Feldt correction) in the output. More information about these adjustments can be found in Chapter 27 of Milliken and Johnson (1992).

The ANOVA tables for the six polynomial contrasts (**D**) are labelled YEAR.1 (linear trend), YEAR.2 (quadratic trend), etc., in Figure 15. As discussed previously, the test for a MEAN effect is incorrect in these tables and should be replaced with the corresponding MANOVA test (i.e., the INTERCEPT test for POLY1=YEAR.1, POLY2=YEAR.2, etc., which is given in **E**). For example, the correct F -ratio for testing the overall mean of the linear transformation (YEAR.1 in **D** and POLY1 in **E**) is 1018.74 (**E.1**), not 2455.33 (**D.1**). Notice that treatment has no significant effect

on the sixth-order polynomial contrasts (i.e., there is no significant difference between β_6 for the three site-preparation treatments in **D.3**). Moreover, the average of these three coefficients is not significantly different from zero (**E.3**). This suggests that the growth curves for the three treatment groups are adequately described by polynomials of order no greater than five (i.e., $\beta_6 = 0$). Examining the results for the fifth-order polynomial transformation (YEAR.5 and POLY5) reveals that there is a fifth-order term in the trends (**E.2**), but it is not significantly different for the three treatment groups (**D.2**). The same conclusion holds for the fourth-order term. Site-preparation does, however, have a significant effect on the cubic and lower-order terms. Thus the evidence suggests that the growth curves can be described by fifth-order polynomials with the effects of site-preparation treatment limited to the cubic and lower-order terms.

4.1.3 Multivariate analysis The REPEATED statement automatically performs a multivariate repeated-measures analysis (unless the NOM option is used). The SAS statements for carrying out a multivariate repeated-measures analysis of the seedling data set are:

```
PROC GLM DATA=MVDATA;
  TITLE 'Multivariate Repeated-Measures Analysis';
  CLASS BLK TRT;
  MODEL PHT HT84-HT89=BLK TRT/SS3;
  REPEATED YEAR 7 (1983 1984 1985 1986 1987 1988 1989)/NOU;
RUN;
```

In this case, the NOU (no univariate analysis) option is used to suppress the univariate tests of the within-units effects. These are replaced by multivariate tests, the results of which are displayed in Figure 16. Recall that SAS does not recognize that block is a random effect. Therefore, the multivariate test of the main effect of year (labelled **A** in Figure 16), like the corresponding univariate test, is incorrect. If there were seven or more blocks, then the following MANOVA statement would produce the correct test:

```
MANOVA H=INTERCEPT E=BLK M=( 1 -1 0 0 0 0 0, 0 1 -1 0 0 0 0,
                                0 0 1 -1 0 0 0, 0 0 0 1 -1 0 0,
                                0 0 0 0 1 -1 0, 0 0 0 0 0 1 -1);
```

However, for this particular design, which has fewer blocks than years, there are too few degrees of freedom to carry out the test. Fortunately, the test for no time effect is not very important in this example.

The results of the multivariate test for parallel growth curves (**B** in Figure 16) should be compared with the corresponding results of the univariate tests (**A** in Figure 14 and **C.1** in Figure 15). Notice that the univariate tests suggest that the growth curves are not parallel for the three groups (i.e., the YEAR*TRT effect is significant), but that two of the three multivariate tests (the sample sizes are too small to compute the

Multivariate Repeated-Measures Analysis

Manova Test Criteria and Exact F Statistics for
the Hypothesis of no YEAR Effect

(A)

H = Type III SS&CP Matrix for YEAR E = Error SS&CP Matrix

S=1 M=2 N=-0.5

Statistic	Value	F	Num DF	Den DF	Pr F
Wilks' Lambda	0.00025957	641.9264	6	1	0.0302
Pillai's Trace	0.99974043	641.9264	6	1	0.0302
Hotelling-Lawley Trace	3851.5581633	641.9264	6	1	0.0302
Roy's Greatest Root	3851.5581633	641.9264	6	1	0.0302

Manova Test Criteria and F Approximations for
the Hypothesis of no YEAR*BLK Effect

H = Type III SS&CP Matrix for YEAR*BLK E = Error SS&CP Matrix

S=3 M=1 N=-0.5

Statistic	Value	F	Num DF	Den DF	Pr > F
Wilks' Lambda	0.01032163	0.7433	18	3.313708	0.7084
Pillai's Trace	1.65054752	0.6116	18	9	0.8210
Hotelling-Lawley Trace	40.36355908	.	18	-1	.
Roy's Greatest Root	39.16883650	19.5844	6	3	0.0167

NOTE: F Statistic for Roy's Greatest Root is an upper bound.

Manova Test Criteria and F Approximations for
the Hypothesis of no YEAR*TRT Effect

(B)

H = Type III SS&CP Matrix for YEAR*TRT E = Error SS&CP Matrix

S=2 M=1.5 N=-0.5

Statistic	Value	F	Num DF	Den DF	Pr > F
Wilks' Lambda	0.00172637	3.8446	12	2	0.2248
Pillai's Trace	1.73572577	2.1893	12	4	0.2340
Hotelling-Lawley Trace	151.08093445	.	12	0	.
Roy's Greatest Root	148.19852119	49.3995	6	2	0.0200

NOTE: F Statistic for Roy's Greatest Root is an upper bound.

NOTE: F Statistic for Wilks' Lambda is exact.

FIGURE 16 Multivariate repeated measures analysis of seedling data.

fourth statistic) fail to reject the hypothesis of parallel trends. The discrepancy can probably be explained by the less restrictive assumptions of the multivariate model and the resultant loss of power. In such cases, it is usually advisable to use the adjusted p -values based on the univariate tests.

For more information and examples of repeated-measures analysis with SAS, the reader is referred to Cody and Smith (1987), Hand and Taylor (1987), Littell (1989), Tabachnick and Fidell (1989), Milliken and Johnson (1992), and Nemeč (1992).

4.2 Time-series Analysis

The SAS/BASE (SAS Institute 1990a) and SAS/STAT (SAS Institute 1989) libraries are not very useful for time-series analysis. Apart from low resolution time plots (PROC TIMEPLOT) and the Durbin-Watson statistic (an option in PROC REG), all time-series procedures are in the SAS/ETS module (SAS Institute 1991a).

The SAS/ETS library has three main procedures for analyzing a time series in the time domain: PROC ARIMA, PROC AUTOREG, and PROC FORECAST. The ARIMA procedure computes the sample ACF, PACE, IACF, and CCF (refer to Sections 3.2.2 and 3.5.1 for definitions). It also fits and forecasts AR, MA, ARMA, and ARIMA models (Section 3.5.1), as well as transfer function and intervention models (Section 3.5.2). The AUTOREG procedure fits and forecasts multiple regression models with autocorrelated errors (i.e., errors generated by an autoregression). It can also be used to fit pure AR models. The last of the three procedures, PROC FORECAST, uses an autoregressive method or exponential smoothing (nonseasonal and seasonal versions) to forecast time series. The SAS/ETS library also includes a useful procedure, PROC EXPAND, for collapsing, interpolating, or otherwise manipulating time-series data.

The following sections give a brief description of PROC EXPAND, PROC ARIMA, PROC AUTOREG, and PROC FORECAST. Simple programs, with accompanying output, are provided to demonstrate their application. For more details of these and other time-series procedures, and additional examples, the reader is referred to the SAS/ETS User's Guide (SAS Institute 1991a) and the SAS/ETS Software: Applications Guide 1 (SAS Institute 1991b).

4.2.1 Time-series data sets Time-series data sets consist of a time variable and the corresponding values of one or more response variables. The data set for the daily soil temperatures depicted in Figure 9a is as follows:

Daily Soil Temperatures			
OBS	DAY		TSOIL
1	Sun,	May 1, 1988	6.000
2	Mon,	May 2, 1988	5.800
3	Tue,	May 3, 1988	6.000
4	Wed,	May 4, 1988	6.500
5	Thu,	May 5, 1988	7.200
6	Fri,	May 6, 1988	7.200
7	Sat,	May 7, 1988	7.800
8	Sun,	May 8, 1988	8.800

9	Mon, May 9, 1988	9.500
10	Tue, May 10, 1988	10.500
11	Wed, May 11, 1988	10.900
12	Thu, May 12, 1988	11.500

.	.	.
.	.	.
.	.	.

(more data)

In this case, the time index, `DAY`, is an SAS date variable (printed in `WEEKDATE17.` format) and the response variable, `TSOIL`, is the soil temperature in degrees Celsius. Indexing a time series with an SAS date variable facilitates data verification and manipulation with SAS/ETS procedures but is not essential (e.g., the variable `DAY` could be omitted from the previous data set).

For `PROC ARIMA`, `PROC AUTOREG`, and `PROC FORECAST`, the input data set must be sorted in order of ascending time and the series must be equally spaced in time. If the data are not equally spaced, then `PROC EXPAND` can be used to convert the series, by interpolation, to equal time intervals. The same procedure can also be used to collapse or expand the sampling interval (e.g., by summing or averaging data points), to replace missing values with interpolated values, or to perform other operations on the series. As a simple illustration, consider the soil temperature series. To convert the daily series to weekly averages, `PROC EXPAND` can be used as follows:

```
PROC EXPAND DATA=DAILY OUT=WEEKLY FROM=DAY TO=WEEK ;
  CONVERT TSOIL/METHOD=AGGREGATE OBSERVED=AVERAGE ;
  ID DATE ;
PROC PRINT DATA=WEEKLY ;
  TITLE 'Weekly Averages' ;
RUN ;
```

The input series (`DAILY`), part of which was listed at the beginning of this section, and the output series (`WEEKLY`) are specified by the `DATA` and `OUT` options in the `PROC EXPAND` line. The change of intervals from daily to weekly is defined with the `FROM` and `TO` options, and the type of conversion is selected with the `METHOD` and `OBSERVED` options of the `CONVERT` statement (i.e., `METHOD=AGGREGATE` specifies that the output series is an aggregation of the input series and `OBSERVED=AVERAGE` specifies that the aggregation is by averaging). The resultant weekly series is listed below.

Weekly Averages			
OBS	DATE		TSOIL
1	Sun, 1	May 88	6.6429
2	Sun, 8	May 88	10.2143
3	Sun, 15	May 88	9.0000
4	Sun, 22	May 88	9.5286

5	Sun, 29	May 88	9.3714
6	Sun, 5	Jun 88	10.2571
7	Sun, 12	Jun 88	12.0000
8	Sun, 19	Jun 88	12.1286
9	Sun, 26	Jun 88	12.6000
10	Sun, 3	Jul 88	12.1000
11	Sun, 10	Jul 88	12.3857
12	Sun, 17	Jul 88	12.5286
13	Sun, 24	Jul 88	12.9143
14	Sun, 31	Jul 88	12.4143
15	Sun, 7	Aug 88	12.3286
16	Sun, 14	Aug 88	12.0286
17	Sun, 21	Aug 88	12.2857
18	Sun, 28	Aug 88	12.5286
19	Sun, 4	Sep 88	11.0714
20	Sun, 11	Sep 88	9.3857
21	Sun, 18	Sep 88	7.0000
22	Sun, 25	Sep 88	7.2000
23	Sun, 2	Oct 88	8.0571
24	Sun, 9	Oct 88	7.2429
25	Sun, 16	Oct 88	5.6714
26	Sun, 23	Oct 88	3.6429

The same procedure can be used to replace missing values in the daily series with interpolated values as follows:

```
PROC EXPAND DATA=DAILY OUT=INTERPOL FROM=DAY ;
  CONVERT TSOIL ;
  ID DATE ;
RUN ;
```

Since no `TO` option is supplied, the sampling frequency (i.e., one measurement per day) for the output series is the same as the input series and the only effect is the replacement of missing values. In this case, the interpolated values are obtained by fitting a cubic spline to the data (which is the default method). Other methods of interpolation can be selected with the `METHOD` option. After the missing points have been replaced, the data should be plotted (or re-plotted) to ensure that the interpolated values are consistent with the rest of the series. For more information about `PROC EXPAND` and its numerous capabilities, the reader is referred to the *SAS/ETS User's Guide* (SAS Institute 1991a).

The `TIMEPLOT` procedure in `SAS/BASE` is useful for making quick, low-resolution plots of time series. It is similar to `PROC PLOT`, except that the axes are reversed (i.e., the time axis runs vertically along the page, instead of horizontally). This is usually more convenient for plotting time series because they tend to have extended x -axes. The SAS statements for plotting a time series are illustrated below for the average weekly soil temperatures:

```

PROC TIMEPLOT DATA=WEEKLY;
  TITLE 'Average Weekly Soil Temperatures';
  PLOT TSOIL='*';
  ID YEAR;
RUN;

```

The output is shown in Figure 17.

4.2.2 PROC ARIMA The ARIMA procedure has three main statements: IDENTIFY, ESTIMATE, and FORECAST. The IDENTIFY statement appears first. It defines the time series to be analyzed and calculates summary statistics that are useful for model identification, including the sample ACF, IACF, and PACF, and if applicable, the sample CCF (refer to Sections 3.2.2 and 3.5.1 for definitions). Once a model has been identified, it can be fitted to the series by including an ESTIMATE statement

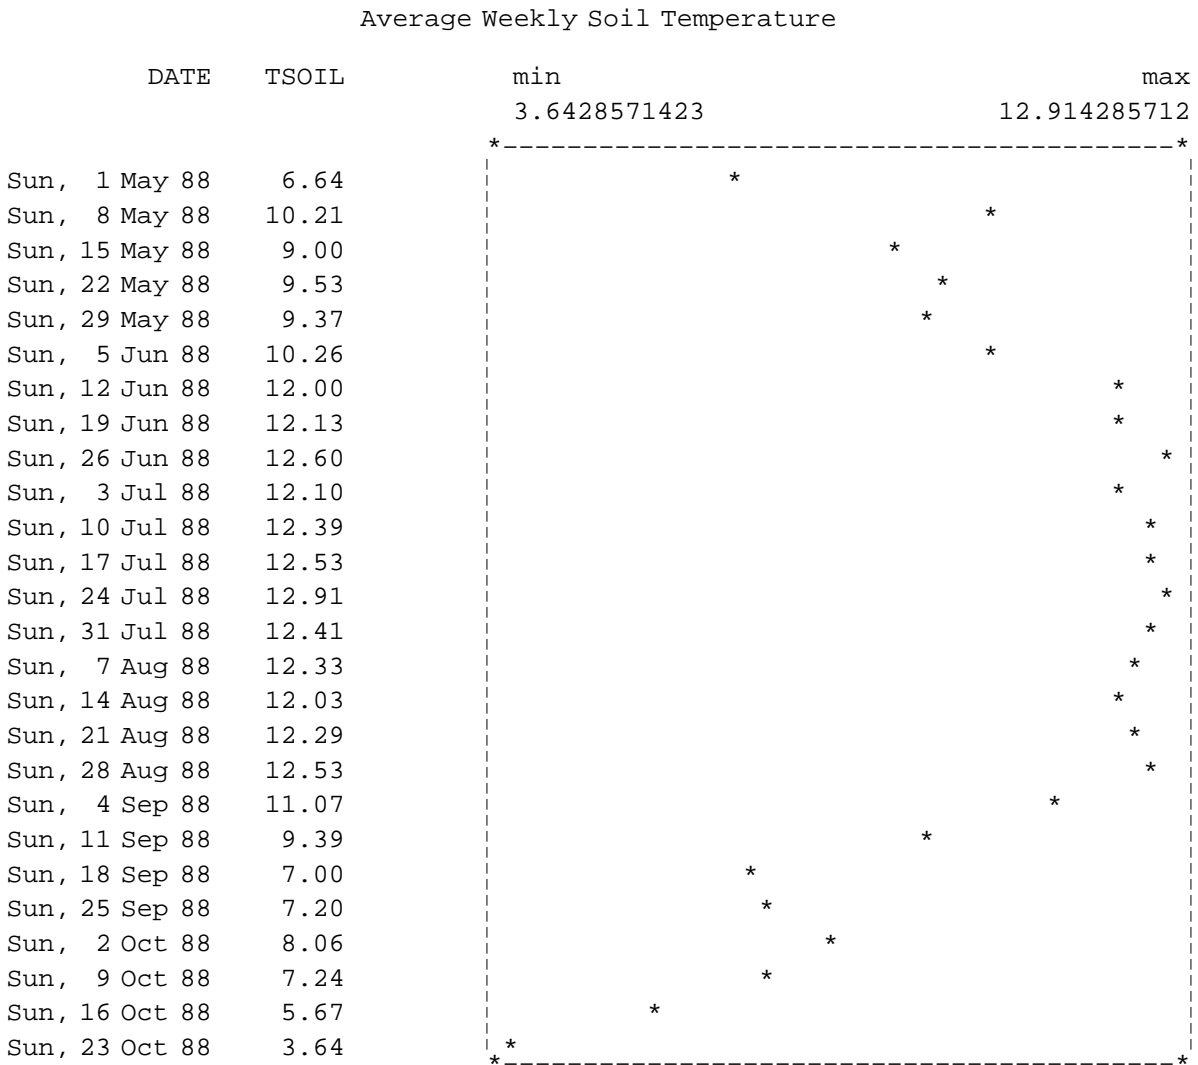


FIGURE 17 Time plot of weekly soil temperatures created with PROC TIMEPLOT.

after the IDENTIFY statement. The FORECAST statement generates predicted values (or forecasts) and occurs after the IDENTIFY and ESTIMATE statements.

To illustrate how the IDENTIFY, ESTIMATE, and FORECAST statements are used, consider the ring-index and rainfall example of Section 1.1.3. The sample ACF, IACF, and PACF of the ring-index series are requested with an IDENTIFY statement as follows:

```
DATA INDEXRAIN;
  INFILE 'INDEX.DAT' MISSOVER;
  INPUT YEAR INDEX RAIN;
  LABEL INDEX='Ring index' RAIN='Spring rainfall';
PROC ARIMA;
  TITLE1 'Time-Series Analysis: PROC ARIMA';
  TITLE2 'Identify model for ring-index series.';
  IDENTIFY VAR=INDEX;
RUN;
```

The VAR part of the IDENTIFY statement specifies the time series (INDEX) to be analyzed and is always required.

Figure 18 shows the resultant autocorrelations (A), inverse autocorrelations (B), and partial autocorrelations (C). Each set of autocorrelations is tabulated and plotted for lags 0 through 24 (for brevity, only lags 0–15 are listed in B and C). Also plotted are two dotted lines—one to the left and one to the right of zero. If the series is white noise, then individual values that lie outside these boundaries are more than two standard errors away from zero. Such values can be considered statistically significant at an approximately 5% level of significance. Notice that the sample ACF (A), which is the same as that illustrated in Figure 10a, tails off but shows no evidence of nonstationarity (because the ring index is a detrended value). Observe also that the sample IACF and PACF have a large, statistically significant spike at lag 1, a smaller spike at lag 6 or 7, and are approximately zero for all other lags. This behaviour is indicative of an autoregression with $p = 1, 6, \text{ or } 7$.

The last part (D) of the output shown in Figure 18, which is labelled “Autocorrelation Check for White Noise,” summarizes the results of a series of tests of randomness. The first row in the table is a test of the null hypothesis that the autocorrelation coefficients from lag 1 up to lag 6 are zero, the second row is a test of the same hypothesis for lags 1–12, and so on. The chi-squared statistic for each test is listed in the second column of the table, the degrees of freedom are in the third column, and the p -values are in the fourth. The last six columns are the estimated autocorrelation coefficients for lags 1–6 (first row), lags 7–12 (second row), etc., which are the same as the values given in the first part (A) of the output. Each null hypothesis is rejected if the corresponding p -value is small (i.e., no greater than 0.05 for a 5% level of significance). If the series is white noise, then none of the hypotheses should be rejected. In this case, all four p -values are less than 0.001 for the ring-index series, which suggests that there is some autocorrelation.

Time-Series Analysis: PROC ARIMA
Identify model for ring-index series.

ARIMA Procedure

Name of variable = INDEX.

Mean of working series = 0.99646
Standard deviation = 0.191788
Number of observations = 463

Autocorrelations

(A)

Lag	Covariance	Correlation	-1	9	8	7	6	5	4	3	2	1	0	1	2	3	4	5	6	7	8	9	1	
0	0.036782	1.00000												*****	*****	*****	*****	*****	*****	*****	*****	*****	*****	*****
1	0.026179	0.71173										.		*****	*****	*****	*****	*****	*****	*****	*****	*****	*****	*****
2	0.020030	0.54455										.		*****	*****	*****	*****	*****	*****	*****	*****	*****	*****	*****
3	0.015960	0.43391										.		*****	*****	*****	*****	*****	*****	*****	*****	*****	*****	*****
4	0.011755	0.31958										.		*****	*****	*****	*****	*****	*****	*****	*****	*****	*****	*****
5	0.009626	0.26170										.		*****	*****	*****	*****	*****	*****	*****	*****	*****	*****	*****
6	0.0092158	0.25055										.		*****	*****	*****	*****	*****	*****	*****	*****	*****	*****	*****
7	0.0057674	0.15680										.		***	***	***	***	***	***	***	***	***	***	***
8	0.0033671	0.09154										.		**	**	**	**	**	**	**	**	**	**	**
9	0.0024593	0.06686										.		*	*	*	*	*	*	*	*	*	*	*
10	0.0012791	0.03477										.		*	*	*	*	*	*	*	*	*	*	*
11	0.0014994	0.04076										.		*	*	*	*	*	*	*	*	*	*	*
12	0.0017450	0.04744										.		*	*	*	*	*	*	*	*	*	*	*
13	-0.000112	-0.00304									
14	-0.0014171	-0.03853									.	*
15	-0.0026482	-0.07200									.	*
16	-0.0047892	-0.13020									.	***
17	-0.0066828	-0.18168									***	***
18	-0.0076178	-0.20710									***	***
19	-0.0081882	-0.22261									***	***
20	-0.0079820	-0.21700									***	***
21	-0.0079143	-0.21516									***	***
22	-0.0089012	-0.24200									****	****
23	-0.0084721	-0.23033									****	****
24	-0.0093843	-0.25513									****	****

``.`` marks two standard errors

FIGURE 18 Time-series analysis of ring-index series: model identification.

Inverse Autocorrelations

(B)

Lag	Correlation	-1	9	8	7	6	5	4	3	2	1	0	1	2	3	4	5	6	7	8	9	1
1	-0.42101										*****											
2	-0.02881										*		.									
3	-0.05024										*		.									
4	0.03697										.		.									
5	0.07171										.		.									
6	-0.14301										***		.									
7	0.03444										.		*									
8	0.04396										.		*									
9	-0.00244										.		.									
10	0.03660										.		*									
11	-0.01465										.		.									
12	-0.06493										*		.									
13	0.02966										.		*									
14	0.01278										.		.									
15	-0.01721										.		.									

Partial Autocorrelations

(C)

Lag	Correlation	-1	9	8	7	6	5	4	3	2	1	0	1	2	3	4	5	6	7	8	9	1
1	0.71173										.		*****									
2	-0.07700										.		**									
3	0.04358										.		*									
4	-0.04365										*		.									
5	0.04527										.		*									
6	0.08370										.		**									
7	-0.12827										***		.									
8	-0.03453										*		.									
9	0.02383										.		.									
10	-0.00429										.		.									
11	0.04314										.		*									
12	0.00104										.		.									
13	-0.07924										**		.									
14	-0.03113										*		.									
15	-0.04762										*		.									

Autocorrelation Check for White Noise

(D)

To Lag	Chi Square	DF	Prob	Autocorrelations					
6	572.34	6	0.000	0.712	0.545	0.434	0.320	0.262	0.251
12	592.47	12	0.000	0.157	0.092	0.067	0.035	0.041	0.047
18	640.53	18	0.000	-0.003	-0.039	-0.072	-0.130	-0.182	-0.207
24	796.46	24	0.000	-0.223	-0.217	-0.215	-0.242	-0.230	-0.255

FIGURE 18 (continued).

The IDENTIFY statement generates summary statistics for differenced series, as well as untransformed series. To analyze a differenced series ($y_t - y_{t-k}$), the interval k is enclosed in parentheses after the variable name. If the series is differenced more than once (e.g., $[y_t - y_{t-1}] - [y_{t-1} - y_{t-2}]$), then the corresponding intervals must be separated by commas as illustrated below:

```
PROC ARIMA DATA=INDXRAIN;
  IDENTIFY VAR=INDEX(1);      /* First difference of INDEX series */
  IDENTIFY VAR=INDEX(1,1);    /* Second difference of INDEX series */
RUN;
```

Alternatively, the same results can be achieved by creating, with the LAG function (SAS Institute 1990b), the differenced series in a DATA step:

```
DATA DIFFS;
  INFILE 'INDEX.DAT' MISSOVER;
  INPUT YEAR INDEX RAIN;
  DIFF1=INDEX-LAG1(INDEX);
  DIFF2=INDEX-2*LAG1(INDEX)+LAG2(INDEX);
PROC ARIMA DATA=DIFFS;
  IDENTIFY VAR=DIFF1;        /* First difference of INDEX series */
  IDENTIFY VAR=DIFF2;        /* Second difference of INDEX series */
RUN;
```

The IDENTIFY statement has various options. The CROSSCOR option requests a cross-correlation analysis of the main series (i.e., the series defined by VAR) and one or more other series. The latter are identified by variable names enclosed in parentheses after CROSSCOR, with differenced series specified in the same manner as described above for the main series. The CROSSCOR option is demonstrated in the following example:

```
PROC ARIMA DATA=PWSERIES;
  IDENTIFY VAR=PWINDEX CROSSCOR(PWRAIN) NLAG=15
  OUTCOV=STATS;
RUN;
```

Here a prewhitened ring-index series (PWINDEX) is cross-correlated with a prewhitened rainfall series (PWRAIN). (Details of the computation of the prewhitened series are deferred until the discussion of FORECAST.) The NLAG and OUTCOV options are also illustrated. The NLAG option defines the maximum lag for which the ACF, IACF, PACF, and CCF are to be calculated (a warning message is printed if NLAG exceeds 25% of the series length); and the OUTCOV option creates a data set (STATS) containing the summary statistics generated by the IDENTIFY statement.

Output from the last example includes the sample ACF, IACF, and PACF of the prewhitened index series for lags 0 through 15, and a check for white noise. These results have the same interpretation as Figure 18 and will not be discussed further. The rest of the output is displayed in Figure 19. It shows the cross-correlogram for lags -15 through 15, which was also plotted

in Figure 11. If a cross-correlation value lies outside the pair of dotted lines it can be considered statistically significant at an approximately 5% level of significance (i.e., the null hypothesis that the cross-correlation is zero can be rejected at the 5% level of significance). Notice that in this case, only one significant spike occurs at lag 0, suggesting that the ring index is (weakly) correlated with the rainfall during the spring of the same year, but is uncorrelated with the rainfall in previous years.




FIGURE 19 Cross-correlation of prewhitened ring-index and rainfall series.

The sample ACF, IACF, and PACF (Figure 18) suggest that an AR model (refer to Section 3.5.1) might be appropriate for the ring-index series. Since the order of the autoregression is uncertain, three candidate models are fitted as follows:

```
PROC ARIMA DATA=INDXRAIN;
  TITLE1 `Time-Series Analysis: PROC ARIMA` ;
  TITLE2 `Fit models to ring-index series.` ;
  IDENTIFY VAR=INDEX NOPRINT NLAG=10;
  ESTIMATE P=1      METHOD=ML;
  ESTIMATE P=7      METHOD=ML;
  ESTIMATE P=(1 6 7) METHOD=ML PLOT;
RUN;
```

Each model is defined by a separate `ESTIMATE` statement and is fitted to the series specified in the preceding `IDENTIFY` statement. Since all three models are fitted to the same series (`INDEX`), only one `IDENTIFY` statement is needed. The `NOPRINT` option of `IDENTIFY` suppresses the printing of the sample ACF, IACF, and PACF (Figure 18), which were useful for model identification but are not required here. The `P` option of the `ESTIMATE` statements defines the autoregressive terms that are to be included in the fitted models. The option `P=k` includes all terms up to lag k . Thus the first two `ESTIMATE` statements fit the models

$$y_t = \nu + \phi_1 y_{t-1} + \varepsilon_t$$

and

$$y_t = \nu + \phi_1 y_{t-1} + \phi_2 y_{t-2} + \dots + \phi_7 y_{t-7} + \varepsilon_t$$

respectively. To limit the model to a specific set of autoregressive terms, the corresponding lags must be enclosed in parentheses as illustrated by the last `ESTIMATE` statement, which fits the model

$$y_t = \nu + \phi_1 y_{t-1} + \phi_6 y_{t-6} + \phi_7 y_{t-7} + \varepsilon_t$$

Moving-average models are defined in a similar manner with the `Q` option: for instance,

```
PROC ARIMA;
  IDENTIFY VAR=INDEX;
  ESTIMATE Q=(1 3);
RUN;
```

fits the model $y_t = \nu + \varepsilon_t - \phi_1 \varepsilon_{t-1} - \phi_3 \varepsilon_{t-3}$. Mixed models require that `P` and `Q` be specified: thus,

```
PROC ARIMA;
  IDENTIFY VAR=INDEX;
  ESTIMATE P=2 Q=(1 3);
RUN;
```


fits the ARMA model $y_t = \nu + \phi_1 y_{t-1} + \phi_2 y_{t-2} + \varepsilon_t - \phi_1 \varepsilon_{t-1} - \phi_3 \varepsilon_{t-3}$. To fit an ARIMA model, the appropriate differenced series is specified in the IDENTIFY statement (or created in a previous data step). The corresponding ARMA model is selected with the P and Q options of ESTIMATE. The following example demonstrates how PROC ARIMA can be used to fit an ARIMA(1, 1, 1) model:

```
PROC ARIMA;
  IDENTIFY VAR=INDEX(1);
  ESTIMATE P=1 Q=1;
RUN;
```

The results of fitting the three AR models are displayed in **A**, **B**, and **C** of Figure 20. For each model, the parameter estimates, approximate standard errors, *T*-statistics (with degrees of freedom equal to the number of observations minus the number of free parameters in the model), and associated lags are summarized in a table labelled “Maximum Likelihood Estimation.” The title reflects the method of estimation, which is selected with the METHOD option of ESTIMATE. In this case, maximum likelihood estimation (METHOD=ML) was used. Other methods are available, but maximum likelihood estimation is usually preferred because the resultant estimates have desirable statistical properties (see Chapter 6 of Diggle 1991).

Each set of parameter estimates includes an estimate of the mean of the response variable (*MU*), estimates of any autoregressive parameters $\{\phi_i\}$ (which have labels with the prefix “AR”), and estimates of any moving-average parameters $\{\theta_i\}$ (prefix “MA”). The estimate of the constant ν is labelled “Constant Estimate” and the estimated variance (standard deviation) of ε_t is labelled “Variance Estimate” (“Std Error Estimate”). All AR and MA estimates are numbered consecutively, with the associated lags listed in the last column. For example, reading the relevant parameter estimates (**C.1**) from Figure 20 yields the third and final fitted model:

$$y_t = 0.26800 + 0.70855 y_{t-1} + 0.16708 y_{t-6} - 0.14479 y_{t-7} + \varepsilon_t$$

Similarly, the estimates for the first and second models can be read from **A.1** and **B.1** of the output.

The constant and variance estimates are followed by two numbers labelled “AIC” and “SBC,” which stand for “Akaike’s information criterion” and “Schwartz’s Bayesian criterion.” Models are sometimes compared on the basis of these criteria, which take into account both the quality of the fit and the number of estimated parameters. Models with smaller values are judged to be “better” (i.e., provide a better fit without substantially increasing number parameters) than models with larger values. Comparison of AIC and SBC for the three AR models (refer to **A.2**, **B.2**, and **C.2** in Figure 20) suggests that the last model is best. The other two models are rejected because the first does not provide a good fit (see below) and the second has too many terms (notice that the AR parameters for lags 2–5 are not statistically significant according to their *T*-statistics).

Time-Series Analysis: PROC ARIMA
Fit models to ring-index series.

ARIMA Procedure

Maximum Likelihood Estimation (A)

Parameter	Estimate	Approx. Std Error	T Ratio	Lag	
MU	0.99608	0.02187	45.55	0	
AR1,1	0.71594	0.03252	22.02	1	(A.1)

Constant Estimate = 0.28294583

Variance Estimate = 0.01807707 (A.2)

Std Error Estimate = 0.134451

AIC = -541.41897

SBC = -533.14351

Number of Residuals= 463

Correlations of the Estimates

Parameter	MU	AR1,1
MU	1.000	0.009
AR1,1	0.009	1.000

Autocorrelation Check of Residuals (A.3)

To Lag	Chi Square	DF	Prob	Autocorrelations						
6	17.48	5	0.004	-0.055	0.007	0.077	-0.029	-0.026	0.163	
12	24.26	11	0.012	-0.019	-0.047	0.024	-0.046	0.001	0.095	
18	28.50	17	0.039	-0.016	-0.013	0.022	-0.031	-0.064	-0.053	
24	39.04	23	0.020	-0.066	-0.036	0.010	-0.094	0.013	-0.083	
30	53.72	29	0.003	-0.088	0.034	-0.092	0.041	-0.055	-0.088	
36	62.06	35	0.003	-0.051	-0.063	-0.041	0.015	-0.036	-0.083	
42	68.37	41	0.005	0.036	-0.007	-0.090	0.009	0.054	-0.010	

Model for variable INDEX

Estimated Mean = 0.99608307

Autoregressive Factors

Factor 1: 1 - 0.71594 B**(1)

FIGURE 20 Time-series analysis of ring-index series: model estimation.

Time-Series Analysis: PROC ARIMA
Fit models to ring-index series.

ARIMA Procedure

Maximum Likelihood Estimation

(B)

Parameter	Estimate	Approx.		Lag
		Std Error	T Ratio	
MU	0.99591	0.02334	42.68	0
AR1,1	0.67213	0.04653	14.45	1
AR1,2	0.04622	0.05572	0.83	2
AR1,3	0.05693	0.05564	1.02	3
AR1,4	-0.06529	0.05577	-1.17	4
AR1,5	-0.00958	0.05587	-0.17	5
AR1,6	0.17429	0.05574	3.13	6
AR1,7	-0.13755	0.04665	-2.95	7

(B.1)

Constant Estimate = 0.26176758

Variance Estimate = 0.01765436

(B.2)

Std Error Estimate = 0.1328697

AIC = -546.23093

SBC = -513.12911

Number of Residuals = 463

Correlations of the Estimates

Parameter	MU	AR1,1	AR1,2	AR1,3	AR1,4	AR1,5
MU	1.000	0.004	0.002	0.002	-0.001	0.000
AR1,1	0.004	1.000	-0.554	-0.037	-0.060	0.060
AR1,2	0.002	-0.554	1.000	-0.439	-0.000	-0.085
AR1,3	0.002	-0.037	-0.439	1.000	-0.433	0.000
AR1,4	-0.001	-0.060	-0.000	-0.433	1.000	-0.438
AR1,5	0.000	0.060	-0.085	0.000	-0.438	1.000
AR1,6	-0.000	0.017	0.042	-0.091	0.006	-0.439
AR1,7	0.001	-0.084	0.014	0.065	-0.062	-0.041

Parameter	AR1,6	AR1,7
MU	-0.000	0.001
AR1,1	0.017	-0.084
AR1,2	0.042	0.014
AR1,3	-0.091	0.065
AR1,4	0.006	-0.062
AR1,5	-0.439	-0.041
AR1,6	1.000	-0.552
AR1,7	-0.552	1.000

FIGURE 20 (continued)

Autocorrelation Check of Residuals

(B.3)

To Lag	Chi Square	DF	Prob	Autocorrelations					
6	0.00	0	0.000	-0.007	0.007	-0.003	-0.001	0.002	-0.013
12	4.64	5	0.461	0.025	-0.035	-0.004	-0.029	0.024	0.079
18	9.59	11	0.567	0.006	0.002	0.011	-0.004	-0.073	-0.069
24	18.73	17	0.344	-0.049	-0.042	0.022	-0.094	0.019	-0.069
30	31.47	23	0.112	-0.080	0.041	-0.083	0.053	-0.060	-0.067
36	39.43	29	0.094	-0.060	-0.077	-0.005	0.004	-0.043	-0.067
42	44.88	35	0.122	0.039	-0.008	-0.075	-0.016	0.056	0.007

Model for variable INDEX

Estimated Mean = 0.99590623

Autoregressive Factors

Factor 1: $1 - 0.67213 B^{**}(1) - 0.046223 B^{**}(2) - 0.056927 B^{**}(3)$
 $+ 0.065289 B^{**}(4) + 0.009579 B^{**}(5) - 0.17429 B^{**}(6)$
 $+ 0.13755 B^{**}(7)$

FIGURE 20 (continued)

Time-Series Analysis: PROC ARIMA
Fit models to ring-index series.

ARIMA Procedure

Maximum Likelihood Estimation (C)

Parameter	Estimate	Approx.		Lag	
		Std Error	T Ratio		
MU	0.99568	0.02278	43.71	0	
AR1,1	0.70855	0.03339	21.22	1	(C.1)
AR1,2	<u>0.16708</u>	0.04655	3.59	6	
AR1,3	<u>-0.14479</u>	0.04637	-3.12	7	

Constant Estimate = 0.26799846

Variance Estimate = 0.01763278 (C.2)

Std Error Estimate = 0.13278848

AIC = -550.7613

SBC = -534.21039

Number of Residuals= 463

Correlations of the Estimates

Parameter	MU	AR1,1	AR1,2	AR1,3
MU	1.000	0.007	0.000	0.001
AR1,1	0.007	1.000	-0.125	-0.092
AR1,2	0.000	-0.125	1.000	-0.698
AR1,3	0.001	-0.092	-0.698	1.000

Autocorrelation Check of Residuals (C.3)

To Lag	Chi Square	DF	Prob	Autocorrelations					
6	3.54	3	0.316	-0.043	0.024	0.066	-0.019	-0.014	-0.014
12	8.09	9	0.525	0.017	-0.033	0.016	-0.033	0.017	0.081
18	12.70	15	0.625	0.002	0.002	0.014	-0.009	-0.074	-0.062
24	23.26	21	0.330	-0.053	-0.044	0.019	-0.104	0.026	-0.070
30	37.40	27	0.088	-0.080	0.047	-0.096	0.050	-0.060	-0.069
36	45.46	33	0.073	-0.050	-0.080	-0.015	0.004	-0.041	-0.072
42	50.90	39	0.096	0.040	-0.007	-0.076	-0.011	0.056	-0.006

FIGURE 20 (continued)

Autocorrelation Plot of Residuals

(C.4)

Lag	Covariance	Correlation	-1	9	8	7	6	5	4	3	2	1	0	1	2	3	4	5	6	7	8	9	1	
0	0.017633	1.00000												*****										
1	-0.0007632	-0.04328										.*												
2	0.00042402	0.02405										.												
3	0.0011637	0.06600										.	*											
4	-0.0003339	-0.01894										.												
5	-0.0002492	-0.01414										.												
6	-0.0002487	-0.01410										.												
7	0.00029504	0.01673										.												
8	-0.0005798	-0.03288										.	*											
9	0.00028309	0.01605										.												
10	-0.0005744	-0.03257										.	*											

``.`` marks two standard errors

Inverse Autocorrelations

Lag	Correlation	-1	9	8	7	6	5	4	3	2	1	0	1	2	3	4	5	6	7	8	9	1	
1	0.03989										.	*											
2	-0.02780										.	*											
3	-0.07089										.	*											
4	0.01846										.												
5	0.01828										.												
6	0.01898										.												
7	-0.01985										.												
8	0.02563										.		*										
9	-0.01135										.												
10	0.03281										.		*										

Partial Autocorrelations

Lag	Correlation	-1	9	8	7	6	5	4	3	2	1	0	1	2	3	4	5	6	7	8	9	1	
1	-0.04328										.	*											
2	0.02222										.												
3	0.06814										.		*										
4	-0.01380										.												
5	-0.01898										.												
6	-0.01931										.												
7	0.01852										.												
8	-0.02874										.	*											
9	0.01426										.												
10	-0.03316										.	*											

Model for variable INDEX

Estimated Mean = 0.995679

Autoregressive Factors

Factor 1: $1 - 0.70855 B^{**}(1) - 0.16708 B^{**}(6) = 0.14479 B^{**}(7)$

FIGURE 20 (concluded).

The results of each `ESTIMATE` statement also include the estimated correlations between the parameter estimates (which are usually of less interest than the parameter estimates themselves) and an “Autocorrelation Check of Residuals” (**A.3**, **B.3**, and **C.3**). The latter has the same interpretation as the “Autocorrelation Check for White Noise” of the `IDENTIFY` statement (part **D** of Figure 18), except in this case, the tests are applied to the residuals from the fitted model. If the model is adequate, the residuals are expected to resemble white noise (i.e., all autocorrelations should be near zero). Notice that the residuals for the first model show evidence of autocorrelation, which suggests that more AR or MA terms are required. The other two models show little or no evidence of autocorrelation, although the second model has too many parameters (i.e., no degrees of freedom) to test the first six lags.

Plots of the ACF, IACF, and PACF of the residuals are also useful for checking the adequacy of the fitted model. They are requested with the `PLOT` option of `ESTIMATE`. Figure 20 (**C.4**) shows the residual ACF, IACF, and PACF for the last of the three fitted models. Only lags 0 through 10 are plotted because a maximum lag of 10 (`NLAG=10`) was specified in the `IDENTIFY` statement (this was done mainly for convenience, to limit the amount of output). Notice that none of the plots shows any evidence of autocorrelation, which is consistent with a good fit.

After a model has been successfully fitted to a time series, predicted values, residuals, and forecasts can be obtained with the `FORECAST` statement as demonstrated below:

```
PROC ARIMA DATA=INDXRAIN;
  TITLE1 'Time-Series Analysis: PROC ARIMA';
  TITLE2 'Calculate predicted values and residuals for ring-index series.';
  IDENTIFY VAR=INDEX NOPRINT;
  ESTIMATE P=(1 6 7) METHOD=ML NOPRINT;
  FORECAST LEAD=0 OUT=PWINDEX ID=YEAR;
RUN;
```

Here the `IDENTIFY` and `ESTIMATE` statements are used to fit the model that was previously determined to be the best of the three AR models. Since there is no need to repeat the results from the model identification or estimation (Figures 18 and 20), they are suppressed with the `NOPRINT` options of `IDENTIFY` and `ESTIMATE`. The `FORECAST` statement uses the results of the most recent `ESTIMATE` statement to compute fitted values and residuals (i.e., observed – fitted values). These are saved in the data set (`PWINDEX`) specified by the `OUT` option. Forecasts (i.e., predicted values for future observations) are requested with the `LEAD` option, which specifies the lead time or number of forecasts. In this application, only the fitted values and residuals are requested since `LEAD=0`. If forecasts were required for the next k years, then `LEAD=0` would be replaced with `LEAD=k`. The last option, `ID=YEAR`, adds the variable `YEAR` to the output data set.

The output data set (`PWINDEX`) contains the fitted ring-index values and, if applicable, forecasted values (both types of predicted values have

the variable name FORECAST), the corresponding estimated standard errors (STD), upper (U95) and lower (L95) 95% confidence limits, the residuals (RESIDUAL), and the identification variable YEAR. The residuals are the prewhitened values that were used to calculate the cross-correlogram shown in Figure 19. The rainfall series was prewhitened in the same manner, by fitting a suitable time-series model and computing the residuals.

The ARIMA procedure is extremely versatile. It can be used for exploratory analyses (i.e., computation of the ACF, CCF, etc.), time-series modelling (including transfer function models and intervention models, as well as ARIMA models), and forecasting. For more information about the IDENTIFY, ESTIMATE, and FORECAST statements of PROC ARIMA, and their various other options, the reader should consult the SAS/ETS User's Guide (SAS Institute 1991a).

4.2.3 PROC AUTOREG The model on which PROC REG is based—the ordinary least-squares regression model—assumes that the data are independent; that is,

$$y_t = \beta_0 + \beta_1 x_{1t} + \beta_2 x_{2t} + \dots + \beta_m x_{mt} + \varepsilon_t$$

where the errors $\{\varepsilon_t\}$ are independent and identically distributed normal random variables. When observations are made over time, the independence assumption may not be realistic. To allow for this possibility, PROC AUTOREG assumes that the errors are generated by an autoregression; that is,

$$\varepsilon_t = \delta_t - \alpha_1 \varepsilon_{t-1} - \alpha_2 \varepsilon_{t-2} - \dots - \alpha_p \varepsilon_{t-p}$$

where $\{\delta_t\}$ are independent and identically distributed normal random variables. Notice that when there are no independent variables (i.e., $y_t = \beta_0 + \varepsilon_t$), the model reduces to a simple autoregression:

$$y_t = \alpha_0 - \alpha_1 y_{t-1} - \alpha_2 y_{t-2} - \dots - \alpha_p y_{t-p} + \delta_t$$

with $\alpha_0 = \beta_0$

To illustrate PROC AUTOREG, consider the final AR model that was fitted to ring-index series with PROC ARIMA. The same analysis can be performed with PROC AUTOREG as follows:

```
PROC AUTOREG DATA=INDXRAIN;
  TITLE1 'Time-Series Analysis: PROC AUTOREG';
  TITLE2 'Fit model to ring-index series.';
  MODEL INDEX=/NLAG=7 METHOD=ML;
  LAGLIST 1 6 7;
  OUTPUT OUT=RESID R=PWINDEX;
RUN;
```


The `MODEL` statement identifies the dependent response variable (`INDEX`) and the regressors $x_{1t}, x_{2t}, \dots, x_{mt}$. In this example there are no regressors, but if there were they would be listed after the equals sign (e.g., `MODEL INDEX=X1 X2;`). The options `NLAG` and `METHOD` specify the order of the autoregression for the error process and the method of estimation (`ML` = maximum likelihood estimation). A `LAGLIST` statement is also required since the number of autoregressive parameters is less than the value specified by `NLAG`. It identifies the lagged terms that are to be included in the autoregression (if `LAGLIST` is omitted, all lags up to and including `NLAG` are selected). The `OUTPUT` statement identifies the output data set (`RESID`) and names the residuals (`PWINDEX`).

The output from `PROC AUTOREG` is presented in Figure 21. The first part (**A**) gives the ordinary least-squares estimates of the parameters (which are the same as the values that would be obtained with `PROC REG`) and the results of a diagnostic analysis of the residuals, including the Durbin-Watson statistic and sample ACF up to the lag specified by `NLAG`. Notice that the Durbin-Watson statistic (0.5685) is considerably less than 2, which is the expected value under the null hypothesis of independent errors. This confirms the need for a model with autocorrelated errors.

The second part (**B**) of the output lists preliminary estimates of the mean square error and autoregressive parameters for the correlated-errors model. These so-called “Yule-Walker” estimates are the starting values for the iterative scheme that is used to obtain the maximum likelihood estimates, given in the last part of the output (**C**). The Yule-Walker estimates are easier to calculate and are usually good first approximations to the maximum likelihood estimates. Notice that the standard error (0.02280) for the maximum likelihood estimate (**C.1**) is considerably larger than the corresponding standard error (0.008923) for the ordinary least-squares estimate of the intercept (**A.1**), even though the estimated intercepts are approximately equal. Here, and in general, failure to account for autocorrelation among the observations causes the standard errors of the regression coefficients to be underestimated. Notice also that the maximum likelihood estimates of AR parameters and associated summary statistics (AIC, SBC, etc.) are the same as the values obtained with `PROC ARIMA`, apart from the sign reversal of the AR estimates, which reflects the different parameters used by the two procedures. The Durbin-Watson statistic provides an additional check of the residuals, which is not available in `PROC ARIMA`. After fitting the AR model, the residuals no longer show evidence of autocorrelation (i.e., the Durbin-Watson statistic is close to 2).

4.2.4 PROC FORECAST The emphasis of `PROC ARIMA` and `PROC AUTOREG` is model development and hypothesis testing. In contrast, the primary objective of `PROC FORECAST` is forecasting, without necessarily understanding details of the process by which the data are generated. There are three methods for generating forecasts with `PROC FORECAST`: a stepwise autoregressive method (`STEPAR`), exponential smoothing (`EXPO`), and the Holt-Winters method (`WINTERS`). The choice of method depends on the characteristics of the series—that is, the type of trend, if

Time-Series Analysis: PROC AUTOREG
Fit model to ring-index series.

Autoreg Procedure

Dependent Variable = INDEX Tree ring index

Ordinary Least Squares Estimates (A)

SSE	17.0303	DFE	462
MSE	0.036862	Root MSE	0.191995
SBC	-209.091	AIC	-213.228
Reg Rsq	0.0000	Total Rsq	0.0000
Durbin-Watson	<u>0.5685</u>		

Variable	DF	B Value	Std Error	t Ratio	Approx Prob	
Intercept	1	0.996460043	<u>0.008923</u>	111.676	0.0001	(A.1)

Estimates of Autocorrelations

Lag	Covariance	Correlation	-1	9	8	7	6	5	4	3	2	1	0	1	2	3	4	5	6	7	8	9	1	
0	0.036782	1.000000																						
1	0.026179	0.711727																						
2	0.02003	0.544552																						
3	0.01596	0.433909																						
4	0.011755	0.319577																						
5	0.009626	0.261701																						
6	0.009216	0.250548																						
7	0.005767	0.156796																						

Preliminary MSE = 0.017658

Estimates of the Autoregressive Parameters (B)

Lag	Coefficient	Std Error	t Ratio
1	-0.70299696	0.03365984	-20.885334
6	-0.16280822	0.04638895	-3.509633
7	0.13521321	0.04624735	2.923696

Expected Autocorrelations

Lag	Autocorr
0	1.0000
1	0.7115
2	0.5143
3	0.3832
4	0.3013
5	0.2581
6	0.2481
7	0.1550

FIGURE 21 Time-series analysis of ring-index series: PROC AUTOREG.

Time-Series Analysis: PROC AUTOREG
Fit model to ring-index series.

Maximum Likelihood Estimates

(C)

SSE	8.093447	DFE	459
MSE	0.017633	Root MSE	0.132788
SBC	-534.21	AIC	-550.761
Reg Rsq	0.0000	Total Rsq	0.5248
Durbin-Watson	<u>2.0748</u>		

Variable	DF	B Value	Std Error	t Ratio	Approx Prob
Intercept	1	0.995679409	<u>0.02280</u>	43.669	0.0001
A(1)	1	<u>-0.708546978</u>	0.03342	-21.200	0.0001
A(6)	1	<u>-0.167079541</u>	0.04660	-3.586	0.0004
A(7)	1	<u>0.144794292</u>	0.04642	3.119	0.0019

(C.1)

Expected Autocorrelations

Lag	Autocorr
0	1.0000
1	0.7162
2	0.5209
3	0.3898
4	0.3068
5	0.2616
6	0.2487
7	0.1511

Autoregressive parameters assumed given.

Maximum Likelihood Estimates

SSE	8.093447	DFE	459
MSE	0.017633	Root MSE	0.132788
SBC	-534.21	AIC	-550.761
Reg Rsq	0.0000	Total Rsq	0.5248
Durbin-Watson	<u>2.0748</u>		

Variable	DF	B Value	Std Error	t Ratio	Approx Prob
Intercept	1	0.995679409	<u>0.02280</u>	43.669	0.0001

FIGURE 21 (continued).

there is one, and the presence or absence of a seasonal component. The stepwise AR method is applicable when there is no seasonal component and the trend is adequately approximated by a deterministic linear or quadratic function (e.g., Figure 4a). If the series exhibits a stochastic trend (e.g., Figure 9a), exponential smoothing or the Holt-Winters method can be used. The former is applicable when no seasonal component is present. When a seasonal component is present, the latter (which is a seasonal version of exponential smoothing) should be used.

Forecasts for the ring-index series can be generated with PROC FORECAST, as follows:

```
PROC FORECAST DATA=INDXRAIN METHOD=STEPAR AR=7 TREND=1
                LEAD=5 OUT=FORECAST OUTLIMIT OUTEST=PARMS;

    VAR INDEX;
    ID YEAR;
PROC PRINT DATA=PARMS;
    TITLE1 'Forecast ring-index series with PROC FORECAST.';
    TITLE2 'Parameter estimates.';
PROC PRINT DATA=FORECAST;
    TITLE 'Forecast ring-index series with PROC FORECAST';
    TITLE2 'Forecasts for 1993-1997.';
RUN;
```

The PROC FORECAST line sets the options for generating forecasts and controls the output (refer to SAS/ETS User's Guide [SAS Institute 1991a], for a complete list of options). In this case, the stepwise AR method is chosen (METHOD=STEPAR). The AR (or NLAGS) option selects the maximum order for the fitted autoregression. Based on the previous analyses, AR=7 is chosen. If no estimate of the order is available, a large value should be assigned (the default is AR=13) and the order of the fitted autoregression determined automatically. The order selection process is controlled by the options SLENTY (significance level for entering a term) and SLSTAY (significance level for removing a term), as well as AR. In this case, SLSTAY and SLENTY are omitted so the default values of 0.05 and 0.2 are used. The TREND option describes the type of the trend (1 = constant, 2 = linear, 3 = quadratic) and LEAD specifies the number of time steps into the future to be forecasted. Here LEAD=5 requests forecasts for the five years following the last observation (i.e., forecasts for the years 1993–1997). The three “OUT” options control the output: OUT names the output data set for the forecasts, OUTLIMIT requests confidence limits for the forecasts, and OUTEST names the data set for the parameter estimates.

The output from the preceding program is displayed in Figure 22. The first part (A) lists the parameter estimates. Notice that an autoregression with lags 1, 6, and 7 is fitted (the missing values for lags 2–5 imply that those terms were dropped) and that estimated parameters are the same as the preliminary Yule-Walker estimates calculated by PROC AUTOREG (part B of Figure 21). The forecasts for the years 1993 through 1997 (B), and the corresponding 95% confidence limits, are printed after the parameter estimates.

Forecast ring-index series with PROC FORECAST. (A)
Parameter estimates.

OBS	__TYPE__	YEAR	INDEX
1	N	1992	463
2	SIGMA	1992	0.1334615
3	CONSTANT	1992	0.99646
4	AR1	1992	0.702997
5	AR2	1992	.
6	AR3	1992	.
7	AR4	1992	.
8	AR5	1992	.
9	AR6	1992	0.1628082
10	AR7	1992	-0.135213

Forecast ring-index series with PROC
FORECAST (B)
Forecasts for 1993-1997.

OBS	YEAR	__TYPE__	__LEAD__	INDEX
1	1993	FORECAST	1	1.15013
2	1993	L95	1	0.88769
3	1993	U95	1	1.41258
4	1994	FORECAST	2	1.09827
5	1994	L95	2	0.77782
6	1994	U95	2	1.41873
7	1995	FORECAST	3	1.04307
8	1995	L95	3	0.69753
9	1995	U95	3	1.38862
10	1996	FORECAST	4	1.05039
11	1996	L95	4	0.69309
12	1996	U95	4	1.40769
13	1997	FORECAST	5	1.04565
14	1997	L95	5	0.68269
15	1997	U95	5	1.40862

FIGURE 22 *Ring-index forecasts generated with PROC FORECAST (METHOD=STEPAR).*

5 SAS EXAMPLES

5.1 Repeated-measures Analysis of Seedling Height Growth

The first example is an expanded analysis of the seedling data (Section 1.1.1). In the previous analysis (Section 4.1), only seedlings grown from lodgepole pine plugs were considered. The following SAS program carries out both univariate and multivariate repeated-measures analyses of the

complete data set, and the corresponding analyses for each of the two species: Douglas-fir (FD) and lodgepole pine (PL). The analysis of the individual species includes an ANOVA of successive height increments (i.e., the change in height from 1983 to 1984, 1984 to 1985, etc.). This is easily generated with the PROFILE transformation and SUMMARY option of the REPEATED statement. A MANOVA statement is added to test the statistical significance of the overall mean of each increment. Part of the output is shown in Figure 23.

```

DATA SEEDLING;                                /* Read data in multivariate form.      */
  INFILE 'SURVIV.DAT';
  INPUT SPP STK TRT BLK ROW TRENO PHT HT84-HT89;

PROC SORT DATA=SEEDLING;                     /* Sort data set and calculate row      */
  BY SPP STK TRT BLK ROW;                     /* averages.                            */
PROC SUMMARY DATA=SEEDLING;
  BY SPP STK TRT BLK ROW;
  VAR PHT HT84-HT89;
  OUTPUT OUT=ROWMEANS MEAN=;

                                                /* Repeated measures analysis of row   */
                                                /* means - both species.                */

PROC GLM DATA=ROWMEANS;
  TITLE 'Repeated-Measures Analysis of Seedling Growth (Section 1.1.1)';
  CLASS BLK TRT SPP STK;
  MODEL PHT HT84-HT89=BLK TRT|SPP|STK/SS3;
  REPEATED YEAR 7 (1983 1984 1985 1986 1987 1988 1989);

                                                /* Repeated measures analysis of row   */
                                                /* means - by species.                  */

PROC GLM DATA=ROWMEANS;
  TITLE1 'Repeated-Measures Analysis of Seedling Growth - by species';
  BY SPP;
  CLASS BLK TRT STK;
  MODEL PHT HT84-HT89=BLK TRT|STK/SS3;
  REPEATED YEAR 7 (1983 1984 1985 1986 1987 1988 1989) PROFILE/SUMMARY;
  MANOVA H=INTERCEPT E=BLK M=( 1 -1  0  0  0  0  0, 0  1 -1  0  0  0  0,
                                   0  0  1 -1  0  0  0, 0  0  0  1 -1  0  0,
                                   0  0  0  0  1 -1  0, 0  0  0  0  0  1 -1)
  MNames=INC1 INC2 INC3 INC4 INC5 INC6/SUMMARY;

RUN;

```

The results of the analysis of the complete data set (not shown in Figure 23) suggest that site-preparation treatment, species, and stock type have a nonadditive effect on the growth of seedlings (i.e., the YEAR*TRT*SPP*STK interaction is significant, which implies that the effect of one factor depends on the other two). Analyzing each species separately provides insight into the various effects.

The results of the multivariate within-rows analyses (Parts A and C of Figure 23) suggest that treatment and stock type have a significant

Repeated-Measures Analysis of Seedling Growth - by species

----- SPP=FD -----

Manova Test Criteria and F Approximations for (A)
the Hypothesis of no YEAR*BRT Effect
H = Type III SS&CP Matrix for YEAR*BRT E = Error SS&CP Matrix

Statistic	S=3	M=1	N=4			
	Value	F	Num DF	Den DF	Pr > F	
Wilks' Lambda	0.35327581	0.7107	18	28.76955	0.7734	
Pillai's Trace	0.78357175	0.7071	18	36	0.7816	
Hotelling-Lawley Trace	1.45319704	0.6997	18	26	0.7814	
Roy's Greatest Root	1.13784264	2.2757	6	12	0.1062	

NOTE: F Statistic for Roy's Greatest Root is an upper bound.

Manova Test Criteria and F Approximations for
the Hypothesis of no YEAR*TRT Effect
H = Type III SS&CP Matrix for YEAR*TRT E = Error SS&CP Matrix

Statistic	S=2	M=1.5	N=4			
	Value	F	Num DF	Den DF	Pr > F	
Wilks' Lambda	0.08931022	3.9103	12	20	0.0035	
Pillai's Trace	1.39068228	4.1843	12	22	0.0018	
Hotelling-Lawley Trace	4.82248581	3.6169	12	18	0.0070	
Roy's Greatest Root	3.07430517	5.6362	6	11	0.0068	

NOTE: F Statistic for Roy's Greatest Root is an upper bound.

NOTE: F Statistic for Wilks' Lambda is exact.

Manova Test Criteria and Exact F Statistics for
the Hypothesis of no YEAR*STK Effect
H = Type III SS&CP Matrix for YEAR*STK E = Error SS&CP Matrix

Statistic	S=1	M=2	N=4			
	Value	F	Num DF	Den DF	Pr > F	
Wilks' Lambda	0.09360369	16.1389	6	10	0.0001	
Pillai's Trace	0.90639631	16.1389	6	10	0.0001	
Hotelling-Lawley Trace	9.68333917	16.1389	6	10	0.0001	
Roy's Greatest Root	9.68333917	16.1389	6	10	0.0001	

FIGURE 23 Repeated-measures analysis of the growth of Douglas-fir (FD) and lodgepole pine (PL) seedlings.

Manova Test Criteria and F Approximations for
the Hypothesis of no YEAR*TRT*STK Effect
H = Type III SS&CP Matrix for YEAR*TRT*STK E = Error SS&CP Matrix

	S=2	M=1.5	N=4		
Statistic	Value	F	Num DF	Den DF	Pr > F
Wilks' Lambda	0.26693746	1.5592	12	20	0.1835
Pillai's Trace	0.96038574	1.6936	12	22	0.1370
Hotelling-Lawley Trace	1.89459867	1.4209	12	18	0.2429
Roy's Greatest Root	1.16125919	2.1290	6	11	0.1315

NOTE: F Statistic for Roy's Greatest Root is an upper bound.

NOTE: F Statistic for Wilks' Lambda is exact.

Tests of Hypotheses for Between Subjects Effects

Source	DF	Type III SS	Mean Square	F Value	Pr > F
BLK	3	100.3616	33.4539	0.33	0.8005
TRT	2	18.6268	9.3134	0.09	0.9116
STK	1	5734.5531	5734.5531	57.36	0.0001
TRT*STK	2	129.7455	64.8728	0.65	0.5367
Error	15	1499.5926	99.9728		

YEAR.N represents the nth successive difference in YEAR (B)

Contrast Variable: YEAR.1

Source	DF	Type III SS	Mean Square	F Value	Pr > F
INTERCEPT	1	694.1780666	694.1780666	5245.83	0.0001
BLK	3	0.3969884	0.1323295	0.19	0.9019
TRT	2	7.5455373	3.7727687	5.40	0.0171
STK	1	31.1908783	31.1908783	44.67	0.0001
TRT*STK	2	3.8294677	1.9147339	2.74	0.0966
Error	15	10.4747898	0.6983193		

Contrast Variable: YEAR.2

Source	DF	Type III SS	Mean Square	F Value	Pr > F
INTERCEPT	1	196.9523432	196.9523432	439.44	0.0002
BLK	3	1.3445730	0.4481910	0.29	0.8337
TRT	2	2.7786218	1.3893109	0.89	0.4308
STK	1	51.3021147	51.3021147	32.91	0.0001
TRT*STK	2	6.2404211	3.1202105	2.00	0.1696
Error	15	23.3838857	1.5589257		

FIGURE 23 (continued)

Contrast Variable: YEAR.3

Source	DF	Type III SS	Mean Square	F Value	Pr > F
INTERCEPT	1	340.8068656	340.8068656	427.76	0.0002
BLK	3	2.3901996	0.7967332	0.17	0.9177
TRT	2	15.0366612	7.5183306	1.57	0.2413
STK	1	61.9420632	61.9420632	12.90	0.0027
TRT*STK	2	11.0474349	5.5237175	1.15	0.3430
Error	15	72.0473290	4.8031553		

Contrast Variable: YEAR.4

Source	DF	Type III SS	Mean Square	F Value	Pr > F
INTERCEPT	1	720.9847437	720.9847437	124.62	0.0015
BLK	3	17.3561446	5.7853815	0.88	0.4722
TRT	2	6.2875278	3.1437639	0.48	0.6281
STK	1	82.3489709	82.3489709	12.57	0.0029
TRT*STK	2	21.4928604	10.7464302	1.64	0.2269
Error	15	98.2998740	6.5533249		

Contrast Variable: YEAR.5

Source	DF	Type III SS	Mean Square	F Value	Pr > F
INTERCEPT	1	2643.323212	2643.323212	471.55	0.0002
BLK	3	16.816883	5.605628	0.84	0.4949
TRT	2	20.564462	10.282231	1.53	0.2478
STK	1	40.449551	40.449551	6.03	0.0267
TRT*STK	2	2.419067	1.209534	0.18	0.8368
Error	15	100.594149	6.706277		

Contrast Variable: YEAR.6

Source	DF	Type III SS	Mean Square	F Value	Pr > F
INTERCEPT	1	4158.138937	4158.138937	3529.16	0.0001
BLK	3	3.534670	1.178223	0.10	0.9568
TRT	2	104.916963	52.458481	4.60	0.0277
STK	1	73.653859	73.653859	6.46	0.0226
TRT*STK	2	5.287465	2.643733	0.23	0.7958
Error	15	171.016228	11.401082		

FIGURE 23 (continued)

Repeated-Measures Analysis of Seedling Growth - by species

----- SPP=PL -----

Manova Test Criteria and F Approximations for (C)
the Hypothesis of no YEAR*BLK Effect

H = Type III SS&CP Matrix for YEAR*BLK E = Error SS&CP Matrix

Statistic	S=3	M=1	N=4			
	Value	F	Num DF	Den DF	Pr > F	
Wilks' Lambda	0.19910807	1.2296	18	28.76955	0.3025	
Pillai's Trace	1.16471225	1.2692	18	36	0.2640	
Hotelling-Lawley Trace	2.40797089	1.1594	18	26	0.3576	
Roy's Greatest Root	1.62683040	3.2537	6	12	0.0389	

NOTE: F Statistic for Roy's Greatest Root is an upper bound.

Manova Test Criteria and F Approximations for
the Hypothesis of no YEAR*TRT Effect

H = Type III SS&CP Matrix for YEAR*TRT E = Error SS&CP Matrix

Statistic	S=2	M=1.5	N=4			
	Value	F	Num DF	Den DF	Pr > F	
Wilks' Lambda	0.03787338	6.8974	12	20	0.0001	
Pillai's Trace	1.55376510	6.3836	12	22	0.0001	
Hotelling-Lawley Trace	9.78228222	7.3367	12	18	0.0001	
Roy's Greatest Root	7.77242156	14.2494	6	11	0.0001	

NOTE: F Statistic for Roy's Greatest Root is an upper bound.

NOTE: F Statistic for Wilks' Lambda is exact.

Manova Test Criteria and Exact F Statistics for
the Hypothesis of no YEAR*STK Effect

H = Type III SS&CP Matrix for YEAR*STK E = Error SS&CP Matrix

Statistic	S=1	M=2	N=4			
	Value	F	Num DF	Den DF	Pr > F	
Wilks' Lambda	0.18244512	7.4685	6	10	0.0031	
Pillai's Trace	0.81755488	7.4685	6	10	0.0031	
Hotelling-Lawley Trace	4.48110016	7.4685	6	10	0.0031	
Roy's Greatest Root	4.48110016	7.4685	6	10	0.0031	

FIGURE 23 (continued)

Manova Test Criteria and F Approximations for
the Hypothesis of no YEAR*TRT*STK Effect
H = Type III SS&CP Matrix for YEAR*TRT*STK E = Error SS&CP Matrix

Statistic	S=2 M=1.5 N=4			Den DF	Pr > F
	Value	F	Num DF		
Wilks' Lambda	0.40010042	0.9682	12	20	0.5072
Pillai's Trace	0.68713772	0.9595	12	22	0.5119
Hotelling-Lawley Trace	1.28133196	0.9610	12	18	0.5154
Roy's Greatest Root	1.07931424	1.9787	6	11	0.1547

NOTE: F Statistic for Roy's Greatest Root is an upper bound.

NOTE: F Statistic for Wilks' Lambda is exact.

Tests of Hypotheses for Between Subjects Effects

Source	DF	Type III SS	Mean Square	F Value	Pr > F
BLK	3	310.8867	103.6289	0.80	0.5144
TRT	2	5176.6388	2588.3194	19.91	0.0001
STK	1	2069.3402	2069.3402	15.92	0.0012
TRT*STK	2	35.1474	17.5737	0.14	0.8746
Error	15	1949.8385	129.9892		

YEAR.N represents the nth successive difference in YEAR **(D)**

Contrast Variable: YEAR.1

Source	DF	Type III SS	Mean Square	F Value	Pr > F
INTERCEPT	1	1148.722693	1148.722693	278.69	0.0005
BLK	3	12.365716	4.121905	3.61	0.0384
TRT	2	12.682957	6.341478	5.55	0.0157
STK	1	32.713768	32.713768	28.65	0.0001
TRT*STK	2	2.137003	1.068502	0.94	0.4140
Error	15	17.128055	1.141870		

Contrast Variable: YEAR.2

Source	DF	Type III SS	Mean Square	F Value	Pr > F
INTERCEPT	1	3306.602159	3306.602159	1882.02	0.0001
BLK	3	5.270843	1.756948	0.73	0.5488
TRT	2	28.463329	14.231665	5.93	0.0127
STK	1	47.177524	47.177524	19.66	0.0005
TRT*STK	2	4.344947	2.172474	0.91	0.4254
Error	15	35.996519	2.399768		

FIGURE 23 (continued)

Contrast Variable: YEAR.3

Source	DF	Type III SS	Mean Square	F Value	Pr > F
INTERCEPT	1	6335.366338	6335.366338	9195.68	0.0001
BLK	3	2.066852	0.688951	0.11	0.9544
TRT	2	81.064714	40.532357	6.33	0.0102
STK	1	20.249735	20.249735	3.16	0.0957
TRT*STK	2	25.405195	12.702598	1.98	0.1722
Error	15	96.094158	6.406277		

Contrast Variable: YEAR.4

Source	DF	Type III SS	Mean Square	F Value	Pr > F
INTERCEPT	1	10941.23603	10941.23603	12593.09	0.0001
BLK	3	2.60649	0.86883	0.06	0.9817
TRT	2	716.78647	358.39323	23.27	0.0001
STK	1	129.45717	129.45717	8.40	0.0110
TRT*STK	2	1.71725	0.85863	0.06	0.9460
Error	15	231.04136	15.40276		

Contrast Variable: YEAR.5

Source	DF	Type III SS	Mean Square	F Value	Pr > F
INTERCEPT	1	23787.47161	23787.47161	1770.44	0.0001
BLK	3	40.30773	13.43591	1.11	0.3754
TRT	2	740.50850	370.25425	30.63	0.0001
STK	1	221.90586	221.90586	18.36	0.0007
TRT*STK	2	3.33963	1.66982	0.14	0.8721
Error	15	181.31448	12.08763		

Contrast Variable: YEAR.6

Source	DF	Type III SS	Mean Square	F Value	Pr > F
INTERCEPT	1	39922.99305	39922.99305	1577.06	0.0001
BLK	3	75.94464	25.31488	2.26	0.1233
TRT	2	794.08672	397.04336	35.46	0.0001
STK	1	279.89157	279.89157	25.00	0.0002
TRT*STK	2	9.53943	4.76972	0.43	0.6608
Error	15	167.96039	11.19736		

FIGURE 23 (concluded).

additive effect on the growth of both species (i.e., the YEAR*STK and YEAR*TRT interactions are significant for both species but the YEAR*TRT*STK interaction is not). The analysis of the annual increments (**B** and **D**), which are labelled YEAR.1 (1983 height–1984 height), YEAR.2 (1984 height–1985 height), etc., are even more informative. To reduce the amount of output, the ANOVA tables for the increments have been modified by replacing the default test for the MEAN (i.e., the test produced by the REPEATED statement, which is incorrect for a randomized block design) with the correct MANOVA test (labelled INTERCEPT). For convenience, the *p*-values for the corrected analyses have been extracted from the output and are summarized in Table 2. Notice that stock type appears to have a more significant and persistent effect on the Douglas-fir seedlings than the method of site preparation. However, both factors have an important and apparently prolonged effect on the growth of lodgepole pine seedlings.

TABLE 2 Analysis of annual height increments: summary of *p*-values (*p*-values less than 0.05 are underlined)

Source	1983–84	1984–85	1985–86	1986–87	1987–88	1988–89
Douglas-fir						
INTERCEPT	<u>0.0001</u>	<u>0.0002</u>	<u>0.0002</u>	<u>0.0015</u>	<u>0.0002</u>	<u>0.0001</u>
BLK	0.9019	0.8337	0.9177	0.4722	0.4949	0.9568
TRT	<u>0.0171</u>	0.4308	0.2413	0.6281	0.2478	<u>0.0277</u>
STK	<u>0.0001</u>	<u>0.0001</u>	<u>0.0027</u>	<u>0.0029</u>	<u>0.0267</u>	<u>0.0226</u>
TRT*STK	0.0966	0.1696	0.3430	0.2269	0.8368	0.7958
Lodgepole pine						
INTERCEPT	<u>0.0005</u>	<u>0.0001</u>	<u>0.0001</u>	<u>0.0001</u>	<u>0.0001</u>	<u>0.0001</u>
BLK	<u>0.0384</u>	0.5488	0.9544	0.9817	0.3754	0.1233
TRT	<u>0.0157</u>	<u>0.0127</u>	<u>0.0102</u>	<u>0.0001</u>	<u>0.0001</u>	<u>0.0001</u>
STK	<u>0.0001</u>	<u>0.0005</u>	0.0957	<u>0.0110</u>	<u>0.0007</u>	<u>0.0002</u>
TRT*STK	0.4140	0.4254	0.1722	0.9460	0.8721	0.6608

5.2 Cross-correlation
Analysis of Missing
Tree Rings

The second example is a cross-correlation analysis of the tree-ring series shown in Figure 2a and 2b. The objective is to find the number of years by which the lower series must be shifted relative to the upper series to achieve the best alignment, or maximum cross-correlation, of the growth patterns. Before calculating the cross-correlation function, the confounding effects of trend and autocorrelation must be eliminated.

An SAS program for carrying out the preliminary detrending and prewhitening of the individual series and the final cross-correlation analysis is listed below. First, a log transformation is applied to each series to convert the exponential trend into a linear trend. This has the added effect of stabilizing the variance. Next a trend line is fitted to each series with PROC REG (PROC AUTOREG could also be used). Subtraction of the trend produces a detrended ring index or, in this case, the log of the ring index. Prewhitening is accomplished by fitting a first-order autoregression to each ring-index series. (A preliminary model identification step suggested that both ring-index series are adequately described by a first-order auto-

regression.) The prewhitened series are calculated by including a FORECAST statement in the PROC ARIMA step (refer to Section 4.2.2). This saves the residuals from the fitted autoregressions, which are the required prewhitened values. Finally, the two prewhitened series are cross-correlated in a third application of PROC ARIMA.

```

                                                    /* Read data from input file and   */
                                                    /* apply log transformation.         */
DATA TREE94;
  TITLE 'Example 1.1.2: Cross-Correlation Analysis of Missing Tree Rings';
  INFILE 'TREE94.DAT' MISSOVER;
  INPUT DISKHT 21-25 YEAR 29-32 RINGNO 34-36 RNGWIDTH 39-43;
  LOGRW=LOG(RNGWIDTH);
                                                    /* Fit trend line to each series. */

PROC SORT DATA=TREE94;
  BY DISKHT;
PROC REG NOPRINT;
  BY DISKHT;
  MODEL LOGRW=RINGNO;
OUTPUT OUT=TREND P=RNGTREND;
                                                    /* Subtract trend from each series*/
                                                    /* and convert data to multi-      */
                                                    /* variate form.                   */

PROC SORT DATA=TREND;
  BY YEAR DISKHT;
PROC SORT DATA=TREE94;
  BY YEAR DISKHT;
DATA RINGS94 (KEEP=YEAR RW1-RW13 LR1-LR13 RI1-RI13 RT1-RT13);
  MERGE TREE94 TREND;
  BY YEAR DISKHT;
  RETAIN;
  ARRAY DISK HT1-HT13 (20 18 16 14 12 10 8 5.9 4 2 1.3 0.7 0.3);
  ARRAY RING RW1-RW13;
  ARRAY LOGRING LR1-LR13;
  ARRAY TREND RT1-RT13;
  ARRAY RNGINDX RI1-RI13;
  IF FIRST.YEAR THEN DO I=1 TO 13;
    RING(I)=.;
    LOGRING(I)=.;
    RNGINDX(I)=.;
    TREND(I)=.;
  END;
  DO I=1 TO 13;
    IF DISKHT=DISK(I) THEN DO;
      RING(I)=RNGWIDTH;
      LOGRING(I)=LOGRW;
      TREND(I)=RNGTREND;
      RNGINDX(I)=LOGRING(I)-TREND(I);
    END;
  END;

```

```

END;
IF LAST.YEAR THEN OUTPUT;
/* Prewhiten detrended series for */
/* disks at 2.0 and 1.3 meters. */

PROC SORT DATA=RINGS94;
  BY YEAR;
PROC ARIMA DATA=RINGS94;
  TITLE 'Prewhiten first series (disk at 2.0m)';
  IDENTIFY VAR=RI10 NOPRINT;
  ESTIMATE P=1 NOPRINT;
  FORECAST LEAD=0 BACK=0 OUT=RES10 ID=YEAR;
PROC ARIMA DATA=RINGS94;
  TITLE 'Prewhiten second series (disk at 1.3m)';
  IDENTIFY VAR=RI11 NOPRINT;
  ESTIMATE P=1 NOPRINT;
  FORECAST LEAD=0 BACK=0 OUT=RES11 ID=YEAR;
/* Merge residual data sets. */

DATA PWINDX94;
  MERGE RES10(RENAME=(RESIDUAL=PW10)) RES11(RENAME=(RESIDUAL=PW11));
  BY YEAR;
/* Cross-correlate detrended, */
/* prewhitened series. */

PROC ARIMA DATA=PWINDX94;
  TITLE 'Cross-correlation analysis of prewhitened ring-index series';
  IDENTIFY VAR=PW11 CROSSCOR=PW10;
RUN;

```

The results of the preceding analysis are displayed in Figure 24. The fitted first-order autoregressions that were used to prewhiten the two ring-index series are given in the first two sections (A and B). Output from the cross-correlation analysis (third PROC ARIMA statement) is given next (C). This includes the ACF, IACF, and PACF (C.1–C.3) of the prewhitened ring-index series for the disk at 1.3 m (i.e., the variable named in the VAR option of the IDENTIFY statement) and the corresponding autocorrelation test for randomness (C.4). Notice that all of these checks are consistent with a white noise or a purely random series. This confirms the effectiveness of the detrending and prewhitening procedure. The same checks can be applied to the other series by including another IDENTIFY statement with VAR=PW10.

The key part of the output is the cross-correlation function for the two series (C.5). Inspection of the cross-correlation function reveals a clear spike at a lag of –5 years. This suggests that the best alignment of the two series occurs when the series for the lower disk (PW11) is shifted to the left by 5 years. In other words, the five outer rings of the upper disk are missing in the lower disk and the year of the innermost ring of that disk must be adjusted accordingly. The same analysis can be applied to all pairs of disks and the results used to obtain a corrected estimate of the year of the innermost ring of each disk. The effect of these corrections is illustrated in Figures 2c and 2d.

Prewhiten first series (disk at 1.3m)

ARIMA Procedure

Model for variable RI10

(A)

Estimated Mean = -0.0957115

Autoregressive Factors

Factor 1: 1 0.68642 B**(1)

Prewhiten second series (disk at 2.0m)

ARIMA Procedure

Model for variable RI11

(B)

Estimated Mean = -0.1387615

Autoregressive Factors

Factor 1: 1 - 0.75888 B**(1)

FIGURE 24 *Cross-correlation analysis of missing tree rings.*

Cross-correlation analysis of prewhitened ring-index series (C)

ARIMA Procedure

Name of variable = PW11.

Mean of working series = 0.02743

Standard deviation = 0.372727

Number of observations = 106

Autocorrelations (C.1)

Lag	Covariance	Correlation	-1	9	8	7	6	5	4	3	2	1	0	1	2	3	4	5	6	7	8	9	1
0	0.138925	1.00000												*****									
1	-0.0072260	-0.05201									.	*											
2	-0.0027024	-0.01945									.												
3	0.019128	0.13769									.			***.									
4	-0.0009828	-0.00707									.												
5	-0.0019721	-0.01420									.												
6	-0.0062541	-0.04502									.	*											
7	-0.0014004	-0.01008									.												
8	-0.016739	-0.12049									.	**											
9	0.0010574	0.00761									.												
10	0.0029933	0.02155									.												
11	-0.029235	-0.21043									.	***											
12	-0.0062555	-0.04503									.	*											
13	0.0050624	0.03644									.			*									
14	-0.0076410	-0.05500									.	*											
15	0.0040052	0.02883									.			*									
16	-0.015233	-0.10965									.	**											
17	0.017330	0.12474									.			**									
18	0.009667	0.06958									.			*									
19	-0.0025763	-0.01854									.												
20	-0.0060426	-0.04350									.	*											
21	-0.012409	-0.08932									.	**											
22	0.011612	0.08358									.			**									
23	-0.015190	-0.10934									.	**											
24	-0.0064761	-0.04662									.	*											

``.`` marks two standard errors

FIGURE 24 (continued)

Inverse Autocorrelations

(C.2)

Lag	Correlation	-1	9	8	7	6	5	4	3	2	1	0	1	2	3	4	5	6	7	8	9	1	
1	0.06895											.	*	.									
2	0.00465											.		.									
3	-0.08517										.	**		.									
4	0.04779										.		*	.									
5	0.04779										.		*	.									
6	-0.00019										.			.									
7	0.02669										.		*	.									
8	0.12002										.		**	.									
9	0.00915										.			.									
10	0.02568										.		*	.									
11	0.16505										.		***	.									
12	0.11763										.		**	.									
13	-0.00310										.			.									
14	0.04324										.		*	.									
15	0.02690										.		*	.									
16	0.11930										.		**	.									
17	-0.10980										.	**		.									
18	-0.06740										.	*		.									
19	0.04166										.		*	.									
20	0.06931										.		*	.									
21	0.08527										.		**	.									
22	-0.01543										.			.									
23	0.09676										.		**	.									
24	0.04462										.		*	.									

FIGURE 24 (continued)

Partial Autocorrelations

(C.3)

Lag	Correlation	-1	9	8	7	6	5	4	3	2	1	0	1	2	3	4	5	6	7	8	9	1	
1	-0.05201										.	*		.									
2	-0.02222																	
3	0.13593									.			***	.									
4	0.00676									.			.	.									
5	-0.00971									.			.	.									
6	-0.06641									.	*		.	.									
7	-0.01644																	
8	-0.12308									.	**		.	.									
9	0.01049																	
10	0.02302																	
11	-0.18193									.	****		.	.									
12	-0.07550									.	**		.	.									
13	0.01537																	
14	-0.01694																	
15	0.04037									.	.		*	.									
16	-0.13835									.	***		.	.									
17	0.10932									.	.		**	.									
18	0.06625									.	.		*	.									
19	-0.02336																	
20	-0.09580									.	**		.	.									
21	-0.11291									.	**		.	.									
22	0.03130									.	.		*	.									
23	-0.10579									.	**		.	.									
24	-0.05140									.	*		.	.									

Autocorrelation Check for White Noise

(C.4)

To	Chi	Autocorrelations							
Lag	Square	DF	Prob						
6	2.70	6	0.845	-0.052	-0.019	0.138	-0.007	-0.014	-0.045
12	10.06	12	0.611	-0.010	-0.120	0.008	0.022	-0.210	-0.045
18	14.86	18	0.671	0.036	-0.055	0.029	-0.110	0.125	0.070
24	19.14	24	0.745	-0.019	-0.043	-0.089	0.084	-0.109	-0.047

FIGURE 24 (continued)

Correlation of PW11 and PW10
 Variance of input = 0.112964
 Number of observations = 99

(C.5)

Crosscorrelations

Lag	Covariance	Correlation	-1	9	8	7	6	5	4	3	2	1	0	1	2	3	4	5	6	7	8	9	1	
-21	-0.014955	-0.12510										***												
-20	0.0034147	0.02856										.	*											
-19	-0.0016700	-0.01397										.												
-18	-0.0021975	-0.01838										.												
-17	-0.019465	-0.16283										***												
-16	-0.018676	-0.15623										***												
-15	0.017216	0.14401										.	***											
-14	-0.014613	-0.12224										**												
-13	-0.010219	-0.08549										**												
-12	0.0020309	0.01699										.												
-11	-0.012335	-0.10319										**												
-10	0.0018776	0.01571										.												
-9	0.0066333	0.05549										.	*											
-8	0.0031654	0.02648										.	*											
-7	0.0068038	0.05692										.	*											
-6	0.0064057	0.05359										.	*											
-5	0.060589	0.50684										.	*****											
-4	-0.0094345	-0.07892										**												
-3	0.020050	0.16772										.	***											
-2	0.022023	0.18423										.	****											
-1	0.00031661	0.00265										.												
0	-0.0004457	-0.00373										.												
1	0.0049424	0.04134										.	*											
2	0.0018182	0.01521										.												
3	-0.012707	-0.10630										**												
4	-0.0063191	-0.05286										.	*											
5	0.0056615	0.04736										.	*											
6	-0.029310	-0.24519										****												
7	0.020149	0.16855										.	***											
8	-0.026690	-0.22327										****												
9	0.0019742	0.01651										.												
10	-0.0020306	-0.01699										.												
11	-0.012346	-0.10327										**												
12	0.0063665	0.05326										.	*											
13	0.011237	0.09400										.	**											
14	-0.0013013	-0.01089										.												
15	-0.010942	-0.09153										**												
16	-0.0004128	-0.00345										.												
17	-0.0037202	-0.03112										.	*											
18	-0.022580	-0.18889										****												
19	-0.00034	-0.00284										.												
20	0.0064959	0.05434										.	*											
21	-0.013566	-0.11348										**												

``.`` marks two standard errors

FIGURE 24 (concluded).

Repeated measurements of the same experimental unit or process are often correlated and require special methods of analysis. If the number of observation times is relatively small and the objective is to compare trends for several treatment groups, then repeated-measures ANOVA is applicable. There are two types of repeated-measures ANOVA—one is based on a split-plot univariate model and the other on a multivariate model. The multivariate model is less restrictive than the univariate model, but the associated statistical tests may lack power because there are more parameters to estimate and therefore fewer degrees of freedom. When the data do not satisfy the correlation assumptions of the univariate model, another alternative is to apply the Greenhouse-Geisser or Huynh-Feldt correction to the univariate tests. Both univariate and multivariate repeated-measures analyses are available in `PROC GLM` and `PROC ANOVA` of the SAS system. They are requested with the `REPEATED` statement, which also provides a range of options for comparing polynomial trends and other contrasts over time.

Time series have various properties that are of interest to researchers, including the presence or absence of trend, seasonality, and other types of nonstationarity, and for stationary time series, the presence or absence of short-term autocorrelation. Many statistical methods are available for the analysis of univariate time series. The time plot, correlogram, and sample PACF and IACF are useful for describing a time series, and are valuable tools for outlier detection, model identification, and verification. Trend and cyclic components can be described or removed by regression methods, or by smoothing the series. If the objective is model development or forecasting, then the Box-Jenkins procedure can be used to fit and predict stationary and nonstationary ARIMA processes. Alternative forecasting methods include stepwise autoregression, exponential smoothing, and the Holt-Winters method. The SAS/ETS library provides a full range of time-series analyses. There are three main procedures for analyzing time series in the time domain: `PROC ARIMA`, `PROC AUTOREG`, and `PROC FORECAST`, and a fourth procedure, `PROC EXPAND`, can be used to aggregate, interpolate, or otherwise manipulate a time series in preparation for analysis.

APPENDIX 1 Average height of seedlings

Spp.	Stk.	Trt.	Blk.	Row	Average height (cm) ^a						
					Initial	1984	1985	1986	1987	1988	1989
FD	B	S	1	12	18.78	22.89	22.44	22.89	22.44	27.56	33.56
FD	B	S	2	10	15.92	20.08	21.50	24.75	28.42	39.67	53.67
FD	B	S	3	10	20.80	26.60	28.20	27.90	36.90	48.30	59.80
FD	B	S	4	4	18.60	22.60	21.40	23.00	22.20	30.60	39.80
FD	B	U	1	1	14.10	18.40	20.80	24.70	29.00	40.80	57.70
FD	B	U	2	12	17.00	21.00	22.00	18.00	20.00	22.00	25.00
FD	B	U	3	9	18.00	22.22	25.22	27.56	29.33	38.67	49.67
FD	B	U	4	9	18.14	23.29	25.36	28.36	30.07	38.00	48.50
FD	B	V	1	10	16.14	19.81	22.10	25.95	33.43	45.76	59.19
FD	B	V	2	2	14.89	19.28	21.11	25.78	30.28	39.89	53.83
FD	B	V	3	5	15.08	18.08	19.77	23.08	27.08	37.38	52.08
FD	B	V	4	8	19.00	23.06	24.24	28.12	34.47	45.35	58.12
FD	P	S	1	5	18.94	26.71	31.06	33.24	42.00	54.12	65.94
FD	P	S	2	8	22.82	28.91	34.05	39.91	47.32	59.50	75.64
FD	P	S	3	3	22.90	28.15	32.95	39.05	49.20	62.60	74.85
FD	P	S	4	5	20.59	25.65	30.12	33.35	39.53	49.29	61.65
FD	P	U	1	8	21.56	29.22	31.17	38.00	43.83	54.72	66.78
FD	P	U	2	5	20.47	27.11	33.16	39.95	46.74	57.37	70.79
FD	P	U	3	6	19.05	27.52	31.33	37.95	46.62	56.95	68.52
FD	P	U	4	3	16.29	24.71	31.05	35.76	43.48	55.62	70.86
FD	P	V	1	6	18.08	23.63	28.54	37.08	47.83	64.75	86.75
FD	P	V	2	6	20.88	26.58	30.50	35.88	42.83	53.71	70.58
FD	P	V	3	2	20.19	25.94	28.38	32.38	37.06	48.63	67.63
FD	P	V	4	11	20.40	26.25	30.00	34.25	38.35	49.05	65.30
PL	B	S	1	11	12.44	20.28	34.67	49.83	65.78	90.83	125.17
PL	B	S	2	7	16.33	22.00	32.00	47.33	58.33	81.73	113.07
PL	B	S	3	4	11.32	20.05	32.45	47.45	66.77	96.32	132.55
PL	B	S	4	7	11.11	16.37	28.74	44.11	65.79	94.05	130.74
PL	B	U	1	7	12.40	17.40	26.50	36.80	49.60	75.80	109.90
PL	B	U	2	11	13.44	18.44	24.75	41.00	54.31	80.00	111.56
PL	B	U	3	1	14.19	19.44	29.63	48.38	62.25	86.69	120.31
PL	B	U	4	10	15.22	20.61	29.61	41.83	55.06	70.72	97.78
PL	B	V	1	4	12.31	17.08	26.38	41.46	64.08	100.38	147.62
PL	B	V	2	4	13.94	19.00	28.41	46.41	81.94	118.41	166.18
PL	B	V	3	8	11.53	18.58	29.68	45.47	70.89	105.47	150.32
PL	B	V	4	2	12.63	16.63	27.06	43.75	67.38	103.06	146.75
PL	P	S	1	3	12.43	23.19	36.71	55.29	75.71	109.48	155.76
PL	P	S	2	9	10.23	18.59	33.91	53.59	74.09	108.27	150.64
PL	P	S	3	12	9.59	17.82	32.05	49.86	69.50	97.59	133.55
PL	P	S	4	1	13.48	21.70	34.26	48.22	73.39	103.83	141.48
PL	P	U	1	2	12.00	22.86	34.38	49.00	71.10	105.05	148.71
PL	P	U	2	3	9.43	17.14	30.10	43.33	60.95	87.24	125.67
PL	P	U	3	7	8.15	15.95	28.60	39.65	58.75	89.00	129.40
PL	P	U	4	12	8.75	15.70	27.45	42.55	58.45	85.55	123.85
PL	P	V	1	9	12.28	19.52	33.12	55.12	89.24	136.16	193.56
PL	P	V	2	1	9.57	17.13	28.74	46.65	74.00	114.22	163.13
PL	P	V	3	11	10.25	17.83	29.38	48.00	78.88	116.29	161.50
PL	P	V	4	6	7.83	13.58	30.00	53.42	84.71	130.38	186.21

^a Data provided by T. Newsome and N. Daintith, B.C. Ministry of Forests, Cariboo Forest Region.

APPENDIX 2 Ring widths

Year	Ring width (cm) ^a		Year	Ring width (cm)	
	Disk at 2.0 m	Disk at 1.3 m		Disk at 2.0 m	Disk at 1.3 m
1888	.	0.090	1941	0.110	0.067
1889	.	0.049	1942	0.063	0.067
1890	.	0.132	1943	0.075	0.046
1891	.	0.137	1944	0.063	0.042
1892	.	0.175	1945	0.072	0.062
1893	.	0.135	1946	0.079	0.024
1894	.	0.114	1947	0.067	0.032
1895	0.090	0.148	1948	0.059	0.091
1896	0.153	0.243	1949	0.051	0.101
1897	0.096	0.211	1950	0.066	0.100
1898	0.157	0.215	1951	0.039	0.126
1899	0.168	0.222	1952	0.029	0.125
1900	0.270	0.211	1953	0.080	0.108
1901	0.264	0.173	1954	0.076	0.101
1902	0.209	0.141	1955	0.059	0.093
1903	0.243	0.203	1956	0.068	0.036
1904	0.261	0.196	1957	0.045	0.033
1905	0.196	0.281	1958	0.059	0.053
1906	0.186	0.268	1959	0.062	0.020
1907	0.260	0.279	1960	0.082	0.025
1908	0.304	0.246	1961	0.023	0.023
1909	0.323	0.164	1962	0.027	0.019
1910	0.305	0.240	1963	0.032	0.015
1911	0.265	0.148	1964	0.024	0.009
1912	0.330	0.127	1965	0.016	0.019
1913	0.225	0.222	1966	0.014	0.031
1914	0.219	0.207	1967	0.013	0.067
1915	0.223	0.127	1968	0.024	0.036
1916	0.163	0.138	1969	0.017	0.085
1917	0.211	0.136	1970	0.032	0.068
1918	0.166	0.134	1971	0.025	0.062
1919	0.143	0.108	1972	0.062	0.061
1920	0.132	0.093	1973	0.042	0.059
1921	0.145	0.108	1974	0.048	0.051
1922	0.159	0.110	1975	0.051	0.038
1923	0.136	0.121	1976	0.072	0.032
1924	0.065	0.102	1977	0.037	0.045
1925	0.094	0.103	1978	0.032	0.019
1926	0.098	0.089	1979	0.040	0.029
1927	0.093	0.102	1980	0.035	0.023
1928	0.073	0.072	1981	0.029	0.016
1929	0.097	0.067	1982	0.044	0.026
1930	0.070	0.099	1983	0.022	0.029
1931	0.047	0.109	1984	0.022	0.026
1932	0.073	0.085	1985	0.026	0.022
1933	0.068	0.101	1986	0.023	0.021
1934	0.067	0.105	1987	0.015	0.009
1935	0.080	0.118	1988	0.009	0.006
1936	0.092	0.083	1989	0.012	0.005
1937	0.088	0.051	1990	0.011	0.004
1938	0.103	0.078	1991	0.013	0.007
1939	0.083	0.062	1992	0.008	0.012
1940	0.138	0.089	1993	0.005	0.006

^a Data provided by I. Cameron, B.C. Ministry of Forests, Research Branch.

APPENDIX 3 Ring index and rainfall

Year	Ring index ^a	Spring rainfall (mm)	Year	Ring index ^a	Spring rainfall (mm)
1891	1.153	209.2	1942	1.299	67.7
1892	1.296	232.1	1943	1.168	244.9
1893	1.428	343.4	1944	1.228	108.7
1894	1.432	356.0	1945	1.020	175.1
1895	1.019	156.3	1946	1.316	201.2
1896	0.976	137.5	1947	1.112	133.6
1897	0.893	191.6	1948	1.425	221.5
1898	1.032	96.0	1949	1.229	96.5
1899	1.133	264.1	1950	1.077	212.1
1900	1.286	219.8	1951	0.912	112.8
1901	1.201	187.4	1952	0.975	104.9
1902	1.139	164.3	1953	0.885	80.7
1903	1.228	192.2	1954	0.842	81.2
1904	1.114	194.3	1955	1.079	160.2
1905	1.084	150.7	1956	0.697	111.3
1906	1.014	99.7	1957	0.977	152.6
1907	0.825	106.9	1958	1.111	119.6
1908	0.785	237.6	1959	0.748	183.1
1909	0.700	71.4	1960	0.927	180.3
1910	0.718	164.9	1961	0.891	138.7
1911	0.655	147.0	1962	0.831	137.9
1912	0.681	145.9	1963	1.086	88.1
1913	0.759	116.3	1964	0.916	116.7
1914	0.912	111.4	1965	0.859	105.2
1915	0.743	114.4	1966	0.803	98.3
1916	0.895	251.2	1967	0.968	149.6
1917	0.836	280.6	1968	0.745	203.2
1918	0.616	189.8	1969	0.881	185.2
1919	0.771	158.0	1970	0.836	210.6
1920	0.873	166.3	1971	1.043	164.6
1921	0.847	129.4	1972	1.122	64.0
1922	0.837	102.3	1973	0.802	60.2
1923	0.739	131.3	1974	0.945	149.2
1924	0.769	63.7	1975	0.755	76.6
1925	0.715	100.3	1976	1.043	57.3
1926	0.617	108.8	1977	1.083	168.2
1927	0.678	123.2	1978	0.937	157.4
1928	0.816	160.2	1979	0.820	93.7
1929	0.888	185.0	1980	1.209	163.0
1930	0.861	153.2	1981	1.387	172.0
1931	1.100	159.4	1982	1.035	89.7
1932	1.203	194.6	1983	1.240	164.2
1933	1.241	92.3	1984	1.451	239.9
1934	1.373	212.8	1985	1.152	140.7
1935	1.021	139.1	1986	0.969	162.2
1936	1.208	110.8	1987	0.948	177.9
1937	1.427	139.4	1988	0.918	241.9
1938	1.144	114.7	1989	0.778	190.8
1939	1.264	116.2	1990	0.945	177.6
1940	1.326	162.6	1991	1.023	199.2
1941	0.991	116.9	1992	1.221	105.2

^a Data provided by D. Spittlehouse, Research Branch.

- Bergerud, W. 1988. Determining polynomial contrast coefficients. B.C. Min. For., Res. Branch, Victoria, B.C. Biom. Inf. Pam. No. 12.
- . 1991. Analysing a split-plot in time with the proper repeated measures ANOVA. B.C. Min. For., Res. Branch, Victoria, B.C. Biom. Inf. Pam. No. 32.
- Bloomfield, P. 1976. Fourier analysis of time series: an introduction. Wiley, New York, N.Y.
- Box, G.P. and G.M. Jenkins. 1976. Time series analysis: forecasting and control. Revised ed. Holden-Day, Oakland, Calif.
- Cameron, I.R. 1993. Correcting ring counts for missing rings in paper birch. Ecology and Management of B.C. Hardwoods Workshop, Dec. 1–2, Richmond, B.C. Poster pap.
- Chatfield, C. 1979. Inverse autocorrelations. J. R. Stat. Soc. A. 142:363–77.
- . 1992. The analysis of time series. 4th ed. Chapman and Hall, London.
- Cody, R.P. and J.K. Smith. 1987. Applied statistics and the SAS programming language. 2nd ed. North-Holland, New York, N.Y.
- Diggle, P.J. 1991. Time series: A biostatistical introduction. Oxford University Press, New York, N.Y.
- Durbin, J. and G.S. Watson. 1951. Testing for serial correlation in least squares regression: II. *Biometrika* 38:159–77.
- Gumpertz, M.L. and C. Brownie. 1993. Repeated measures in randomized block and split-plot experiments. *Can. J. For. Res.* 23:625–39.
- Hand, D.J. and C.C. Taylor. 1987. Multivariate analysis of variance and repeated measures. Chapman and Hall, London.
- International Tree-Ring Data Bank. 1993. International tree-ring data bank (ITRDB) programs. Lab. Tree-Ring Res., Univ. Ariz., Tucson, Ariz. (Dr. Harold C. Fritts, Chairman of the ITRDB).
- Jenkins, G.M. and D.G. Watts. 1968. Spectral analysis and its applications. Holden-Day, Oakland, Calif.
- Kendall, M. and J.K. Ord. 1990. Time series. 3rd ed. Oxford University Press, New York, N.Y.
- Kendall, M., A. Stuart, and J.K. Ord. 1983. The advanced theory of statistics. 4th ed. Griffin & Co., London. Vol. 3.
- Keppel, G. 1973. Design and analysis: a researcher's handbook. Prentice-Hall, Englewood Cliffs, N.J.
- Levenbach, H. and J.P. Cleary. 1981. The beginning forecaster: the forecasting process through data analysis. Lifetime Learning Publ., Belmont, Calif.
- Littell, R.C. 1989. Statistical analysis of experiments with repeated measures. *HortScience* 24:37–40.

- Looney, S.W. and W.B. Stanley. 1989. Exploratory repeated measures analysis for two or more groups. *Am. Stat.* 43:220–25.
- McCleary, R. and R.A. Hay, Jr. 1980. *Applied time series analysis for the social sciences*. Sage Publ., London.
- Meredith, M.P. and S.V. Stehman. 1991. Repeated measures experiments in forestry: focus on analysis of response curves. *Can. J. For. Res.* 21:957–65.
- Milliken, G.A. and D.E. Johnson. 1992. *Analysis of messy data. Vol. I: Designed experiments*. Chapman and Hall, New York, N.Y.
- Monserud, R.A. 1986. Time-series analysis of tree-ring chronologies. *For. Sci.* 32:349–72.
- Morrison, D.F. 1976. *Multivariate statistical methods*. 2nd ed. Macmillan (Griffin), New York, N.Y.
- Moser, E.B., A.M. Saxton, and S.R. Pezeshki. 1990. Repeated measures analysis of variance: application to tree research. *Can. J. For. Res.* 20:524–35.
- Nemec, A.F.L. 1992. Guidelines for the statistical analysis of forest vegetation management data. B.C. Min. For., Res. Br., Victoria, B.C., *Biom. Inf. Handb.* No. 3.
- Rossi, R.E., D.J. Mulla, A.G. Journel, and E.H. Franz. 1992. Geostatistical tools for modelling and interpreting ecological spatial dependence. *Ecol. Monogr.* 62:277–314.
- SAS Institute Inc. 1989. *SAS/STAT user's guide, Version 6, 4th ed., Vols. 1 and 2*, SAS Institute Inc., Cary, N.C.
- . 1990a. *SAS procedures guide, Version 6, 3rd ed.* SAS Institute Inc., Cary, N.C.
- . 1990b. *SAS language: reference, Version 6, 1st ed.* SAS Institute Inc., Cary, N.C.
- . 1991a. *SAS/ETS user's guide, Version 6, 1st ed.* SAS Institute Inc., Cary, N.C.
- . 1991b. *SAS/ETS software: applications guide 1—time series modelling and forecasting, financial reporting, and loan analysis, Version 6, 1st ed.* SAS Institute Inc., Cary, N.C.
- Sit, V. 1992a. A repeated measures example. B.C. Min. For., Res. Branch, Victoria, B.C., *Biom. Inf. Pam.* No. 39.
- . 1992b. Finding the expected mean squares and the proper error terms with SAS. B.C. Min. For., Res. Branch, Victoria, B.C., *Biom. Inf. Pam.* No. 40.
- Tabachnick, B.G. and L.S. Fidell. 1989. *Using multivariate statistics*. 2nd ed. Harper Collins, New York, N.Y.