CHAPTER 18 Time-Series Analysis

18.1 General Purpose and Description

Time-series analysis is used when observations are made repeatedly over 50 or more time periods. Sometimes the observations are from a single case, but more often they are aggregate scores from many cases. For example, the scores might represent the daily number of temper tantrums of a two-year-old, the weekly output of a manufacturing plant, the monthly number of traffic tickets issued in a municipality, or the yearly GNP for a developing country, all of these tracked over considerable time. One goal of the analysis is to identify patterns in the sequence of numbers over time, which are correlated with themselves, but offset in time. Another goal in many research applications is to test the impact of one or more interventions (IVs). Time-series analysis is also used to forecast future patterns of events or to compare series of different kinds of events.

As in other regression analyses, a score is decomposed into several potential elements. One of the elements is a random process, called a shock. Shocks are similar to the error terms in other analyses. Overlaying this random element are numerous potential patterns. One potential pattern is trends over time: linear (where the mean is steadily increasing or decreasing over time), quadratic (where the mean first increases and then decreases over time, or the reverse), or something more complicated. A second potential pattern is lingering effects of earlier scores, and a third potential pattern is lingering effects of earlier shocks. These patterns are not mutually exclusive; two or all three can be superimposed on the random process.

The model described in this chapter is *a*uto-*r*egressive, *i*ntegrated, *m*oving *a*verage, called an ARIMA (p, d, q) model. The auto-regressive element, p, represents the lingering effects of preceding scores. The integrated element, d, represents trends in the data, and the moving average element, q, represents the lingering effects of preceding random shocks. A big question is how lingering is lingering? That is, do you have to take into account just the previous score (or shock) or do you get a better model if you take into account two or more of the previous scores (or shocks)?

The first three steps in the analysis, identification, estimation, and diagnosis, are devoted to modeling the patterns in the data. The first step is *identification* in which autocorrelation functions (ACFs) and partial autocorrelation functions (PACFs) are examined to see which of the potential three patterns are present in the data. Autocorrelations are self-correlations of the series of scores with itself, removed one or more periods in time; partial autocorrelations are self-correlations with intermediate autocorrelations partialed out. Various auto-regressive and moving average patterns leave distinctive footprints on the autocorrelation and partial autocorrelation functions.

When the time series is long, there are also tendencies for measures to vary periodically, called seasonality, periodicity, or cyclic in time-series jargon. For example, viral infections peak during the winter months, as do calories and alcohol consumed. Thus, seasonality is another form of autocorrelation frequently seen in data sets. Periodic variation occurs over shorter time periods as well. For example, quality of manufacturing differs over days of the week, peaking in mid-week. And calorie and alcohol consumption also increase over the weekend. These patterns are also identified using ACFs and PACFs and accounted for in the model.

Time-series analysis is more appropriate for data with autocorrelation than, say, multiple regression, for two reasons. The first is that there is explicit violation of the assumption of independence of errors. The errors are correlated due to the patterns over time in the data. Type I error rate is substantially increased if regression is used when there is autocorrelation. The second is that the patterns may either obscure or spuriously enhance the effect of an intervention unless accounted for in the model.

The second step in modeling the series is *estimation* in which the estimated size of a lingering auto-regressive or moving average effect is tested against the null hypothesis that it is zero. The third step is *diagnosis*, in which residual scores are examined to determine if there are still patterns in the data that are not accounted for. Residual scores are the differences between the scores predicted by the model and the actual scores for the series. If all patterns are accounted for in the model, the residuals are random. In many applications of time series, identifying and modeling the patterns in the data are sufficient to produce an equation, which is then used to predict the future of the process. This is called *forecasting,* the goal of many applications of time series in the economic arena.

However, often the goal is to assess the impact of an intervention. The intervention is occasionally experimental, featuring random assignment of cases to levels of treatment, control of other variables, and manipulation of the levels of the IV by the researcher. However, true experiments typically assess the DV far fewer than 50 times and therefore cannot use time-series analysis. More likely with experiments, there is a single measure before the intervention (a pretest) and, perhaps, a few measures after the intervention to assess immediate and follow-up effects. In this situation, repeated measures ANOVA is more appropriate as long as heterogeneity of covariance (sphericity) is taken into account, as discussed in Chapter 8, or MLM if complications arise (cf. Chapter 15).

In time series, the intervention is more likely to be either naturally occurring or quasi-experimental. With naturally occurring interventions, such as an assassination or a natural disaster, “nature” manipulates the timing and delivery of the intervention. In quasi-experiments, the timing and delivery of the intervention are under the control of the researcher, but there is often insufficient control of extraneous variables, including cases, to consider the study a true experiment. Similarly, there is often no control group with cases randomly assigned. For example, a profit-sharing plan is introduced to employees at a single company at a specific time and quality of computers produced is compared before and after the plan. Numerous observations of the quality of the product are made both before and after the introduction of the profit-sharing plan. It is better if the time for introduction of the plan is randomly chosen from among the population of possible times.

An intervention is tested for significance in the traditional manner, once the endogenous patterns in the data have been reduced to error through modeling. The effects of interventions considered in this chapter vary in both onset and duration; the onset of the effect may be abrupt or gradual and the duration of the effect may be either permanent or temporary. Effects of interventions are superimposed on the other patterns in the data. There are numerous other possible types of interventions and tests for them, as discussed in McCleary and Hay (1980) and McDowall, McCleary, Meidinger, and Hay (1980), but beyond the scope of this book.

There may be more than a single intervention. For example, a profit-sharing plan is implemented at one point, and a child-care program implemented at another point in time. Or, the same intervention might be implemented and withdrawn at different time points. For example, background music is introduced into the assembly room for some time period, withdrawn for another time period, reintroduced, withdrawn again, and so forth, in the typical ABAB reversal pattern. There might also be a covariate measured at all time periods along with the DV. For example, temperature inside the plant is measured along with product quality, both before and after implementation of the profit-sharing program. The trick is to find a covariate that is not affected by the intervention. Or, additional DVs might be measured, creating a multivariate time-series analysis. For example, productivity might be assessed as well as quality.

Nunn (1993) investigated the impact of putting a mobile digital terminal, which allows patrol officers to communicate directly with remote crime information databases, into police vehicles. DVs were theft recoveries and clearance rates over a ten-year period from January 1980 to January 1990. The terminals were installed from March 1985 through January 1986. There was little evidence that the intervention affected auto-theft clearances. Recoveries did not directly show the effect of the intervention (i.e., did not increase significantly), but the trend changed after intervention, reflecting a greater number of auto thefts nationwide. Nunn suggests that the mobile digital terminal may have helped police hold the line against a rapid drop in percentage of recoveries expected with the greater number of thefts; they were able to continue to recover a constant percentage of stolen vehicles.

Lin and Crawford (1983) compared yearly mortality patterns in three communities from 1890 to 1967. They found that populations that were closer geographically showed remarkable similarities in mortality patterns. They also found evidence to suggest that mortality patterns began cycling due to pandemic diseases at the turn of the twentieth century with greater communication and transportation, but patterns again became distinctive later in the century with advances in medical technology.

Time-series analysis has its own unique jargon and sometimes uses familiar terms in ways that are different from uses in other statistical techniques. Table 18.1 defines some time-series terms as they are used in this chapter. Many of the terms are defined algebraically in Section 18.4.

This chapter provides only a simplified overview of the complicated data analysis strategy that is time-series analysis. A recent update of the classic reference for time series analysis is available (Box, Jenkins, & Reinsel, 1994), supplying a comprehensive treatment of the topic. Another recent resource for more advanced applications is provided by Hershberger, Molenaar, and Corneal (1996).

18.2 Kinds of Research Questions

The major research questions involve the patterns in the series, the predicted value of the scores in the near future, and the effect of an intervention (an IV). Less common questions address the relationships among time series. It should be understood that this chapter barely scratches the surface of the complex world of time-series analysis. Only those questions that are relatively easily addressed in SPSS and SAS are discussed.

TABLE 18.1 Some Time-Series Terminology

Term	Definition
Observation	The DV score at one time period. The score can be from a single case or an aggregate score from numerous cases.
Random shock	The random component of a time series. The shocks are reflected by the residuals (or errors) after an adequate model is identified.
ARIMA (p, d, q)	The acronym for an auto-regressive integrated moving average model. The three terms to be estimated in the model are auto-regressive (p), integrated (trend—d), and moving average (q).
Auto-regressive terms (p)	The number of terms in the model that describe the dependency among successive observations. Each term has an associated correlation coefficient that describes the magnitude of the dependency. For example, a model with two auto-regressive terms ($p = 2$) is one in which an observation depends on (is predicted by) two previous observations.
Moving average terms (q)	The number of terms that describe the persistence of a random shock from one observation to the next. A model with two moving average terms ($q = 2$) is one in which an observation depends on two preceding random shocks.
Lag	The time periods between two observations. For example, lag 1 is between Y_t and Y_{t-1}. Lag 2 is between Y_t and Y_{t-2}. Time series can also be lagged forward, Y_t and Y_{t+1}.
Differencing	Calculating differences among pairs of observations at some lag to make a nonstationary series stationary.
Stationary and nonstationary series	Stationary series vary around a constant mean level, neither decreasing nor increasing systematically over time, with constant variance. Nonstationary series have systematic trends, such as linear, quadratic, and so on. A nonstationary series that can be made stationary by differencing is called "nonstationary in the homogenous sense."
Trend terms (d)	The terms needed to make a nonstationary time series stationary. A model with two trend terms ($d = 2$) has to be differenced twice to make it stationary. The first difference removes linear trend, the second difference removes quadratic trend, and so on.
Autocorrelation	Correlations among sequential scores at different lags. The lag 1 autocorrelation coefficient is similar to correlation between the pairs of scores at adjacent points in time, $r_{Yt,Yt-1}$ (e.g., the pair at time 1 and time 2, the pair at time 2 and time 3, and so on). The lag 2 autocorrelation coefficient is similar to correlation between the pairs of scores two time periods apart, $r_{Yt,Yt-2}$ (e.g., the pair at time 1 and time 3, the pair at time 2 and time 4, and so on).
Autocorrelation function (ACF)	The pattern of autocorrelations in a time series at numerous lags; the correlation at lag 1, then the correlation at lag 2, and so on.
Partial autocorrelation function (PACF)	The pattern of partial autocorrelations in a time series at numerous lags after partialing out the effects of autocorrelations at intervening lags.

18.2.1 Pattern of Autocorrelation

The pattern of autocorrelation is modeled in any time-series study, for itself, in preparation for forecasting, or prior to tests of an intervention. Are there linear or quadratic trends in the data? Does the previous score affect the current one? The previous random shock? How quickly do autocorrelations die out over time? For the example, is the quality of the computer increasing steadily over the time frame? Decreasing? Is the quality of the computer produced in one time frame associated with the quality in the next time frame? How long do the random shocks in the manufacturing processes linger? Section 18.4.1.5 shows how the autocorrelation functions (ACFs) and partial autocorrelation functions (PACFs) are examined to reveal these patterns.

18.2.2 Seasonal Cycles and Trends

Time-series data are also examined for seasonal cycles if such are possible. Are there weekly, quarterly, monthly, or yearly trends in the data? For example, does the quality of the computers produced vary systematically over the days of the week? The weeks of a month? As for auto-regressive and moving average components, ACFs and PACFs are examined to reveal seasonal cycles, as demonstrated in Section 18.5.1.

18.2.3 Forecasting

Based on the known patterns in the data, what is the predicted value of observations in the near future? For example, based on previous patterns in the data, is the quality of the computers likely to increase in the next month? Decrease? Forecasting is a major enterprise in business and economics. It is discussed briefly in Section 18.6.3.

18.2.4 Effect of an Intervention

Has an intervention had an impact, after taking into account patterns in the scores associated with trends, auto-regression, moving averages, and periodicity? For example, is there a difference in quality of computers after introduction of a profit-sharing plan? Intervention is added to a model as an IV. Sections 18.5.2 and 18.7.4 demonstrate tests of intervention in time-series analysis.

Procedures also are available for determining the onset and duration of the effects of an intervention—whether they are abrupt or gradual, permanent or temporary. These are discussed in Section 18.5.2 and also in McCleary and Hay (1980) and McDowall et al. (1980), among others.

18.2.5 Comparing Time Series

Are the patterns over time similar for different variables or populations? For example, do income and consumer prices follow the same time series? Do different populations have the same patterns of mortality? That is, what are the relationships among two or more time series? This often is referred to as multivariate time series; terms also used are cross-correlation functions, transfer function models, models with input series, and dynamic regression. These models are similar to intervention models in that an IV (in this case a continuous IV) is added to a model. These models are discussed by McCleary and Hay (1980) and Cromwell, Hannan, Labys, and Terraza (1994). They also are discussed in the on-disk SAS/ETS User's Guide, with an example shown (Example 11.3). Section 18.5.3 discusses time-series analysis with more than one continuous variable.

18.2.6 Time Series with Covariates

Covariates (often called predictors in time-series jargon) may be measured along with the DV. A first question is: Is the covariate related to the DV, after adjusting both for autocorrelation and periodicity? For example, is the temperature in the manufacturing plant related to product quality? If the answer is positive, then including the covariate in the model may enhance the test of the intervention. Section 18.5.3 discusses models with covariates.

18.2.7 Effect Size and Power

How much of the variability in observations is due to the chosen model? Section 18.6.2 shows two effect size measures used in time-series analysis. Power depends on the accuracy of ARIMA modeling, as well as on the number of observations over time and the impact of the intervention.

18.3 Assumptions of Time-Series Analysis

18.3.1 Theoretical Issues

The usual cautions apply with regard to causal inference in time-series analysis. Only with random assignment to treatment conditions, control of extraneous variables, and manipulation of the intervention(s) can cause reasonably be inferred when differences associated with an intervention are observed. Quasi-experiments are designed to rule out as many alternative sources of influence on the DV as possible, but causal inference is much weaker in any design that falls short of the requirements of a true experiment.

18.3.2 Practical Issues

The random shocks that perturb the system are considered to be independent and normally distributed with mean zero and constant variance over time. Contingencies among *scores* over time are part of the model that is developed during identification and estimation. If the model is good, all sequential contingencies are removed so that you are left with the randomly distributed shocks. The residuals, then, are a reflection of the random shocks: independent and normally distributed, with mean zero and homogeneity of variance. It is expressly assumed that there *are* correlations in the sequence of observations over time that have been adequately modeled. This assumption is tested during the diagnostic phase when remaining, as yet unaccounted for patterns are sought among the residuals. Outliers among scores are sought before modeling and among the residuals once the model is developed.

18.3.2.1 Normality of Distributions of Residuals

A model is developed and then normality of residuals is evaluated in time-series analysis. *Examine the normalized plot of residuals for the model before evaluating an intervention. Transform the DV if residuals are nonnormal.* The normalized plot of residuals is examined as part of the diagnostic phase of modeling, as discussed in Section 18.4.3 and demonstrated in Section 18.7.3. The usual square root, logarithmic, or inverse transformations are appropriate in the event of nonnormally distributed residuals.

18.3.2.2 Homogeneity of Variance and Zero Mean of Residuals

After the model is developed, *examine plots of standardized residuals versus predicted values to assess homogeneity of variance over time. Consider transforming the DV if the width of the plot varies over the predicted values.* McCleary and Hay (1980) recommend a logarithmic transformation to remedy heterogeneity of variance.

18.3.2.3 Independence of Residuals

During the diagnostic phase, once the model is developed and residuals are computed, there should be no remaining autocorrelations or partial autocorrelations at various lags in the ACFs and PACFs. Remaining autocorrelations at various lags signal other possible patterns in the data that have not been properly modeled. *Examine the ACFs and PACFs for other patterns and adjust the model accordingly.*

18.3.2.4 Absence of Outliers

Outliers are observations that are highly inconsistent with the remainder of the time-series data. They can greatly affect the results of the analysis and must be dealt with. They sometimes show up in the original plot of the DV against time, but are often more noticeable after initial modeling is complete. *Examine the time-series plot before and after adjusting for autocorrelation and seasonality to identify obvious outliers.* Unfortunately, there are no concrete guidelines to determine how discrepant a case must be to be labeled an outlier in a time-series analysis (Cryer, 1986, p. 250). An outlier is dealt with in the usual manner by checking the original data for errors, deleting the observation, replacing the observation with an imputed value, and so on. Section 18.7.1 demonstrates a search for outliers.

Outliers also may be sought in the solution. SAS ARIMA has a procedure to detect changes in the level of the response series that are not accounted for by the model. This also is demonstrated in Section 18.7.1.

18.4 Fundamental Equations for Time-Series ARIMA Models

Like most multivariate procedures, time-series analysis is done by computer, not by hand. Many of the computations are not difficult so much as extremely tedious if done by hand. Full and partial autocorrelations between pairs of scores at 25 to 30 different lags? Therefore, the emphasis in this section is on conceptual understanding rather than processing of a data set.

Several texts are devoted to time-series analysis, some of them highly mathematical. The primary reference for ARIMA models—the ones addressed in this chapter—is Box, Jenkins, and Reinsel (1994). Two texts that demonstrate at least some of the equations with numbers are Glass, Wilson, and Gottman (1975) and McDowall et al. (1980). A few less mathematical, more computer-oriented sources are Cryer (1986); McCleary and Hay (1980); and McCain and McCleary (1979). The notation from the latter two sources has been adapted for this section.

Observations are made repeatedly through the duration of a study. It is the order of the observations that is important. If there is an intervention, many observations are made before it and many after. For this small-sample example, computer quality is measured weekly for 20 weeks on a scale where higher numbers indicate higher quality. The observations appear in Table 18.2. You may think

TABLE 18.2 Observations of Computer Quality over 20 Weeks

Week	Quality
1	19
2	21
3	17
4	19
5	20
6	21
7	27
8	28
9	20
10	24
11	31
12	20
13	29
14	21
15	28
16	28
17	29
18	31
19	23
20	34

of the observations as coming from assessment of a single randomly selected computer or as the average assessment of several randomly selected computers each week. Note that twenty observations are inadequate for a normal time-series analysis.

18.4.1 Identification of ARIMA (p, d, q) Models

The ARIMA (auto-regressive, integrated, moving average) model of a time series is defined by three terms (p, d, q). Identification of a time series is the process of finding integer, usually very small (e.g., 0, 1, or 2), values of p, d, and q that model the patterns in the data. When the value is 0, the element is not needed in the model. The middle element, d, is investigated before p and q. The goal is to determine if the process is stationary and, if not, to make it stationary before determining the values of p and q. Recall that a stationary process has a constant mean and variance over the time period of the study.

18.4.1.1 Trend Components, d*: Making the Process Stationary*

The first step in the analysis is to plot the sequence of scores over weeks as seen in Figure 18.1, produced in the SPSS Windows menus by selecting Graphs, and then Sequence. The two relevant features of the plot are central tendency and dispersion. Is the mean apparently shifting over the time period? Is the dispersion increasing or decreasing over the time period?

```
TSPLOT VARIABLES= quality
 /ID= week
 /NOLOG
 /FORMAT NOFILL NOREFERENCE.
```

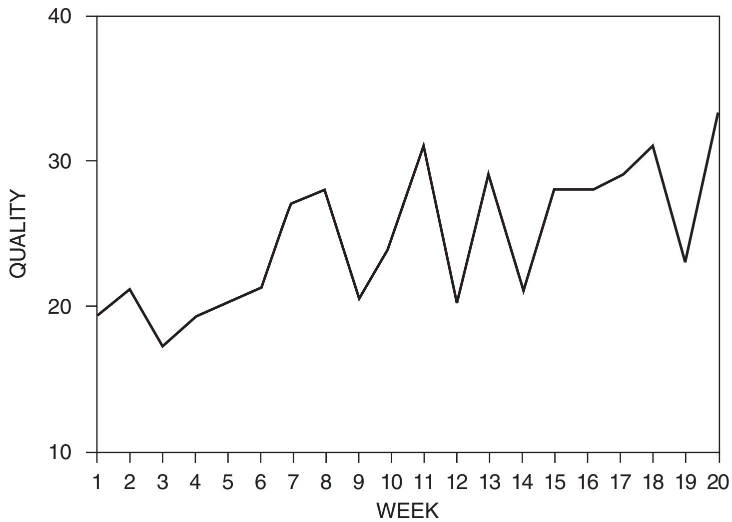

FIGURE 18.1 Plot of quality by week for the data of Table 18.2. SPSS TSPLOT syntax and output.

There are possible shifts in both the mean and the dispersion over time for this series. The mean may be edging upwards, and the variability may be increasing. If the mean is changing, the trend is removed by differencing once or twice. If the variability is changing, the process may be made stationary by logarithmic transformation.

Differencing the scores is the easiest way to make a nonstationary mean stationary (flat). The number of times you have to difference the scores to make the process stationary determines the value of d. If $d = 0$, the model is already stationary and has no trend. When the series is differenced once, $d = 1$ and linear trend is removed. When the difference is then differenced, $d = 2$ and both linear and quadratic trend are removed. For nonstationary series, d values of 1 or 2 are usually adequate to make the mean stationary.

Differencing simply means subtracting the value of an earlier observation from the value of a later observation. The first value in Table 18.2, for example, is 19, and the second value is 21. Thus the lag 1 difference is 2.[1] Table 18.3 shows all of the lag 1 first differences in the column labeled qual_1. These differences represent the *changes* in quality over the 20-week period. Notice that there are only 19 lag 1 first differences for this 20-week series. Table 18.3 also contains the second difference, qual_2; this is the difference of the first difference with both linear and quadratic trend removed, if present. There are 18 lag 1 second differences in this example.

[1]Lags of other lengths may be required when adjusting for seasonality, as discussed in Section 18.5.1.

TABLE 18.3 Lag 1 and Lag 2 Differences for the Data of Table 18.2

week	quality	qual_1	qual_2
1	19	.	.
2	21	21 − 19 = 2	.
3	17	17 − 21 = −4	−4 − 2 = −6
4	19	19 − 17 = 2	2 − (−4) = 6
5	20	20 − 19 = 1	1 − 2 = −1
6	21	21 − 20 = 1	1 − 1 = 0
7	27	27 − 21 = 6	6 − 1 = 5
8	28	28 − 27 = 1	1 − 6 = −5
9	20	20 − 28 = −8	−8 − 1 = −9
10	24	24 − 20 = 4	4 − (−8) = 12
11	31	31 − 24 = 7	7 − 4 = 3
12	20	20 − 31 = −11	−11 − 7 = −18
13	29	29 − 18 = 11	11 − (−13) = 24
14	21	21 − 29 = −8	−8 − 11 = −19
15	28	28 − 21 = 7	7 − (−8) = 15
16	28	28 − 28 = 0	0 − 7 = −7
17	29	29 − 28 = 1	1 − 0 = 1
18	31	31 − 29 = 2	2 − 1 = 1
19	23	23 − 31 = −8	−8 − 2 = −10
20	34	34 − 23 = 11	11 − (−8) = 19
		mean = .79	

In the simplest time series, an observation at a time period simply reflects the random shock at that time period, a_t, that is:

$$Y_t = a_t \tag{18.1}$$

The random shocks are independent with constant mean and variance, and so are the observations. If there is trend in the data, however, the score also reflects that trend as represented by the slope of the process. In this slightly more complex model, the observation at the current time, Y_t, depends on the value of the previous observation, Y_{t-1}, the slope, and the random shock at the current time period:

$$Y_t = \theta_0 (Y_{t-1}) + a_t \tag{18.2}$$

That mean of the first difference, qual_1, is the slope of the linear trend in the series, in this example .79. For the example:

$$Y_t = .79(Y_{t-1}) + a_t$$

To see if the process is stationary after linear trend is removed, the first difference scores at lag 1 are plotted against weeks, as seen in Figure 18.2. If the process is now stationary, the line will be basically horizontal with constant variance.

```
CREATE
QUAL_1 = SDIFF(QUALITY).
TSPLOT VARIABLES= QUAL_1
 /ID= week
 /NOLOG
 /FORMAT NOFILL NOREFERENCE.
```

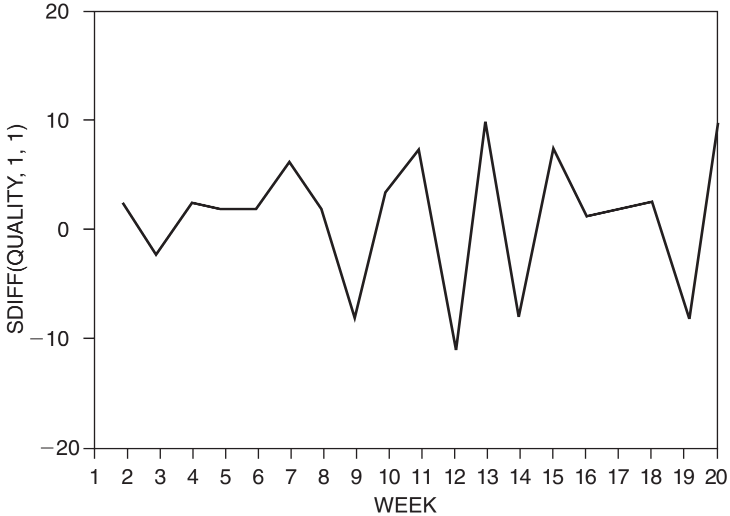

FIGURE 18.2 Plot of lag 1 first differences against week for the data of Table 18.3. SPSS TSPLOT syntax and output.

The series now appears stationary with respect to central tendency, so second differencing does not appear necessary. However, the variability seems to be increasing over time. Transformation is considered for series in which variance changes over time and differencing does not stabilize the variance (cf. Box, Jenkins, & Reinsel, 1994; McCleary & Hay, 1980). The logarithmic transformation is appropriate. Because of the zero and negative values in qual_1, with −11 the largest negative value, 12 is added before computing the log (Table 18.4).

The transformed difference is plotted to see if both mean and variance are now stabilized, as seen in Figure 18.3.

Although the scale has changed, the transformed difference does not appear to have less variability than the untransformed difference. There is also the usual problem of increased difficulty of interpretation of transformed variables. Therefore, the untransformed difference is used in future analyses.

18.4.1.2 Auto-Regressive Components

The auto-regressive components represent the memory of the process for preceding observations. The value of p is the number of auto-regressive components in an ARIMA (p, d, q) model. The value of p is 0 if there is no relationship between adjacent observations. When the value of p is 1, there is a relationship between observations at lag 1 and the correlation coefficient ϕ_1 is the magnitude of the

TABLE 18.4 Log Transformation of Lag 1 First Difference of Table 18.3

week	quality	qual_1	log(qual_1 + 12)
1	19	.	.
2	21	2	1.15
3	17	>4	0.9
4	19	2	1.15
5	20	1	1.11
6	21	1	1.11
7	27	6	1.26
8	28	1	1.11
9	20	>8	0.6
10	24	4	1.2
11	31	7	1.28
12	20	>11	0
13	29	11	1.32
14	21	>8	0.6
15	28	7	1.28
16	28	0	1.08
17	29	1	1.11
18	31	2	1.15
19	23	>8	0.6
20	34	11	1.36

relationship. When the value of p is 2, there is a relationship between observations at lag 2 and the correlation coefficient ϕ_2 is the magnitude of the relationship. Thus p is the number of correlations you need to model the relationship.

For example, a model with $p = 2$, ARIMA (2, 0, 0), is

$$Y_t = \phi_1 Y_{t-1} + \phi_2 Y_{t-2} + a_t \tag{18.3}$$

18.4.1.3 Moving Average Components

The moving average components represent the memory of the process for preceding random shocks. The value q indicates the number of moving average components in an ARIMA (p, d, q). When q is zero, there are no moving average components. When q is 1, there is a relationship between the current score and the random shock at lag 1 and the correlation coefficient θ_1 represents the magnitude of the relationship. When q is 2, there is a relationship between the current score and the random shock at lag 2, and the correlation coefficient θ_2 represents the magnitude of the relationship.

Thus, an ARIMA (0, 0, 2) model is

$$Y_t = a_t - \theta_1 a_{t-1} - \theta_2 a_{t-2} \tag{18.4}$$

```
COMPUTE LGQUAL_1 = LG10(QUAL_1 + 12).
EXECUTE.
TSPLOT VARIABLES= LGQUAL_1
 /ID= WEEK
 /NOLOG
 /FORMAT NOFILL NOREFERENCE.
```

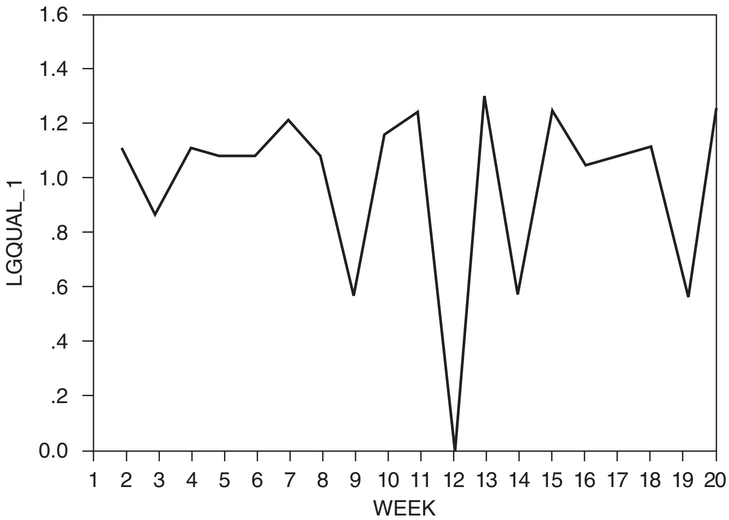

FIGURE 18.3 Plot of log of the lag 1 difference scores of Table 18.4. SPSS TSPLOT syntax and output.

18.4.1.4 Mixed Models

Somewhat rarely, a series has both auto-regressive and moving average components so both types of correlations are required to model the patterns. If both elements are present only at lag 1, the equation is:

$$Y_t = \phi_1 Y_{t-1} - \theta_1 a_{t-1} + a_t \tag{18.5}$$

18.4.1.5 ACFs and PACFs

Models are identified through patterns in their ACFs (autocorrelation functions) and PACFs (partial autocorrelation functions). Both autocorrelations and partial autocorrelations are computed for sequential lags in the series. The first lag has an autocorrelation between Y_{t-1} and Y_t, the second lag has both an autocorrelation and partial autocorrelation between Y_{t-2} and Y_t, and so on. ACFs and PACFs are the functions across all the lags.

The equation for autocorrelation is similar to bivariate r (Equation 3.29) except that the overall mean $\overline{Y}$ is subtracted from each Y_t and from each Y_{t-k}, and the denominator is the variance of the whole series.

$$r_k = \frac{\frac{1}{N-k}\sum_{t=1}^{N-k}(Y_t - \overline{Y})(Y_{t-k} - \overline{Y})}{\frac{1}{N-1}\sum_{t=1}^{N}(Y_t - \overline{Y})^2} \tag{18.6}$$

where N is the number of observations in the whole series, k is the lag. $\overline{Y}$ is the mean of the whole series and the denominator is the variance of the whole series.

The standard error of an autocorrelation is based on the squared autocorrelations from all previous lags. At lag 1, there are no previous autocorrelations, so r_0^2 is set to 0.

$$\mathrm{SE}_{r_k} = \sqrt{\frac{1 + 2\sum_{l=0}^{k-1} r_l^2}{N}} \tag{18.7}$$

The equations for computing partial autocorrelations are much more complex, and involve a recursive technique (cf. Dixon, 1992, p. 619). However, the standard error for a partial autocorrelation is simple and the same at all lags:

$$\mathrm{SE}_{pr} = \frac{1}{\sqrt{N}} \tag{18.8}$$

The series used to compute the autocorrelation is the series that is to be analyzed. Because differencing is used with the small-sample example to remove linear trend, it is the differenced scores that are analyzed. For the small-sample example of Table 18.2 the mean of the differenced series is 0.79, that is $\overline{Y} = 0.79$. The autocorrelation at lag 1 is computed between pairs of differenced scores at Y_t and Y_{t-1}, as seen in Table 18.5. Notice that there are only 18 such pairs, one pair lost to differencing and a second to lag. Also notice that the scores have been inverted so that Week 20 is first.

For the first lag, the autocorrelation is found by applying Equation 18.6 to the differenced data, recalling that in the original series $N = 20$:

$$r_1 = \frac{\frac{1}{19}[(11 - .79)(-8 - .79) + (-8 - .79)(2 - .79) + \cdots + (-4 - .79)(2 - .79)]}{\frac{1}{19}[(11 - .79)^2 + (-8 - .79)^2 + \cdots + (2 - .79)^2]} = -.61$$

Using Equation 18.7, the standard error of the autocorrelation for the first lag is

$$\mathrm{SE}_{r_1} = \sqrt{\frac{1 + 2(0)}{20}} = 0.22$$

Autocorrelations and standard errors for other lags are calculated by the same procedure.

TABLE 18.5 Pairs of Differenced Scores for Lag 1 Autocorrelation from Data of Table 18.3

Week	Differenced Score at t	Differenced Score at $t - 1$
20	11	−8
19	−8	2
18	2	1
17	1	0
16	0	7
15	7	−8
14	−8	9
13	9	−11
12	−11	7
11	7	4
10	4	−8
9	−8	1
8	1	6
7	6	1
6	1	1
5	1	2
4	2	−4
3	−4	2

Using Equation 18.8, the standard error of all the partial correlations is

$$\mathrm{SE}_{pr} = \frac{1}{\sqrt{20}} = 0.22$$

Calculation of the partial autocorrelations after the first few is labor intensive. However, McCleary and Hay (1980) provide equations showing the following relationships between ACF and PACF for the first three lags.

$$\mathrm{PACF}(1) = \mathrm{ACF}(1) \tag{18.9}$$

$$\mathrm{PACF}(2) = \frac{\mathrm{ACF}(2) - (\mathrm{ACF}(1))^2}{1 - [\mathrm{ACF}(1)]^2} \tag{18.10}$$

$$\mathrm{PACF}(3) = \frac{-2(\mathrm{ACF}(1))\mathrm{ACF}(2) - [\mathrm{ACF}(1)]^2\mathrm{ACF}(3)}{1 + 2[\mathrm{ACF}(1)]^2\mathrm{ACF}(2) - [\mathrm{ACF}(2)]^2 - 2[\mathrm{ACF}(1)]^2} \tag{18.11}$$

If an autocorrelation at some lag is significantly different from zero, the correlation is included in the ARIMA model. Similarly, if a partial autocorrelation at some lag is significantly different from zero, it, too, is included in the ARIMA model. The significance of full and partial autocorrelations is assessed using their standard errors.

Although you can look at the autocorrelations and partial autocorrelations numerically, it is standard practice to plot them. The center vertical (or horizontal) line for these plots represents full or partial autocorrelations of zero; then symbols such as * or _ are used to represent the size and direction of the autocorrelation and partial autocorrelation at each lag. You compare these obtained plots with standard, and somewhat idealized, patterns that are shown by various ARIMA models, as discussed more completely in Section 18.6.1.

The ACF and PACF for the first 10 lags of the differenced scores of Table 18.3 are seen in Figure 18.4, as produced by SPSS ACF.

The boundary lines around the functions are the 95% confidence bounds. The pattern here is a large, negative autocorrelation at lag 1 and a decaying PACF, suggestive of an ARIMA (0, 0, 1) model, as illustrated in Section 18.6.1. Recall, however, that the series has been differenced, so the ARIMA model is actually (0, 1, 1). The series apparently has both linear trend and memory for the preceding random shock. That is, the quality of the computers is generally increasing, however, the quality in one week is influenced by random events in the manufacturing process from both the current and preceding weeks. The q value of 1 indicates that, with a differenced series, only the first of the two correlations in Equation 18.4 needs to be estimated, the correlation coefficient θ_1:

$$Y_t = a_t - \theta_1 a_{t-1}$$

18.4.2 Estimating Model Parameters

Estimating the values of parameters in models consists of estimating the ϕ parameter(s) from an auto-regressive model or the θ parameter(s) from a moving average model. As indicated by McDowall et al. (1980) and others, the following rules apply:

1. Parameters must differ significantly from zero and all significant parameters must be included in the model.
2. Because they are correlations, all auto-regressive parameters, ϕ, must be between -1 and 1. If there are two such parameters ($p = 2$) they must also meet the following requirements:

$$\phi_1 + \phi_2 < 1 \text{ and}$$
$$\phi_2 - \phi_1 < 1$$

   These are called the *bounds of stationarity* for the auto-regressive parameter(s).

3. Because they are also correlations, all moving average parameters, θ, must be between -1 and 1. If there are two such parameters ($q = 2$) they must also meet the following requirements:

$$\theta_1 + \theta_2 < 1 \text{ and}$$
$$\theta_2 - \theta_1 < 1.$$

   These are called the *bounds of invertibility* for the moving average parameter(s).

```
ACF
/VARIABLES = QUAL_1
/NOLOG
/MXAUTO 10
/SERROR = MA
/PACF.
```

```
MODEL:  MOD_1.

Variable:  QUAL_1      Missing cases:  1     Valid cases:  19

Autocorrelations:  QUAL_1   SDIFF(QUALITY,1,1)

     Auto- Stand.
Lag  Corr.   Err. -1 -.75  -.5  -.25  0   .25  .5   .75   1  Box-Ljung  Prob.
                  ├────┼────┼────┼────┼────┼────┼────┼────┤
  1  -.615   .229        ***.********|         .                 8.377   .004
  2   .128   .304        .           |***         .              8.762   .013
  3   .047   .307        .           |*           .              8.817   .032
  4  -.175   .307        .        ***|            .              9.630   .047
  5   .259   .312        .           |*****       .             11.540   .042
  6  -.213   .323       .        ****|            .             12.934   .044
  7   .203   .331       .            |****        .             14.299   .046
  8  -.264   .337       .       *****|            .             16.830   .032
  9   .137   .348       .            |***         .             17.579   .040
 10   .070   .351       .            |*           .             17.799   .058

Plot Symbols:     Autocorrelations *     Two Standard Error Limits.
Standard errors are based on the Bartlett (MA) approximation.

Total cases:   20      Computable first lags:   18

Partial Autocorrelations:   QUAL_1   SDIFF(QUALITY,1,1)

   Pr-Aut- Stand.
Lag  Corr.   Err. -1 -.75  -.5  -.25  0   .25  .5   .75   1
                  ├────┼────┼────┼────┼────┼────┼────┼────┤
  1  -.615   .229        ***.********|         .
  2  -.402   .229           .********|         .
  3  -.172   .229           .     ***|         .
  4  -.329   .229           . *******|         .
  5  -.045   .229           .       *|         .
  6  -.091   .229           .      **|         .
  7   .171   .229           .        |***      .
  8  -.154   .229           .     ***|         .
  9  -.174   .229           .     ***|         .
 10  -.058   .229           .       *|         .

Plot Symbols:      Autocorrelations *    Two Standard Error Limits .

Total cases:  20     Computable first lags: 18
```

(continued)

FIGURE 18.4 ACF and PACF for differenced scores of Table 18.3. SPSS ACF syntax and output.

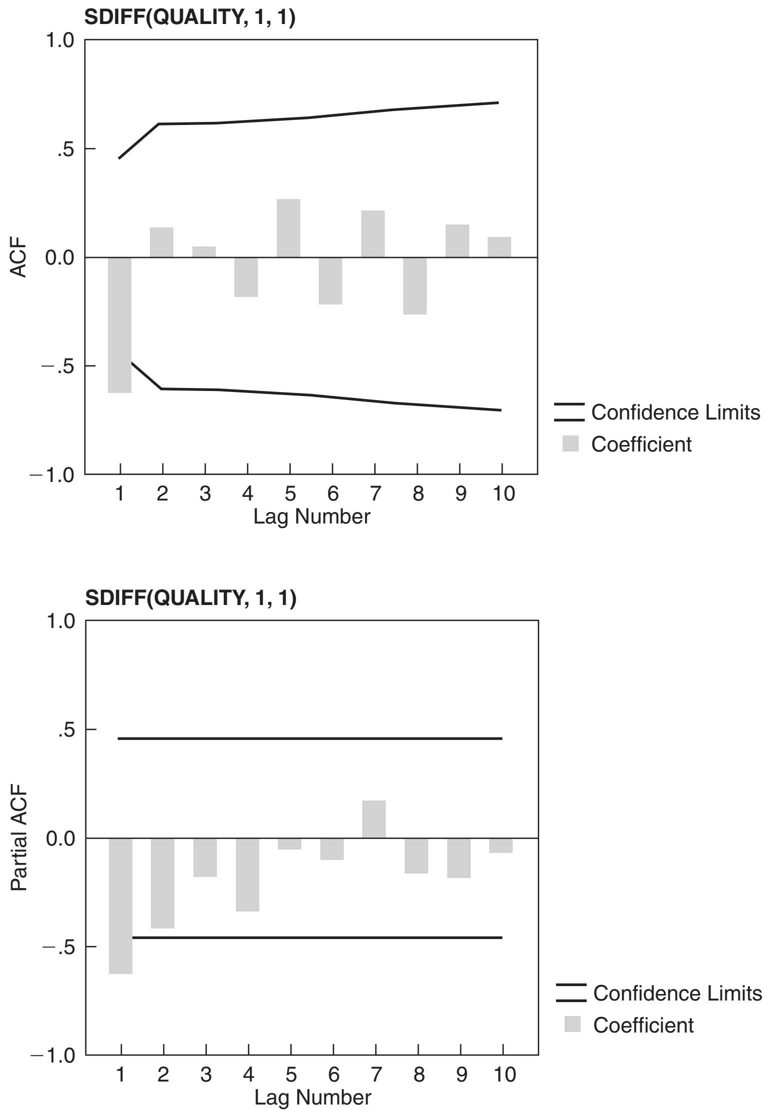

FIGURE 18.4 Continued

Complex and iterative maximum likelihood procedures are used to estimate these parameters. The reason for this is seen in an equation for θ_1:

$$\theta_1 = \frac{-\mathrm{cov}(a_t a_{t-1})}{\sigma_N^2} \tag{18.12}$$

The problem, of course, is that the sizes of the random shocks at various time periods are not known. Rather, this part of the process has to be estimated. For the ARIMA (0, 1, 1) model of the small-sample example data of Table 18.3, SAS estimates that $\theta_1 = .73$, and SPSS estimates it at .69. Either value is significant. The equation for the differenced series is, then, about

$$Y_t = a_t - .71a_{t-1}$$

18.4.3 Diagnosing a Model

How well does the model fit the data? Are the values of observations predicted from the model close to actual ones?

If the model is good, the residuals (differences between actual and predicted values) of the model are a series of random errors. These residuals form a set of observations that are examined the same way as any time series.

ACFs and PACFs for the residuals of the model are examined. As recommended by Pankrantz (1983), if the residuals represent only random error, the absolute value of $t\ (= r_t/\text{SE}_r)$ for autocorrelations at each of the first three lags[2] should be less than 1.25 and for later lags less than 1.60. However, as McCain and McCleary (1979) point out, if there are many lags (say, 30 or more) one or two of the higher-order lags may exceed these criteria by chance even if the residuals are essentially random. An alternative method is available through SAS ARIMA, in which there is a check to see if the first several lags deviate from random error (called white noise).

SPSS ARIMA adds residuals into the data set in a column called `ERR_1` or `ERR#1`. ACFs and PACFs are requested for this added variable and examined for patterns. If the residuals represent only random error, there should be no sizeable full and partial autocorrelations remaining in the data. All of the autocorrelations should fall within their 95% confidence intervals, and the criteria proposed by Pankrantz should be met.

Figure 18.5 shows the ACFs and PACFs for the `ERR_1` column created after specification of the ARIMA (0, 1, 1) model (note that sometimes SPSS labels the column `ERR#1`). The standard method is specified for estimating standard errors (`SERROR=MA`), rather than the default SPSS method.

At lag 1, the absolute value of t for the first autocorrelation is 1.92 $(-.440/.229)$, larger than the value of 1.25 recommended by Pankrantz. The Box-Ljung statistic for this lag—another test that the ACF does not differ from zero—is also significant. All of the other full and partial autocorrelations in the first 10 lags are acceptable. To try to improve the fit, it is decided to try an ARIMA (1, 1, 0) model. The results, including the ACF and PACF for the residuals, appear in Table 18.6.

At this point, all of the autocorrelations are small and nonsignificant, as are the partial autocorrelations, except for lag 8, which is assumed to be random. The t value for the first autocorrelation is $(.322/.229 =)$ 1.41, larger than the 1.25 value recommended by Pankrantz, but not significant. The parameter for the model is $\phi_1 =$ `-.68624562` and is statistically significant, `APPROX. PROB. = .00205822`. Therefore, the decision is to use this ARIMA (1, 1, 0) model. With this model, there is both trend (the quality of the computers is generally improving over time) and auto-regression in the model (the quality at one week depends on the quality of computers produced in the preceding week).

18.4.4 Computer Analysis of Small-Sample Time-Series Example

This section and some of the previous tables and figures show runs for identification, estimation, and diagnosis through SPSS and SAS for the small-sample time-series data of Table 18.2. The diagnosis

[2]And seasonal lags.

```
ARIMA QUALITY
 /MODEL=(0 1 1 )CONSTANT.
ACF
/VARIABLES = err_1
/NOLOG
/MXAUTO 10
/SERROR = MA
/PACF.
```

```
MODEL:  MOD_3.

Variable:  ERR_1        Missing cases:  1     Valid cases:  19

Autocorrelations:  ERR_1    Error for QUALITY from ARIMA, MOD_8 CON

     Auto- Stand.
Lag  Corr.   Err. -1 -.75  -.5  -.25  0   .25  .5   .75   1  Box-Ljung  Prob.
                  |----|----|----|----|----|----|----|----|
  1  -.401   .229            .********|        .                3.568   .059
  2   .033   .264          .          |***       .              3.594   .166
  3   .006   .264          .          *          .              3.595   .309
  4  -.159   .264          .       ***|          .              4.269   .371
  5   .168   .269          .          |***       .              5.071   .407
  6  -.179   .274          .      ****|          .              6.053   .417
  7   .098   .281          .          |**        .              6.369   .497
  8  -.295   .282          .    ******|          .              9.533   .299
  9   .200   .298         .           |****      .             11.121   .268
 10   .100   .305         .           |**        .             11.567   .315

Plot Symbols:     Autocorrelations *     Two Standard Error Limits .
Standard errors are based on the Bartlett (MA) approximation.

Total cases:  20     Computable first lags:  18

Partial Autocorrelations:   ERR_1    Error for QUALITY from ARIMA, MOD_8 CON

   Pr-Aut- Stand.
Lag  Corr.   Err. -1 -.75  -.5  -.25  0   .25  .5   .75   1
                  |----|----|----|----|----|----|----|----|
  1  -.401   .229            .********|        .
  2  -.152   .229            .     ***|        .
  3  -.049   .229            .       *|        .
  4  -.209   .229            .    ****|        .
  5  -.018   .229            .        *        .
  6  -.151   .229            .     ***|        .
  7  -.040   .229            .       *|        .
  8  -.407   .229            .********|        .
  9  -.113   .229            .      **|        .
 10   .037   .229            .        |*       .

Plot Symbols:      Autocorrelations *    Two Standard Error Limits .

Total cases:  20     Computable first lags: 18
```

FIGURE 18.5 ACF and PACF for the residuals. SPSS ACF syntax and output.

TABLE 18.6 Estimation and Diagnosis of an ARIMA (1, 1, 0) Model for the Data of Table 18.2 (SPSS ARIMA and ACF Syntax and Selected Output)

```
ARIMA QUALITY
 /MODEL=( 1 1 0)NOCONSTANT
 /MXITER 10
 /PAREPS .001
 /SSQPCT .001
 /FORECAST EXACT.
ACF
/VARIABLES = ERR#1
/NOLOG
/MXAUTO 10
/SERROR=MA
/PACF.
```

```
MODEL: MOD_2

Model Description:

Variable: QUALITY
Regressors: NONE

Non-seasonal differencing: 1
No seasonal component in model.

FINAL PARAMETERS:

Number of residuals  19
Standard error        4.6624974
Log likelihood      -56.017981
AIC                 114.03596
SBC                 114.9804

           Analysis of Variance:

             DF  Adj. Sum of Squares   Residual Variance

Residuals    18            404.63345           21.738882

          Variables in the Model:

               B         SEB      T-RATIO   APPROX. PROB.

AR1   -.68624562   .19075163   -3.5975873      .00205822

The following new variables are being created:

  Name        Label

  FIT#1       Fit for QUALITY from ARIMA, MOD_2 NOCON
  ERR#1       Error for QUALITY from ARIMA, MOD_2 NOCON
  LCL#1       95% LCL for QUALITY from ARIMA, MOD_2 NOCON
  UCL#1       95% UCL for QUALITY from ARIMA, MOD_2 NOCON
  SEP#1       SE of fit for QUALITY from ARIMA, MOD_2 NOCON
```

(continued)

TABLE 18.6 Continued

```
MODEL:  MOD_3.

Variable:  ERR#1        Missing cases:   1     Valid cases:   19

Autocorrelations:   ERR#1   Error for QUALITY from ARIMA, MOD_2 NOCO

     Auto- Stand.
Lag  Corr.   Err. -1  -.75  -.5 -.25   0   .25  .5   .75   1    Box-Ljung  Prob.
                   |----|----|----|----|----|----|----|----|

  1  -.322   .229             .  ******|        .                  2.304   .129
  2  -.217   .252            .     ****|         .                 3.404   .182
  3   .040   .262            .         |*        .                 3.443   .328
  4  -.083   .262            .       **|         .                 3.624   .459
  5   .129   .263           .          |***       .                4.102   .535
  6  -.016   .267           .          *          .                4.109   .662
  7   .004   .267           .          *          .                4.110   .767
  8  -.286   .267           .    ******|          .                7.085   .527
  9   .258   .283           .          |*****     .                9.738   .372
 10   .154   .295          .           |***        .              10.787   .374

Plot Symbols:      Autocorrelations *     Two Standard Error Limits.
Standard errors are based on the Bartlett (MA) approximation.

Total cases:  20     Computable first lags:  18


Partial Autocorrelations:   ERR#1   Error for QUALITY from ARIMA, MOD_2 NOCO

   Pr-Aut- Stand.
Lag  Corr.   Err.   -1  -.75  -.5 -.25  0  .25   .5  .75    1
                   |----|----|----|----|----|----|----|----|

  1  -.322   .229             .  ******|        .
  2  -.358   .229             . *******|        .
  3  -.218   .229             .    ****|        .
  4  -.310   .229             .  ******|        .
  5  -.114   .229             .      **|        .
  6  -.134   .229             .     ***|        .
  7  -.041   .229             .       *|        .
  8  -.469   .229             *********|        .
  9  -.178   .229             .    ****|        .
 10  -.082   .229             .      **|        .

Plot Symbols:      Autocorrelations *     Two Standard Error Limits.

Total cases:  20     Computable first lags:  18
```

run assumes an ARIMA (1, 1, 0) model. Programs demonstrated are SAS ARIMA and SPSS ACF and ARIMA. Initial data are arranged in two columns, as seen in Table 18.2; as the analysis proceeds, additional columns may be added.

Using SPSS, the first step is examination of a plot of the series over time, as demonstrated in Figure 18.1 with SPSS syntax and output. If differencing is required, the sequence is reexamined, as seen in Figure 18.2. SPSS ACF is used for the identification phase, as demonstrated in Figure 18.4. SPSS ARIMA is used during the estimation phase, as seen in Table 18.6. Because the DV is differenced, the constant is excluded from the model using NOCONSTANT. Remaining instructions are produced by default in the SPSS Windows menu system.

The estimation for ϕ_1 is from the `Variables in the Model:` segment of output, labeled `B`. The value is `AR1 -.68624562` and the associated `T-RATIO` with its `APPROX. PROB.` of `-3.5975873` and `.00205822`, respectively, indicates that the correlation is significantly different from zero. AIC and SBC are discussed in Section 18.6.4. Residual variance is discussed in Section 18.6.2. Notice in the section labeled `The following new variables are being created:` that `ERR_1` is added to the list of variables in the data set. These are the residuals which are examined in the diagnostic phase using SPSS ACF again, as seen in Table 18.6 and Figure 18.5.

Table 18.7 contains the SAS GPLOT syntax for producing the initial plots of the series over time, with and without differencing. The variable `qual_1` is the first-differenced time series of Table 18.3.

Table 18.8 shows plots of the ACF and PACF generated by SAS ARIMA. First-differencing with lag = 1 is produced in SAS ARIMA by adding `(1)` to the DV in the variable instruction: `var=quality(1)`.

SAS ARIMA prints the zero-lag (perfect) correlation as the first row in the ACF. After that, autocorrelations, standard errors of autocorrelations, and partial autocorrelations match those of SPSS ACF. SAS ARIMA also prints the inverse autocorrelations. These are used for the same purposes as the PACF, but are better in some circumstances (e.g., with subset and seasonal autoregressive models). They may also indicate over-differencing, as noted in the on-disk documentation. The final table, `Autocorrelation Check for White Noise`, shows that, as a whole, the autocorrelations reliably differ from zero, `Prob > ChiSq = 0.0441.`

Table 18.9 shows the modeling run. The `estimate` instruction specifies a model with one auto-regressive component (`p=1`) and no constant (`noint`).

ACF and PACF plots are not shown here, but are the same as those produced by SPSS, with the addition of the inverse autocorrelation plot. The `Autocorrelation Check of White`

TABLE 18.7 SAS Syntax for the First Steps in Identification of a Time Series

Statistics Produced	Syntax

Figure 18.1	`proc gplot data=SASUSER. SSTIMSER;`
	`plot quality * week /`
	`run;`

Figure 18.2	`proc gplot data=SASUSER. SSTIMSER;`
	`plot qual_1 * week /`
	`run;`

TABLE 18.8 ACF and PACF for Time-Series Data of Table 18.2 (SAS ARIMA Syntax and Selected Output)

```
proc arima data=SASUSER.SSTIMSER;
  identify var=quality(1) nlag=10;
run;

                           Name of Variable = quality

                  Period(s) of Differencing                    1
                  Mean of Working Series                0.789474
                  Standard Deviation                    6.005076
                  Number of Observations                      19
                  Observation(s) eliminated by differencing    1

                                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      Std Error

   0     36.060942        1.00000    |                    |********************|              0
   1    -22.168538        -.61475    |        ************|        .           |       0.229416
   2      4.615833        0.12800    |        .           |***        .        |       0.303994
   3      1.696603        0.04705    |        .           |*          .        |       0.306818
   4     -6.305730        -.17486    |        .        ***|           .        |       0.307197
   5      9.334597        0.25886    |        .           |*****      .        |       0.312392
   6     -7.684356        -.21309    |       .        ****|            .       |       0.323485
   7      7.307771        0.20265    |       .            |****        .       |       0.330790
   8     -9.525587        -.26415    |       .       *****|            .       |       0.337261
   9      4.940225        0.13700    |      .             |***          .      |       0.347980
  10      2.541770        0.07049    |      .             |*             .     |       0.350807

                          "." marks two standard errors
```

TABLE 18.8 Continued

```
                          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        0.79213    |           .        |****************    |
   2        0.55365    |           .        |***********         |
   3        0.37607    |           .        |********.           |
   4        0.27422    |           .        |*****   .           |
   5        0.18409    |           .        |****    .           |
   6        0.17765    |           .        |****    .           |
   7        0.17911    |           .        |****    .           |
   8        0.15787    |           .        |***     .           |
   9        0.05555    |           .        |*       .           |

                          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       -0.61475    |        ************|        .           |
   2       -0.40175    |           .********|        .           |
   3       -0.17180    |           .     ***|        .           |
   4       -0.32851    |           . *******|        .           |
   5       -0.04478    |           .       *|        .           |
   6       -0.09101    |           .      **|        .           |
   7        0.17122    |           .        |***     .           |
   8       -0.15420    |           .     ***|        .           |
   9       -0.17380    |           .     ***|        .           |
  10       -0.05815    |           .       *|        .           |

                    Autocorrelation Check for White Noise

  To      Chi-           Pr >
 Lag    Square    DF    ChiSq    --------------- Autocorrelations ---------------

   6     12.93     6   0.0441    -0.615    0.128    0.047   -0.175    0.259   -0.213
```

TABLE 18.9 ARIMA (1, 1, 0) Model for Time-Series Data of Table 18.2 with Lag=1 First-Differencing (SAS ARIMA Syntax and Selected Output)

```
proc arima data=SASUSER.SSTIMSER;
  identify var=quality(1) nlag=10;
  estimate p=1 noint;
run;
```

```
                         The ARIMA Procedure

              Conditional Least Squares Estimation

                              Approx Std
     Parameter      Estimate       Error     t Value    Pr > |t|    Lag

     AR1,1          -0.72743     0.19450       -3.74      0.0015      1

                  Variance Estimate      21.78926
                  Std Error Estimate     4.667896
                  AIC                    113.4393
                  SBC                    114.3837
                  Number of Residuals          19
              * AIC and SBC do not include log determinant.

                  Autocorrelation Check of Residuals

 To        Chi-             Pr >
Lag      Square      DF     ChiSq    --------------------Autocorrelations--------------------

  6        3.04       5    0.6937    -0.261    -0.176     0.088    -0.030     0.143     0.015
 12       10.20      11    0.5123    -0.008    -0.245     0.294     0.148    -0.091    -0.050
 18       11.91      17    0.8053     0.048    -0.005     0.035     0.044    -0.067     0.026
```

`Noise` is useful for model diagnosis, to be discussed in Section 18.4.3. The parameter estimate is −0.72743 with an approximate standard error of 0.19450. This is followed by information on residuals similar to that of SPSS ARIMA.

SAS ARIMA also provides an `Autocorrelation Check of Residuals`, which assesses the departure of model residuals from random error. A significant result suggests discrepancies between the model and the data.

18.5 Types of Time-Series Analyses

There are two major varieties of time-series analysis: time domain (including Box-Jenkins ARIMA analysis) and spectral domain (including Fourier—Spectral—analysis). Time domain analyses deal directly with the DV over time; spectral domain analyses decompose a time series into its sine wave components. Either time or spectral domain analyses can be used for identification, estimation, and diagnosis of a time series. However, current statistical software offers no assistance for intervention analysis using spectral methods. As a result, this chapter is limited to time domain analyses. Numerous complexities are available with these analyses, however: seasonal autocorrelation and one or more interventions (and with different effects), to name a few.

18.5.1 Models with Seasonal Components

Seasonal autocorrelation is distinguished from "local" autocorrelation in that it is predictably spaced in time. Observations gathered monthly, for example, are often expected to have a spike at lag 12 because many behaviors vary consistently from month to month over the year. Similarly, observations made daily could easily have a spike at lag 7, and observations gathered hourly often have a spike at lag 24.[3] These seasonal cycles can often be postulated a priori, while local cycles are inferred from the data.

Like local cycles, seasonal cycles show up in plots of ACFs and PACFs as spikes. However, they show up at the appropriate lag for the cycle. Like local cycles, these autocorrelations can also be auto-regressive or moving average (or both). And, like local autocorrelation, a seasonal auto-regressive component tends to produce a decaying ACF function and spikes on the PACF while a moving average component tends to produce the reverse pattern.

When there are both local and seasonal trends, *d*, multiple differencing is used. For weekly measurement, as in the computer quality example, there could be a local linear trend at lag 1 and also a seasonal linear trend at lag 4 (monthly).

Seasonal models may be either *additive* or *multiplicative.* The additive seasonal model just described would be (0, 2, 0), with differencing at lags 1 and 4. The notation for a multiplicative seasonal model is ARIMA $(p, d, q)(P, D, Q)_s$, where s is the seasonal cycle. Thus, the notation for a seasonal model with a local trend at lag 1 and a seasonal trend component at lag 4 is $(0, 1, 0)(0, 1, 0)_4$. In this model, the interaction between time and seasons also is of interest. For example, the seasonal trend may be stronger at lower (or higher) levels of the series. The additive model is more parsimonious than the multiplicative model, and it is often found that the multiplicative component is very small (McCain & McCleary, 1979). Therefore, an additive model is recommended unless the multiplicative model is required to produce acceptable ACF and PACF, as well as significant parameter estimates.

[3]Notice that the expected lag is one less than the cycle when the first observation is "used up" in differencing.

TABLE 18.10 Data with Both Trend and Seasonality

Week	Quality
1	26
2	21
3	17
4	19
5	28
6	21
7	27
8	28
9	29
10	24
11	31
12	20
13	39
14	21
15	28
16	28
17	40
18	31
19	23
20	34

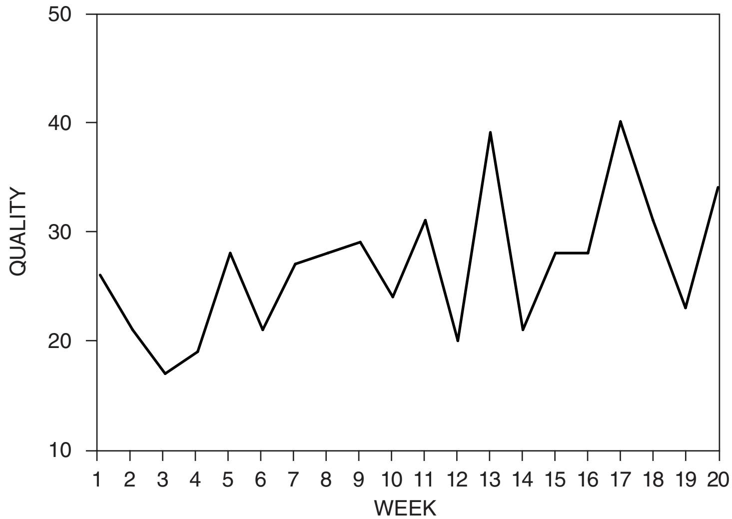

As an example of a seasonal model, reconsider the computer quality data, but this time with a linear trend added every four weeks. Perhaps monthly pep talks are given by the manager, or morale is improved by monthly picnics, or greater effort is made after receiving a monthly paycheck, or whatever. Data appear in Table 18.10 and are plotted using SPSS TSPLOT.

In the plot of quality, the astute observer might notice peaks at the first, fifth, ninth, thirteenth, and seventeenth weeks, indicative of a lag 3 autocorrelation after differencing. However, the pattern is much more evident from the ACF and PACF plots of differenced scores as seen in Figure 18.6, created by SPSS ACF.

```
CREATE
QUAL_1 = SDIFF(QUALITY 1 1).
ACF
/VARIABLES = QUAL_1
/NOLOG
/MXAUTO 10
/SERROR=MA
/PACF.
```

FIGURE 18.6 ACF and PACF for the differenced scores of data with both seasonality and trend from Table 18.12. SPSS ACF syntax and output.

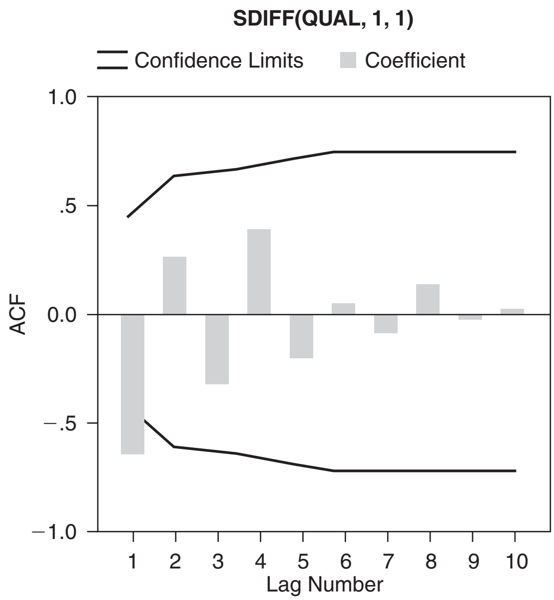

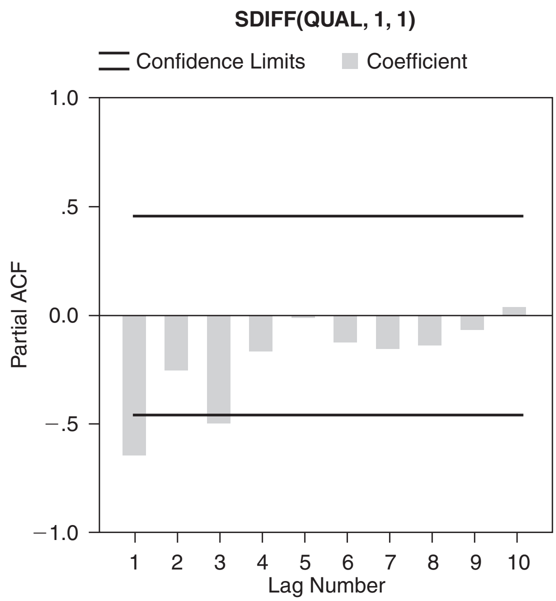

FIGURE 18.6 Continued

The large partial autocorrelation at lag 3 indicates the presence of a monthly trend which has not yet been removed by differencing. SPSS syntax for this differencing is seen in Figure 18.7, together with the ACF and PACF for the results. The large ACF at lag 1 remains in the data to be modeled with a moving average component, an auto-regressive component, or both.

The large PACF at lag 3 is now removed. The remaining autocorrelation can be modeled using ARIMA (1, 0, 0). The final model is ARIMA (1, 1, 0)$(0, 1, 0)_4$. Section 18.7.4 shows a seasonal model with intervention through SAS ARIMA.

18.5.2 Models with Interventions

Intervention, or interrupted, time-series analyses compare observations before and after some identifiable event. In quasi-experiments, the intervention is an attempt at experimental manipulation; however, the techniques are applicable to analysis of any event that occurs during the time series. The goal is to evaluate the impact of the intervention.

Interventions differ in both the onset (abrupt or gradual) and duration (permanent or temporary) of the effects they produce. An impulse intervention occurs once, briefly, and often produces an effect with abrupt onset and temporary duration. An earthquake or a riot is likely to have such an effect. Sometimes discrete interventions occur several times (e.g., assassinations). Quasi-experiments, on the other hand, are more likely to have an intervention that continues over considerable time; the effects produced are often also permanent or of long duration, but the onset can be either gradual or abrupt. For example, the profit-sharing plan applied to the small-sample computer quality example is likely to have a gradual-onset, long-duration effect. However, some impulse interventions also have an abrupt-onset, long-term or permanent effect; an example is bypass surgery.

The connection between an intervention and its effects is called a *transfer function.* In time-series jargon, an impulse intervention is also called an impulse or a pulse indicator and the effect is called an impulse or a pulse function. An effect with abrupt onset and permanent or long duration is called a step function. Because there are two levels of duration (permanent and temporary) and two levels of onset (abrupt and gradual), there are four possible combinations of effects, but the gradual-onset, short-term effect occurs rarely and requires curve fitting (Dixon, 1992). Therefore, it is not covered in this chapter.

Intervention analysis requires a column for the indicator variable that flags the occurrence of the event. With an impulse indicator, a code of 1 is applied in the indicator column at the single specific time period of the intervention and a code of 0 to the remaining time periods. When the intervention is of long duration or the effect is expected to persist, the column for the indicator contains 0s during the baseline time period and 1s at the time of and after the intervention.

Intervention analysis begins with identification, estimation, and diagnosis of the observations before intervention. The model, including the indicator variable, is then re-estimated for the entire series and the results diagnosed. The effect of the intervention is assessed by interpreting the coefficients for the indicator variable.

Transfer functions have two possible parameters. The first parameter, ω, is the magnitude of the asymptotic change in level after intervention. The second parameter, δ, reflects the onset of the change, the rate at which the post-intervention series approaches its asymptotic level. The ultimate change in level is

$$\text{level change} = \frac{\omega}{1 - \delta} \tag{18.13}$$

```
CREATE
QUAL_2 = SDIFF(QUAL_1 1 3).
ACF
/VARIABLES = QUAL_2
/NOLOG
/MXAUTO 10
/SERROR=MA
/PACF.
```

```
MODEL:  MOD_2.

Variable:  QUAL_2      Missing cases:  4     Valid cases:  16


Autocorrelations:   QUAL_2   SDIFF(QUAL_1,1,3)

     Auto- Stand.
Lag  Corr.   Err. -1 -.75  -.5  -.25  0   .25  .5   .75   1  Box-Ljung  Prob.
                   ├────┼────┼────┼────┼────┼────┼────┼────┤
  1  -.790   .250     ******.*********|         .                11.993   .001
  2   .626   .375       .             |************* .           20.049   .000
  3   .619   .435     .    ************|                .        28.546   .000
  4   .549   .487   .                 |***********         .     35.787   .000
  5  -.396   .525             ********|                          39.884   .000
  6   .259   .543                     |*****                     41.821   .000
  7  -.264   .551                *****|                          44.047   .000
  8   .270   .558                     |*****                     46.674   .000
  9  -.204   .567                 ****|                          48.381   .000
 10   .134   .571                     |***                       49.243   .000

Plot Symbols:     Autocorrelations *     Two Standard Error Limits .
Standard errors are based on the Bartlett (MA) approximation.

Total cases:  20     Computable first lags:  15

Partial Autocorrelations:   QUAL_2   SDIFF(QUAL_1,1,3)

   Pr-Aut- Stand.
Lag  Corr.   Err. -1 -.75  -.5  -.25  0   .25  .5   .75   1
                   ├────┼────┼────┼────┼────┼────┼────┼────┤
  1  -.790   .250     ******.*********|         .
  2   .003   .250           .         *         .
  3  -.330   .250           .  *******|         .
  4  -.114   .250           .       **|         .
  5   .160   .250           .         |***      .
  6  -.173   .250           .      ***|         .
  7  -.218   .250           .     ****|         .
  8   .079   .250           .         |**       .
  9  -.041   .250           .        *|         .
 10  -.134   .250           .      ***|         .

Plot Symbols:      Autocorrelations *    Two Standard Error Limits .

Total cases:  20     Computable first lags:  15
```

FIGURE 18.7 ACF and PACF for data of Table 18.10 after differencing both local and seasonal linear trends. SPSS ACF syntax and output.

Both ω and δ are tested for significance. If the null hypothesis that ω is 0 is retained, there is no impact of intervention. If ω is significant, the size of the change is ω. The δ coefficient varies between 0 and 1. When δ is 0 (nonsignificant), the onset of the impact is abrupt. When δ is significant but small, the series reaches its asymptotic level quickly. The closer δ is to 1, the more gradual the onset of change. SAS permits specification and evaluation of both ω and δ but SPSS permits evaluation only of ω.

18.5.2.1 Abrupt, Permanent Effects

Abrupt, permanent effects, called *step functions*, are those that are expected to show an immediate impact and continue over the long term. All post-intervention observations are coded 1 and ω is specified, but not δ (or it is specified as 0).

Table 18.11 contains hypothetical data and plot for an intervention analysis in which another 20 observations are added to the original 20 of Table 18.2. The first 20 observations are the same as Table 18.2 and are adequately modeled by ARIMA (1, 1, 0), as previously. The plot shows an apparent discontinuity at the 21st week when the profit-sharing plan was initiated. Notice that the trend that was evident in the first 20 weeks continues through the second 20.

SPSS ARIMA syntax and output for intervention analysis are in Table 18.12. The intervention is specified by adding WITH profit to the ARIMA paragraph. The auto-regressive model is then specified as MODEL=(1 1 0). The constant is not used because of the differencing. The assumption is, of course, that the same underlying patterns as in the first 20 weeks continue in weeks 21–40, with a possible intervention effect superimposed. The rest of the syntax is produced by default through the SPSS menu system.

The effect of the intervention is assessed by examining the `Variables in the Model` segment of output where the `T-RATIO` and `APPROX. PROB. for PROFIT` are `3.5722935` and `.00100388`, respectively, so that there is a statistically significant intervention effect. The magnitude (slope) of impact is indicated by the regression coefficient `(B)` for `PROFIT: 14.075030`. That is, quality increased by about 14 units as a result of the intervention.

The ACF and PACF for the residuals, `ERR_1`, as produced for this solution by SPSS are acceptable (except for the pesky autocorrelation and PACF at lag 8, an artifact of the way the data set was generated).

Table 18.13 shows the same analysis through SAS. The variable indicating the intervention is identified in two places: `crosscorr=PROFIT(1)` in the `identify` paragraph, and `input=PROFIT` in the `estimate` paragraph. The parenthetical 1 differences the IV.

The intervention variable is labeled `NUM1` (`profit` to the right in the table). Output is consistent with that of SPSS, with small differences in parameter estimates. The `Autocorrelation Check of Residuals` shows no problem with fit up until the 24th lag at $\alpha = .05$, but significant deviation from random autocorrelation for very late lags. Section 18.7.4 shows an abrupt, permanent intervention model with a seasonal component through SAS ARIMA.

18.5.2.2 Abrupt, Temporary Effects

Now suppose that the intervention has a strong, abrupt initial impact that dies out quickly. This is called a pulse effect. Effects like this are often seen with natural or man-made disasters. Or perhaps New Year's resolutions.

TABLE 18.11 Data Set for an Intervention with an Abrupt Onset, Long Duration Effect

Week	Profit	Quality	Week	Profit	Quality
1	0	19	21	1	41
2	0	21	22	1	45
3	0	17	23	1	40
4	0	19	24	1	42
5	0	20	25	1	45
6	0	21	26	1	46
7	0	27	27	1	53
8	0	28	28	1	52
9	0	20	29	1	46
10	0	24	30	1	51
11	0	31	31	1	59
12	0	20	32	1	47
13	0	29	33	1	59
14	0	21	34	1	49
15	0	28	35	1	57
16	0	28	36	1	57
17	0	29	37	1	58
18	0	31	38	1	62
19	0	23	39	1	53
20	0	34	40	1	65

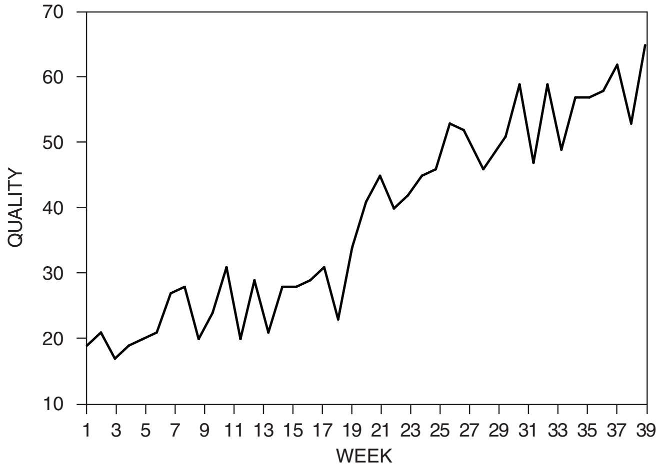

A hypothetical data set and plot for this kind of impact is in Table 18.14 in which the intervention is a pep talk given at the 21st week by the CEO.

The first 20 weeks are modeled as ARIMA (1, 1, 0), as previously, before the effect of the intervention is tested on the differenced variable qualit_1. SPSS syntax and output for the test of the

TABLE 18.12 Intervention Analysis for Abrupt, Permanent Effect (SPSS ARIMA Syntax and Selected Output)

```
* ARIMA.
TSET PRINT=DEFAULT CIN=95 NEWVAR=ALL.
PREDICT THRU END.
ARIMA QUALITY WITH PROFIT
 /MODEL=( 1 1 0) NOCONSTANT
 /MXITER 10
 /PAREPS .001
 /SSQPCT .001
 /FORECAST EXACT.
```

```
MODEL: MOD_1

Model Description:

Variable: QUALITY
Regressors: PROFIT

Non-seasonal differencing: 1
No seasonal component in model.


FINAL PARAMETERS:

Number of residuals  39
Standard error       4.8282746
Log likelihood       -116.07705
AIC                  236.1541
SBC                  239.48123

            Analysis of Variance:

             DF  Adj. Sum of Squares    Residual Variance

Residuals    37           878.09601            23.312235

           Variables in the Model:

                   B          SEB      T-RATIO   APPROX. PROB.

AR1        -.708299    .1232813   -5.7453872       .00000140
PROFIT    14.075030   3.9400542    3.5722935       .00100388
```

intervention appear in Table 18.15. Note that syntax is the same as for an abrupt, permanent effect; only the coding of the IV is different (compare coding in Tables 18.12 and 18.15 for profit sharing and pep talk, respectively).

The section of output labeled `Variables in the Model:` provides the test of the intervention, labeled `PEP_TALK` where the associated `T-RATIO` and `APPROX. PROB.` are `4.2983852` and `.00234735`, respectively. Thus, the pep talk has a significant, abrupt impact on the quality of the computers produced. The size of the impact is given by the `B` value of `14.043981`. That is, the computers improve by approximately 14 units of quality immediately after the pep talk by the CEO. The

TABLE 18.13 Intervention Analysis for Abrupt, Permanent Effect (SAS ARIMA Syntax and Selected Output)

```
proc arima data=SASUSER. INTERVN;
   identify var=QUALITY(1) nlag=7 crosscor=PROFIT(1);
   estimate p=1 noint input=PROFIT;
run;
```

Name of Variable = quality

Period(s) of Differencing	1
Mean of Working Series	1.179487
Standard Deviation	6.436561
Number of Observations	39
Observation(s) eliminated by differencing	1

Conditional Least Squares Estimation

Parameter	Estimate	Approx Std Error	t Value	Pr > |t|	Lag	Variable	Shift
AR1,1	-0.72716	0.12454	-5.84	<.0001	1	quality	0
NUM1	14.13524	3.93806	3.59	0.0010	0	profit	0

Variance Estimate	23.35194
Std Error Estimate	4.832384
AIC	235.5006
SBC	238.8277
Number of Residuals	39

* AIC and SBC do not include log determinant.

Autocorrelation Check of Residuals

To Lag	Chi-Square	DF	Pr > ChiSq	Autocorrelations					
6	5.24	5	0.3876	-0.230	-0.105	0.084	-0.024	0.216	0.003
12	17.10	11	0.1050	0.061	-0.243	0.189	0.310	0.042	-0.151
18	18.76	17	0.3423	0.085	-0.024	0.119	0.009	-0.019	-0.052
24	41.39	23	0.0107	-0.075	0.477	-0.157	-0.047	0.076	-0.025

TABLE 18.14 Data Set and Plot for an Intervention with an Abrupt Onset, Temporary Effect

Week	Quality	Pep Talk	Week	Quality	Pep Talk
1	19	0	21	45	1
2	21	0	22	45	0
3	17	0	23	36	0
4	19	0	24	34	0
5	20	0	25	34	0
6	21	0	26	35	0
7	27	0	27	42	0
8	28	0	28	41	0
9	20	0	29	35	0
10	24	0	30	40	0
11	31	0	31	48	0
12	20	0	32	36	0
13	29	0	33	48	0
14	21	0	34	38	0
15	28	0	35	46	0
16	28	0	36	46	0
17	29	0	37	47	0
18	31	0	38	51	0
19	23	0	39	42	0
20	34	0	40	54	0

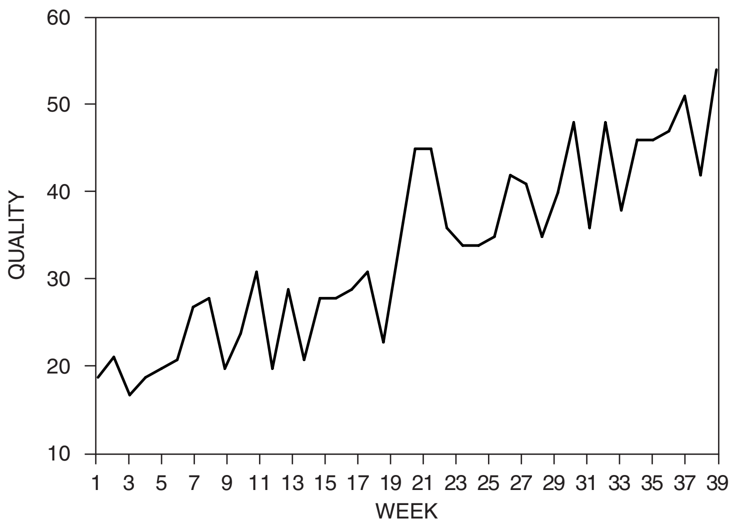

TABLE 18.15 Intervention Analysis for Abrupt, Temporary Effect (SPSS ARIMA Syntax and Output)

```
* ARIMA.
TSET PRINT=DEFAULT CIN=95 NEWVAR=ALL.
PREDICT THRU END.
ARIMA QUALITY WITH PEP_TALK
 /MODEL=( 1 1 0)NOCONSTANT
 /MXITER 10
 /PAREPS .001
 /SSQPCT .001
 /FORECAST EXACT.

MODEL: MOD_1

Model Description:

Variable:   QUALITY
Regressors: PEP_TALK

Non-seasonal differencing: 1
No seasonal component in model.

FINAL PARAMETERS:

Number of residuals  39
Standard error       5.3576408
Log likelihood       -120.07854
AIC                  244.15709
SBC                  247.48421

            Analysis of Variance:

              DF  Adj. Sum of Squares    Residual Variance
Residuals     37            1077.9915            28.704314

           Variables in the Model:

                     B         SEB      T-RATIO   APPROX. PROB.

AR1           -.663692    .1253407   -5.2951063       .00000573
PEP_TALK     14.043981   4.2983852    3.2672690       .00234735
```

ACF and PACF are acceptable for this model, except for the pesky partial autocorrelations at lags 3 and 4.

As for SPSS, syntax for an abrupt, temporary intervention in SAS is the same as for an abrupt, permanent intervention, but with a change in coding of the IV.

18.5.2.3 Gradual, Permanent Effects

Another possible effect of intervention is one that gradually reaches its peak effectiveness and then continues to be effective. Suppose, for example, the employees don't pay much attention to the profit-sharing plan when it is first introduced, but that it gradually becomes more and more popular. The data set is the same as for an abrupt, permanent effect (Table 18.11). However, now we are proposing that there is a linear growth in the effect of the intervention from the time it is implemented. SPSS cannot be used to model a gradual effect because there is no way to specify δ. (You might be able to approximate a gradual intervention effect in SPSS ARIMA by appropriately coding the IV, if you have a very clear notion of the nature of the expected trend, such as linear for the first 5 periods and then leveling off.)

Table 18.16 shows syntax and output for modeling a gradual, permanent intervention in SAS ARIMA, which indicates δ as a denominator function (cf. Equation 18.13). The denominator function, δ, is indicated by `( / (1) PROFIT)`, with the `/` indicating denominator and the `(1)` indicating that the predicted onset is linear.

The variables for ω and δ of Equation 18.13 are NUM1 and DEN1, indicating numerator and denominator functions of Equation 18.13, respectively. Note that the effect of the intervention remains statistically significant, but the test of gradualness of the impact is not. Thus, the step intervention model is supported.

A seasonal component may be added by supplementing the syntax with seasonal syntax. Table 18.17 shows SPSS and SAS syntax for adding monthly differencing to the intervention analysis for a gradual, permanent effect as shown in Tables 18.15 and 18.16.

18.5.2.4 Models with Multiple Interventions

Multiple interventions are modeled by specifying more than one IV. Several interventions can be included in models when using either SPSS ARIMA or SAS ARIMA. Table 18.18 shows syntax for the two programs for simultaneously modeling two interventions; say, an intervention with a step effect and an intervention with a pulse effect, called STEP and PULSE, respectively. The syntax assumes an ARIMA (1, 1, 0) model of the original variable, QUALITY.

18.5.3 Adding Continuous Variables

Continuous variables are added to time-series analyses to serve several purposes. One is to compare time series over populations or over different conditions. For example, time series for computer quality might be compared for two or more plants. Or water consumption might be compared over cities with and without water-conservation campaigns. Here, both continuous variables are considered DVs.

Another use is to add a continuous IV as a predictor. For example, a gas furnace might be observed while measuring input gas rate (the continuous IV) and output carbon dioxide (the DV). The goal here is to examine the relationship between the continuous IV and the DV; is output carbon dioxide predicted by input gas rate? If so, what is the prediction equation? Think of this as a bivariate regression analysis using the power of time-series analysis.

Finally, one or more continuous IVs might be added as covariates to an intervention analysis. For example, does a profit-sharing plan increase quality after adjusting for temperature in the manufacturing plant? All of these are examples of multivariate time-series analysis.

TABLE 18.16 **Intervention Analysis for Gradual, Permanent Effect (SAS ARIMA Syntax and Selected Output)**

```
proc arima data=SASUSER. INTERVN;
   identify var=QUALITY(1) nlag=7 crosscor=PROFIT(1);
   estimate p=1 noint input=( / (1) PROFIT);
run;
```

The ARIMA Procedure

Name of Variable = quality

Period(s) of Differencing	1
Mean of Working Series	1.179487
Standard Deviation	6.436561
Number of Observations	39
Observation(s) eliminated by differencing	1

Conditional Least Squares Estimation

Parameter	Estimate	Approx Std Error	t Value	Pr > |t|	Lag	Variable	Shift
AR1,1	-0.73398	0.12861	-5.71	<.0001	1	quality	0
NUM1	15.13381	4.79083	3.16	0.0033	0	profit	0
DEN1,1	-0.13199	0.33158	-0.40	0.6930	1	profit	0

Variance Estimate	24.69963
Std Error Estimate	4.969872
AIC	232.5722
SBC	237.485
Number of Residuals	38

* AIC and SBC do not include log determinant.

Autocorrelation Check of Residuals

To Lag	Chi-Square	DF	Pr > ChiSq	--------------------Autocorrelations--------------------					
6	4.79	5	0.4416	-0.229	-0.089	0.082	-0.037	0.205	0.001
12	15.36	11	0.1667	0.072	-0.230	0.189	0.288	0.061	-0.133
18	17.01	17	0.4539	0.061	-0.038	0.131	0.018	-0.013	-0.046
24	39.62	23	0.0170	-0.086	0.477	-0.147	-0.052	0.060	-0.045

TABLE 18.17 SAS ARIMA Syntax for a Gradual, Permanent Intervention Analysis with a Seasonal (Weekly) Component

```
proc arima data=SASUSER. INTERVN;
   identify var=QUALITY(1, 7) nlag=7 crosscor=PROFIT(1, 7) WEEK(7));
   estimate p=(1)(7) noint input=( / (1) PROFIT, WEEK );
run;
```

These analyses have their own assumptions and limitations. As usual, use of a covariate assumes that it is not affected by the IV; for example, it is assumed that temperature is not affected by introduction of the profit-sharing plan in the company under investigation. Also, for any multivariate time-series analysis, the CV (input) and DV (output) series must be "prewhitened"—reduced to random noise. That is, an ARIMA model is fit to the CV. The model is then applied to the DV before the two series are cross-correlated to determine whether both ARIMA models are the same. The SAS ETS on-disk manual shows how to do this in the section on transfer function. We recommend McCleary and Hay (1980) and/or Cromwell, Hannan, Labys, and Terraza (1994) if you plan to do multivariate time-series analyses, particularly if you are using SPSS.

Once the ARIMA models are reduced to random noise for both series, the time-series analysis is set up like the multiple intervention analysis of Table 18.18. That is, the CV enters the equation just as if it were any other IV. If both ARIMA models are the same, for example, temp_1 is added to the WITH list instead of pulse for SPSS ARIMA. For SAS ARIMA, `temp_1` is added both to the `crosscorr` list and to the `input` list instead of `pulse`.

TABLE 18.18 Syntax for Multiple Intervention Models

		Name of Parameter in Output	
Program	**Syntax**	*Step Intervention*	*Temporary Intervention*
SPSS ARIMA	ARIMA quality WITH step pulse /MODEL=(1 1 0)NOCONSTANT /MXITER 10 /PAREPS .001 /SSQPCT .001 /FORECAST EXACT.	STEP	PULSE

SAS ARIMA	`proc arima data=SASUSER. TIMESP;`
	`identify var=QUALITY(1) nlag=7`
	`crosscorr=(STEP PULSE);`
	`estimate p=1 noint input=(STEP PULSE);`
	`run;`
		NUM1	NUM2

18.6 Some Important Issues

18.6.1 Patterns of ACFs and PACFs

ARIMA models are identified by matching obtained patterns of ACF and PACF plots with idealized patterns. The best match often indicates which of the (*p, d, q*) parameters need to be included in the model, and at what size (0, 1, or 2). Sometimes, however, more than one pattern is suggested by the plots, or, like the small-sample example, the best pattern match does not reduce the residuals to random error. A first best guess is made on the basis of the ACF, PACF pattern and then, if the model fits the data poorly, another is tried out until the diagnostic phase is satisfactory.

Recall that ACF and PACF plots show deviations from zero autocorrelation. Table 18.19 shows idealized ACF and PACF patterns for the first 10 lags for many of the more common ARIMA models. In the plots of Table 18.19, absolute (positive) values of all autocorrelations (except those of the last model) are shown. The same type of model might produce negative autocorrelations as well, or a mix of positive and negative autocorrelations.

A "classic" auto-regressive model—ARIMA (p, 0, 0)—has an ACF that slowly approaches 0 and a PACF that spikes at lag p. Thus, if there is a spike only at lag 1 of the PACF (partial autocorrelation between Y_t and Y_{t-1}) and the ACF slowly declines, the model is probably ARIMA (1, 0, 0). If there is also a spike at lag 2 of the PACF (an autocorrelation between Y_t and Y_{t-2} with Y_{t-1} partialed out) and the ACF slowly declines, the model is likely to be ARIMA (2, 0, 0). Thus, the number of spikes on the PACF indicates the value of p for auto-regressive models.

A "classic" moving average model—ARIMA (0, 0, q)—has an ACF that spikes on the first q lags and a PACF that declines slowly. If there is a spike at lag 1 of the ACF and the PACF declines, the model may be ARIMA (0, 0, 1). If there also is a spike at lag 2 of the ACF and the PACF declines, the model is probably ARIMA (0, 0, 2). Thus, the number of spikes on the ACF indicate the value of q.

A mixed auto-regressive, moving average model has a slowly declining ACF, from the auto-regressive portion of the model (p), and also a slowly declining PACF, from the moving average portion of the model (q). The sizes of the auto-regressive and moving average parameters are not evident from the ACF and PACF. Therefore, it is prudent to start with $p = 1$ and $q = 1$ and increase these values only if residual ACF and PACF show spikes in the diagnostic phase. If the parameters are set too high, it may not show up in diagnosis, except perhaps as nonsignificant parameter estimates.

A nonstationary model—ARIMA (0, 1, 0)—has a single PACF spike, but there are two common ACF patterns. There may be constant spikes [Table 18.19 (a)], or there may be a "damped sine wave" [Table 18.19 (b)], in which spikes oscillate first on one side and then on the other side of zero autocorrelation.

If either an autocorrelation or a partial autocorrelation between observations k lags apart is statistically significant, the autocorrelation is included in the ARIMA model. The significance of an autocorrelation is evaluated from the 95% confidence intervals printed alongside it, or from the t distribution where the autocorrelation is divided by its standard error, or from the Box-Ljung statistic printed as standard SPSS output. Alternatively, SAS ARIMA provides tests of sets of lags. For example, Table 18.8 shows a χ^2 test for the first 6 lags for the small-sample data set.

McCain and McCleary (1979) point out that sometimes a time series cannot be identified as any ARIMA (p, *d, q*) after many attempts. They suggest trying out a natural log transform of the DV before again attempting identification.

TABLE 18.19 ACF and PACF for Common ARIMA Models. Adapted from Dixon (1992). Used with permission.

<table>
<tr><th rowspan="2">Model</th><th rowspan="2">Lag</th><th>ACF</th><th>PACF</th></tr>
<tr><th>– 0 +</th><th>– 0 +</th></tr>
<tr><td rowspan="10">ARIMA (1, 0, 0)</td><td>1</td><td>|__________</td><td>|________</td></tr>
<tr><td>2</td><td>|________</td><td>|</td></tr>
<tr><td>3</td><td>|_____</td><td>|</td></tr>
<tr><td>4</td><td>|___</td><td>|</td></tr>
<tr><td>5</td><td>|_</td><td>|</td></tr>
<tr><td>6</td><td>|</td><td>|</td></tr>
<tr><td>7</td><td>|</td><td>|</td></tr>
<tr><td>8</td><td>|</td><td>|</td></tr>
<tr><td>9</td><td>|</td><td>|</td></tr>
<tr><td>10</td><td>|</td><td>|</td></tr>
<tr><td rowspan="10">ARIMA (0, 0, 1)</td><td>1</td><td>|________</td><td>|__________</td></tr>
<tr><td>2</td><td>|</td><td>|__________</td></tr>
<tr><td>3</td><td>|</td><td>|_______</td></tr>
<tr><td>4</td><td>|</td><td>|____</td></tr>
<tr><td>5</td><td>|</td><td>|_</td></tr>
<tr><td>6</td><td>|</td><td>|</td></tr>
<tr><td>7</td><td>|</td><td>|</td></tr>
<tr><td>8</td><td>|</td><td>|</td></tr>
<tr><td>9</td><td>|</td><td>|</td></tr>
<tr><td>10</td><td>|</td><td>|</td></tr>
<tr><td rowspan="10">ARIMA (2, 0, 0)</td><td>1</td><td>|__________</td><td>|_________</td></tr>
<tr><td>2</td><td>|___________</td><td>|__________</td></tr>
<tr><td>3</td><td>|_______</td><td>|</td></tr>
<tr><td>4</td><td>|____</td><td>|</td></tr>
<tr><td>5</td><td>|_</td><td>|</td></tr>
<tr><td>6</td><td>|</td><td>|</td></tr>
<tr><td>7</td><td>|</td><td>|</td></tr>
<tr><td>8</td><td>|</td><td>|</td></tr>
<tr><td>9</td><td>|</td><td>|</td></tr>
<tr><td>10</td><td>|</td><td>|</td></tr>
<tr><td rowspan="10">ARIMA (0, 0, 2)</td><td>1</td><td>|__________</td><td>|___________</td></tr>
<tr><td>2</td><td>|___________</td><td>|___________</td></tr>
<tr><td>3</td><td>|_______</td><td>|_______</td></tr>
<tr><td>4</td><td>|____</td><td>|_____</td></tr>
<tr><td>5</td><td>|_</td><td>|___</td></tr>
<tr><td>6</td><td>|</td><td>|_</td></tr>
<tr><td>7</td><td>|</td><td>|</td></tr>
<tr><td>8</td><td>|</td><td>|</td></tr>
<tr><td>9</td><td>|</td><td>|</td></tr>
<tr><td>10</td><td>|</td><td>|</td></tr>
</table>

TABLE 18.19 Continued

<table>
<tr><th></th><th></th><th>ACF</th><th>PACF</th></tr>
<tr><th>Model</th><th>Lag</th><th>– 0 +</th><th>– 0 +</th></tr>
<tr><td rowspan="10">ARIMA (p, 0, q)</td><td>1</td><td>|_________</td><td>|__________</td></tr>
<tr><td>2</td><td>|___________</td><td>|___________</td></tr>
<tr><td>3</td><td>|___________</td><td>|______</td></tr>
<tr><td>4</td><td>|_________</td><td>|__________</td></tr>
<tr><td>5</td><td>|_______</td><td>|______</td></tr>
<tr><td>6</td><td>|_____</td><td>|____</td></tr>
<tr><td>7</td><td>|___</td><td>|_</td></tr>
<tr><td>8</td><td>|</td><td>|</td></tr>
<tr><td>9</td><td>p |</td><td>q |</td></tr>
<tr><td>10</td><td>|</td><td>|</td></tr>
<tr><td rowspan="10">ARIMA (0, 1, 0)
(a)</td><td>1</td><td>|__________</td><td>|___________</td></tr>
<tr><td>2</td><td>|__________</td><td>|</td></tr>
<tr><td>3</td><td>|__________</td><td>|</td></tr>
<tr><td>4</td><td>|__________</td><td>|</td></tr>
<tr><td>5</td><td>|__________</td><td>|</td></tr>
<tr><td>6</td><td>|__________</td><td>|</td></tr>
<tr><td>7</td><td>|__________</td><td>|</td></tr>
<tr><td>8</td><td>|__________</td><td>|</td></tr>
<tr><td>9</td><td>|__________</td><td>|</td></tr>
<tr><td>10</td><td>|</td><td>|</td></tr>
<tr><td rowspan="12">ARIMA (0, 1, 0)
(b)</td><td>1</td><td>|________</td><td>|__________</td></tr>
<tr><td>2</td><td>|______</td><td>|</td></tr>
<tr><td>3</td><td>|____</td><td>|</td></tr>
<tr><td>4</td><td>|__</td><td>|</td></tr>
<tr><td>5</td><td>__|</td><td>|</td></tr>
<tr><td>6</td><td>____|</td><td>|</td></tr>
<tr><td>7</td><td>____|</td><td>|</td></tr>
<tr><td>8</td><td>__|</td><td>|</td></tr>
<tr><td>9</td><td>|__</td><td>|</td></tr>
<tr><td>10</td><td>|____</td><td>|</td></tr>
<tr><td>11</td><td>|____</td><td>|</td></tr>
<tr><td>12</td><td>|__</td><td>|</td></tr>
</table>

Figure 18.4 shows ACF and PACF plots for the small-sample example of Table 18.2. The ACF and PACF plots for the example show a classic moving average pattern, but the data are better modeled by auto-regression. This emphasizes the importance of the estimation and diagnostic phases, and also the fact that ARIMA modeling is usually an exploratory procedure. Especially if there is an intervention, the goal of identification, estimation, and diagnosis of the baseline data is to reduce the

pattern of autocorrelation to random error so that the intervention can be properly tested. If the data are poorly fit by a model, the identification process is revisited until a satisfactory model is found.

18.6.2 Effect Size

Effect size in time-series analysis is based on evaluation of residuals after finding an acceptable model. If there is an intervention, the residuals after the intervention is modeled are compared with the residuals before the intervention is modeled.

Two measures of effect size are available (McCleary and Hay, 1980). One is a goodness-of-fit statistic that is traditionally used with time-series analysis. However, it has no ready interpretation as variance in the DV accounted for by the model. The other, less traditional, statistic is more easily interpreted.

The traditional goodness-of-fit measure for time-series analysis is the residual mean square:

$$RMS = 1/N \sqrt{\sum_{t=1}^{N} \hat{a}_t^2} \tag{18.14}$$

The residual mean square (*RMS*) is the average of the square root of the squared residual values ($\hat{a}_t$) summed over the N time periods.

Both of the ARIMA programs provide *RMS*, as seen in Tables 18.6 and 18.11. However, SPSS ARIMA calls the value `Residual Variance` and SAS ARIMA calls it `Variance Estimate`. They also give slightly different estimates of *RMS* for the small-sample example.

The second, more interpretable, measure of effect size is R^2, which, as usual, reflects variance explained by the model as a proportion of total variance in the DV.

$$R^2 = 1 - \frac{SS_{\hat{a}}}{SS_{y_t}} \tag{18.15}$$

The proportion of systematic variance explained by the model (R^2) is one minus the sum of squared residuals divided by the sum of squared Y_t values, where Y_t is the difference-adjusted DV.

SAS ARIMA provides the standard deviation of the differenced series which is first converted to variance by squaring and then converted to SS by multiplication by $(N - 1)$. Table 18.13 shows that `Standard deviation` = 6.44, so that the variance $\left(s_y^2\right) = (6.44)^2 = 41.47$. With $N = 40$, $SS_{y_t} = (41.47)(39) = 1617.33$ for the small-sample example. Table 18.13 shows the variance for the residuals: `Variance estimate` $\left(s_{\hat{a}}^2\right) = 23.35$, with df = ($N$ − number of parameters − 1) = (40 − 3 − 1) = 36, so that $SS_{\hat{a}} = (23.35)(36) = 840.60$. (In this example, there is one parameter for auto-regression, one for trend, and one for intervention.) When these values are plugged into Equation 18.15:

$$R^2 = 1 - \frac{840.60}{1617.33} = .48$$

Thus, 48% of the variance in the intervention analysis of Table 18.13 is explained by the ARIMA (1, 1, 0) model with an abrupt, permanent effect.

If SPSS ARIMA is used, `Residual Variance` is the post-model $SS_{\hat{a}}$. The variance of the differenced series $(S^2_{y_t})$ is found by hand calculation in the usual manner.

An alternative form of R^2, useful when the intervention is statistically significant, is to use $s^2_{\hat{a}}$ for the model without the intervention parameters (but with the modeling parameters included) instead of SS_{y_t} in the denominator of Equation 18.15. Effect size for a model with a statistically significant intervention is demonstrated in Section 18.7.4.2.

18.6.3 Forecasting

Forecasting refers to the process of predicting future observations from a known time series and is often the major goal in nonexperimental use of the analysis. However, prediction beyond the data is to be approached with caution. The farther the prediction beyond the actual data, the less reliable the prediction. And only a small percentage of the number of actual data points can be predicted before the forecast turns into a straight line.

Table 18.20 shows the forecast for the next 7 weeks of the data of Table 18.2 through SAS ARIMA. The `forecast lead = 7` instruction indicates that the 7 weeks ahead are to be forecast.

The `Obs` column shows the week being forecast; the `Forecast` column shows the predicted `quality` for the week. Standard errors and the 95% confidence intervals for `quality` also are shown and can be seen to increase as the prediction progresses.

Sometimes it is interesting to compare the forecasted values of the DV without an intervention with those that actually occurred after an intervention. Figure 18.8 plots the nonintervention predicted values as well as the actual pre-intervention and first 7 post-intervention values for the data of Table 18.11 (abrupt, permanent intervention).

Figure 18.8 highlights the impact of the intervention by showing graphically the expected values of quality change without intervention, and the difference between those values and the actual values after intervention.

18.6.4 Statistical Methods for Comparing Two Models

Often the identification process suggests two or more candidate models. Sometimes it is obvious which is better because only one meets the diagnostic criteria of Section 18.6.3. Sometimes, however, both models fail to meet the criteria, or both meet all of the criteria. Statistical tests are available to determine whether two models are reliably different.

Say, for example, you wish to compare the abrupt, permanent model of Table 18.14 with the gradual, permanent model of Table 18.16. The difference between the models is that δ, the parameter for gradualness, has been added to the ARIMA (1, 1, 0) model with intervention, ω. Two tests are available through SAS and SPSS ARIMA: AIC (Akaike's Information Criterion) and SBC (Schwarz's Bayesian Criterion). AIC is the more commonly encountered.

The difference in `AIC` for the two models is evaluated as χ^2 with df equal to the difference in the number of parameters for the two models.

$$\chi^2\,(\text{df}) = \text{AIC(smaller model)} - \text{AIC(larger model)} \tag{18.16}$$

TABLE 18.20 Forecasting through SAS ARIMA (Syntax and Selected Output)

```
proc arima data=SASUSER.SSTIMSER;
   identify var=QUALITY(1) nlag=7;
   estimate p=1 noint;
   forecast lead = 7;
run;
```

Forecasts for variable quality

Obs	Forecast	Std Error	95% Confidence Limits	
21	25.9983	4.6679	16.8494	35.1472
22	31.8190	4.8382	22.3363	41.3016
23	27.5848	6.1166	15.5964	39.5732
24	30.6649	6.4186	18.0846	43.2451
25	28.4243	7.1957	14.3210	42.5277
26	30.0542	7.5549	15.2468	44.8615
27	28.8686	8.1263	12.9412	44.7959

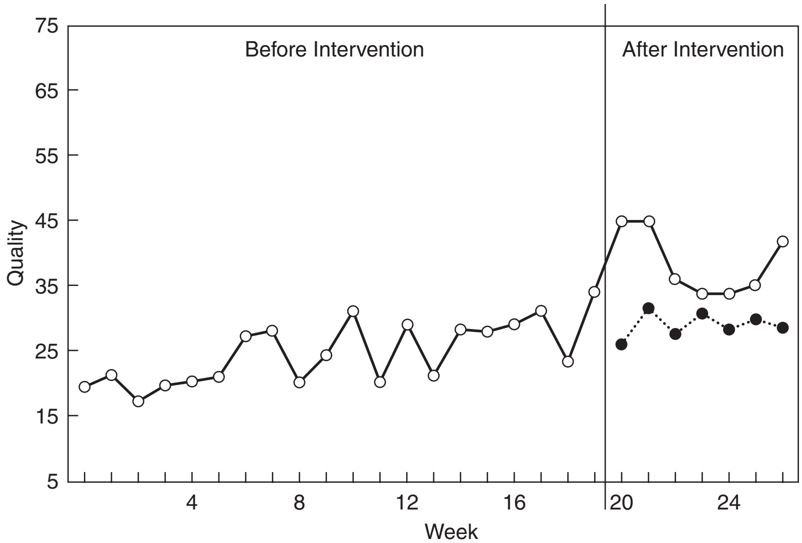

FIGURE 18.8 Actual (solid) and predicted (dotted) quality (first-lag differenced) over 27 weeks, before and after introduction of the profit-sharing plan. Generated in Quattro Pro 8.

For the sample data, the larger model is in Table 18.16 (gradual onset).

$$\chi^2(1) = 235.5006 - 232.5722 = 2.93$$

Critical χ^2 (at $\alpha = .05$ and 1 df) $= 3.84$, as seen in Table C.4. The more parsimonious model (1, 1, 0), without the δ parameter, is not significantly worse than the more complex model of gradual, permanent change, consistent with the nonsignificant δ parameter in Table 18.16. When two models are not statistically different, the more parsimonious model is selected.

AIC and SBC criteria are only used to compare nested models. This means that the smaller model (e.g., 0, 1, 0) must be a subset of the larger model (1, 1, 0). The tests are not legitimately used to compare, for example, (1, 0, 1) and (0, 1, 0). And they could not be used to compare models in which a different value of ω is generated by coding the IV differently.

18.7 Complete Example of a Time-Series Analysis

The data for this example were collected between January 1980 and July 1990 to assess the impact of a seat belt law enacted in 1985 in Illinois. The Illinois Department of Transportation collected monthly data on number of accidents, deaths, and injuries (broken down by level of severity). Data were provided by Rock (1992), who analyzed five DVs using ARIMA models. A more complete description of the data set is in Appendix B. The current example focuses on incapacitating injuries, those identified as “A” level. Data are in TIMESER.*.

The first 66 months serve as a baseline period, and the remaining 66 months are the post-intervention series. Figure 18.9 plots the raw time-series data.

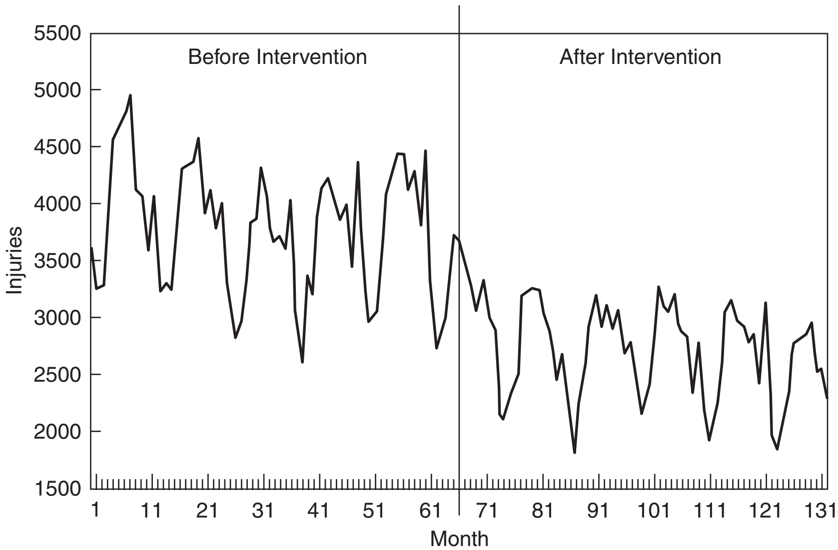

FIGURE 18.9 A-level injuries over 132 months, before and after introduction of seat belt law. Generated in Quattro Pro 8.

18.7.1 Evaluation of Assumptions

18.7.1.1 Normality of Sampling Distributions

Normality of sampling distributions is evaluated by examining residuals for the ARIMA model for the baseline period as part of the diagnostic process before the intervention series is included.

18.7.1.2 Homogeneity of Variance

Homogeneity of variance is usually evaluated after initial model identification by examining plots of standardized residuals. However, Figure 18.9 suggests heterogeneity of variance because swings appear to be diminishing over time, particularly after introduction of the seat belt law. However, this heterogeneity will not show up during baseline modeling because the decrease in variance occurs after intervention. Therefore, the decision is made to use a logarithmic transform of the data before baseline modeling. Figure 18.10 shows the time series after transformation.

18.7.1.3 Outliers

Figure 18.10 reveals no obvious outliers among the log-transformed observations. Presence of outliers is rechecked after modeling.

18.7.2 Baseline Model Identification and Estimation

ACF and PACF plots are produced through SAS in Table 18.21 for the pre-intervention series. The analysis is limited to the pre-intervention series by selecting cases with `BELT = 0`. The variable ana-

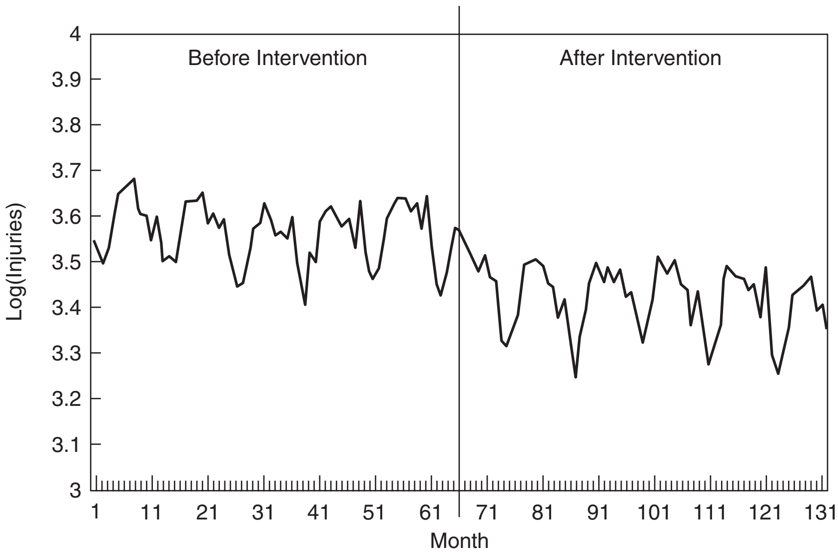

FIGURE 18.10 Logarithmic transform of A-level injuries over 132 months, before and after introduction of seat belt law. Generated in Quattro Pro 8.

lyzed is the transformed one, `log_inj`. Only 25 lags are requested to limit length of output and because later lags are unlikely to show significant autocorrelations.

The plots show a pattern of (0, 1, 0) when compared with Table 18.19, with large positive spikes at lag = 12, not at all surprising with monthly data. The data first are differenced to see if the same pattern emerges, but with the spikes closer together. Table 18.22 shows a run with differencing at lags 1 and 12, requested by adding `(1 12)` to the `var=` instruction.

Table 18.22 shows that first-differencing with lag = 1 and 12 still produces spikes on the ACF and PACF at lag = 1, and the pattern now resembles an ARIMA model (p, 0, q) with both AR and MA parameters. Thus, either an additive or multiplicative seasonal model is suggested with both auto-regressive and moving average parameters at both lags 1 and 12. In addition, either an additive or multiplicative seasonal model is suggested with AR1, AR12, MA1, and MA12 parameters. That is, suggested models are ARIMA (2, 2, 2) or ARIMA $(1, 1, 1)(1, 1, 1)_{12}$. The more parsimonious (2, 2, 2) model, however, produced residual ACF and PACF PLOTS (not shown) that clearly signaled the need for further modeling, strongly resembling the (0, 1, 0) (b) model of Table 18.19.

Table 18.23 shows SAS ARIMA syntax and selected output for the $(1, 1, 1)(1, 1, 1)_{12}$ model. The ARIMA paragraph requests a $(1, 1, 1)(1, 1, 1)_{12}$ model with auto-regressive parameters (p) at lags 1 and 12, and moving average parameters (q) at lags 1 and 12, as well as differencing parameters both at lags 1 and 12. The forecast instruction is necessary to produce residuals, which are saved in a temporary file named RESID. We also ask that residuals go back 51 cases from the end of the data and give us 52 values, so that instead of forecasting we are actually looking at predicted values for existing data and differences between those and actual values (i.e., residuals).

18.7.3 Baseline Model Diagnosis

Table 18.23 shows that only the two moving average parameters (`0.55936` and `0.57948`) are significantly different from zero at $\alpha = .05$, with the t Value for MA1,1 = 2.39, Pr>|t| = 0.0208, and MA2,1 = 2.38, Pr>|t| = 0.0238 for the second. Both parameters are between −1 and 1. Neither of the auto-regressive components is significant. Therefore, the suggested model is ARIMA $(0, 1, 1)(0, 1, 1)_{12}$.

The `Autocorrelation Check of Residuals` shows acceptable residuals; with 6 lags, $\chi^2(2) = 5.54$, below the $\alpha = .05$ criterion of 5.99. Although ACF and PACF for the residuals (the

TABLE 18.21 ACF and PACF Plots for the Transformed Pre-Seat Belt Law Time Series (SAS ARIMA Syntax and Selected Output)

```
proc arima data=SASUSER.TIMESER;
  where BELT=0;
  identify var=LOG_INJ nlag=25;
run;

                    Name of Variable = LOG_INJ

                 Mean of Working Series   3.570007
                 Standard Deviation       0.065699
                 Number of Observations         66
```

(continued)

TABLE 18.21 Continued

```
                                      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    Std Error

  0    0.0043164      1.00000     |                   |*******************|            0
  1    0.0026667      0.61780     |              .    |************       |     0.123091
  2    0.0013585      0.31474     |            .      |******.            |     0.163455
  3   -0.0002082      -.04824     |            .     *|      .            |     0.172393
  4   -0.0011075      -.25658     |            . *****|      .            |     0.172598
  5   -0.0014457      -.33493     |            *******|      .            |     0.178283
  6   -0.0018458      -.42762     |          *********|       .           |     0.187575
  7   -0.0015799      -.36601     |           .*******|       .           |     0.201806
  8   -0.0015283      -.35406     |          . *******|        .          |     0.211625
  9   -0.0007142      -.16545     |          .     ***|        .          |     0.220417
 10   0.00061872      0.14334     |          .        |***     .          |     0.222291
 11    0.0017584      0.40738     |          .        |********.          |     0.223687
 12    0.0028752      0.66612     |          .        |*************      |     0.234660
 13    0.0018368      0.42555     |         .         |********* .        |     0.261747
 14   0.00089938      0.20836     |         .         |****      .        |     0.272027
 15   -0.0001789      -.04146     |         .        *|          .        |     0.274435
 16   -0.0008805      -.20399     |         .     ****|          .        |     0.274530
 17   -0.0010549      -.24438     |         .    *****|          .        |     0.276817
 18   -0.0016480      -.38180     |         . ********|          .        |     0.280067
 19   -0.0014890      -.34497     |        .   *******|           .       |     0.297845
 20   -0.0016951      -.39272     |        .  ********|           .       |     0.294042
 21   -0.0009842      -.22801     |        .     *****|           .       |     0.301885
 22   0.00006003      0.01391     |        .          |           .       |     0.304483
 23    0.0011643      0.26975     |        .          |*****      .       |     0.304493
 24    0.0021433      0.49656     |        .          |********** .       |     0.308092
 25    0.0014855      0.34416     |       .           |*******     .      |     0.319989

                                "." marks two standard errors
```

TABLE 18.21 Continued

```
                    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      0.61780    |                .    |************        |
   2     -0.10826    |                .  **|    .               |
   3     -0.32214    |              ******|    .               |
   4     -0.11241    |                .  **|    .               |
   5     -0.04368    |                .   *|    .               |
   6     -0.28223    |              ******|    .               |
   7     -0.05955    |                .   *|    .               |
   8     -0.20918    |                .****|    .               |
   9      0.03416    |                .    |*    .              |
  10      0.28546    |                .    |******             |
  11      0.17401    |                .    |***  .              |
  12      0.35243    |                .    |*******            |
  13     -0.37758    |            ********|     .              |
  14     -0.10238    |                .  **|     .              |
  15      0.10593    |                .    |**   .              |
  16     -0.05754    |                .   *|     .              |
  17      0.07091    |                .    |*    .              |
  18     -0.17099    |                . ***|     .              |
  19     -0.00659    |                .    |     .              |
  20     -0.08356    |                .  **|     .              |
  21     -0.03547    |                .   *|     .              |
  22     -0.05915    |                .   *|     .              |
  23      0.04632    |                .    |*    .              |
  24      0.05709    |                .    |*    .              |
  25     -0.07699    |                .  **|     .              |
```

TABLE 18.22 ACF and PACF Plots for the Pre-Seat Belt Law Time Series with LAG 1 and 12 Differencing (SAS ARIMA Syntax and Selected Output)

```
proc arima data=SASUSER.TIMESER;
  where BELT=0;
  identify var=LOG_INJ(1 12) nlag=25;
run;
```

```
                    Name of Variable = LOG_INJ

       Period(s) of Differencing                         1,12
       Mean of Working Series                         -0.0001
       Standard Deviation                            0.036002
       Number of Observations                              53
       Observation(s) eliminated by differencing           13
```

(continued)

TABLE 18.22 Continued

```
                                   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     Std Error

  0   0.0012961       1.00000    |                    |********************|            0
  1  -0.0004981       -.38433    |            ********|     .              |     0.137361
  2  -0.0000227       -.01753    |              .     |     .              |     0.156339
  3  -0.0000643       -.04963    |              .    *|     .              |     0.156376
  4  0.00022946       0.17704    |              .     |**** .              |     0.156673
  5  -0.0004657       -.35927    |            ********|      .             |     0.160403
  6  0.00022242       0.17161    |             .      |***   .             |     0.174928
  7  0.00003158       0.02436    |             .      |*     .             |     0.178076
  8  0.00004115       0.03175    |             .      |*     .             |     0.178139
  9  -0.0000145       -.01121    |             .      |      .             |     0.178246
 10  0.00014656       0.11308    |             .      |**    .             |     0.178259
 11  0.00013478       0.10399    |             .      |**    .             |     0.179607
 12  -0.0005353       -.41304    |            ********|      .             |     0.180740
 13  0.00013411       0.10347    |            .       |**     .            |     0.197749
 14  0.00003295       0.02542    |            .       |*      .            |     0.198768
 15  0.00005370       0.04143    |            .       |       .            |     0.198829
 16  -0.0001772       -.13671    |            .    ***|       .            |     0.198992
 17  0.00034685       0.26761    |            .       |****** .            |     0.200756
 18  -0.0002298       -.17729    |            .    ***|       .            |     0.207378
 19  0.00021274       0.16413    |           .        |***     .           |     0.210218
 20  -0.0001927       -.14866    |           .     ***|        .           |     0.212622
 21  0.00012719       0.09813    |           .        |**      .           |     0.214574
 22  -0.0002876       -.22190    |           .    ****|        .           |     0.215419
 23  0.00018294       0.14114    |           .        |***     .           |     0.219690
 24  -0.0001048       -.08088    |           .       *|        .           |     0.221394
 25  0.00024906       0.19216    |           .        |****    .           |     0.221951

                           "." marks two standard errors
```

TABLE 18.22 Continued

```
                     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     -0.38433   |           ********|     .             |
       2     -0.19387   |             . ****|     .             |
       3     -0.16115   |             .  ***|     .             |
       4      0.10775   |             .     |**   .             |
       5     -0.31396   |             ******|     .             |
       6     -0.09584   |             .   **|     .             |
       7     -0.00772   |             .     |     .             |
       8      0.01194   |             .     |     .             |
       9      0.11743   |             .     |**   .             |
      10      0.08823   |             .     |**   .             |
      11      0.32412   |             .     |******             |
      12     -0.30500   |             ******|     .             |
      13     -0.17454   |             .  ***|     .             |
      14     -0.10257   |             .   **|     .             |
      15     -0.01554   |             .     |     .             |
      16      0.01689   |              .    |     .             |
      17     -0.03651   |             .    *|     .             |
      18     -0.09992   |             .   **|     .             |
      19      0.14423   |             .     |***  .             |
      20      0.01614   |             .     |     .             |
      21      0.10803   |             .     |**   .             |
      22     -0.06235   |             .    *|     .             |
      23      0.07362   |             .     |*    .             |
      24     -0.16794   |             .  ***|     .             |
      25      0.01157   |             .     |     .             |
```

variable called RESIDUAL in the RESID file) show significant spikes at lag 5. However, attempts to improve the model were unsuccessful.

Residuals are examined through SAS ARIMA using the file of residuals just created. Figure 18.11 shows SAS UNIVARIATE syntax and output for a normal probability plot of residuals.

There is some departure of the plotted *z*-scores from the diagonal line representing a normal distribution of them; the most negative residuals are a bit too high. A normal probability plot of the

TABLE 18.23 ARIMA (1, 1, 1)(1, 1, 1)$_{12}$ Model for the Transformed Pre-Intervention Time Series

```
proc arima data=SASUSER.TIMESER;
  where BELT=0
  identify var=LOG_INJ(1, 12) nlag=36;
  estimate noint p=(1)(12) q=(1)(12);
  forecast out=RESID back=51 lead=52;
run;
```

(continued)

TABLE 18.23 Continued

Conditional Least Squares Estimation

Parameter	Estimate	Approx Std Error	t Value	Pr > |t|	Lag
MA1,1	0.55936	0.23415	2.39	0.0208	1
MA2,1	0.57948	0.24833	2.33	0.0238	12
AR1,1	0.03823	0.27667	0.14	0.8907	1
AR2,1	0.09832	-0.27755	-0.35	0.7247	12

Variance Estimate	0.000789
Std Error Estimate	0.028089
AIC	-224.425
SBC	-216.544
Number of Residuals	53

* AIC and SBC do not include log determinant.

Correlations of Parameter Estimates

Parameter	MA1,1	MA2,1	AR1,1	AR2,1
MA1,1	1.000	-0.034	0.852	0.072
MA2,1	-0.034	1.000	0.020	0.830
AR1,1	0.852	0.020	1.000	0.104
AR2,1	0.072	0.830	0.104	1.000

Autocorrelation Check of Residuals

To Lag	Chi-Square	DF	Pr > ChiSq	--------------------Autocorrelations--------------------					
6	5.54	2	0.0625	-0.000	0.004	0.025	0.044	-0.272	0.120
12	7.09	8	0.5269	-0.009	0.073	0.119	0.061	-0.007	-0.002
18	8.52	14	0.8606	-0.018	-0.054	0.009	-0.097	0.069	-0.027
24	12.72	20	0.8891	0.112	-0.091	0.079	-0.118	0.073	-0.011
30	21.89	26	0.6944	0.225	-0.059	-0.065	-0.089	-0.039	-0.124
36	27.14	32	0.7111	-0.040	0.034	0.072	0.149	0.050	0.043

```
proc univariate data=WORK.RESID plot;
 var RESIDUAL
RUN
```

```
                The UNIVARIATE Procedure

    Variable:  RESIDUAL  (Residual: Actual-Forecast)

                Normal Probability Plot
 0.185+                                                *
      |                                              ++
      |                                            *+
      |                                          *+
      |                                        ++
      |                                       **
      |                                    ***
 0.115+                                   *+
      |                                ***
      |                              **+
      |                           ***++
      |                            ++
      |                          **
      |                        +**
 0.045+                       +**
      |                     ++**
      |                   ++ *
      |                 ++****
      |                + *
      |              ****
      |          ****
-0.025+  *   * * ++
       +----+----+----+----+----+----+----+----+----+----+
           -2        -1         0        +1        +2
```

FIGURE 18.11 Normal probability plot of residuals for transformed pre-intervention ARIMA $(1, 1, 1)(1, 1, 1)_{12}$ model. SAS univariate syntax and selected output.

(transformed) scores themselves (not shown), however, indicates no such problems. Additional transformation is unlikely to alleviate the problem in the residuals which may have been introduced by the model itself. [Residuals for the $(0, 1, 1)(0, 1, 1)_{12}$ model were not noticeably different.]

18.7.4 Intervention Analysis

18.7.4.1 Model Diagnosis

The introduction of the seat belt law was expected to have an immediate, continuing effect, therefore all post-intervention months are coded 1, and the model chosen is the abrupt, permanent (step) function of Section 18.5.2.1. Table 18.24 shows the SAS ARIMA re-estimation of the entire series using the ARIMA (0, 1, 1)(0, 1, 1)$_{12}$ model, with BELT as the intervention variable. Note that the seasonal differencing requires additional syntax in an intervention analysis. First, the `crosscor` instruction is expanded to include differencing in the intervention: `BELT(1, 12)`. Then, `MONTH (12)` needs to be added as well. Finally, `MONTH` is added to `BELT` in the `input` instruction, to indicate that it is an "intervention" component. The `outlier` syntax checks for changes in level unaccounted for by the model.

The moving average parameters are still statistically significant at $\alpha = .05$ and within limits of 1 and -1. The intervention parameter is also significant at $\alpha = .05$ with `Pr > |t| = -2.92` for `NUM1`. The `Autocorrelation Check of Residuals` indicates no significant deviation of residuals from chance to lag 24.

Remaining output summarizes model information and shows that there are no level shifts that are not accounted for by the model at $\alpha = .001$.

18.7.4.2 Model Interpretation

The negative sign of the intervention parameter (-0.06198) indicates that incapacitating injuries are reduced after introduction of the seat belt law. The anti-log of the intervention parameter is $10^{-0.06198} = 0.87$. In percentage terms, this indicates that A-level injuries are reduced to 87% of their pre-intervention mean, or by about 13% after implementing the seat belt law. Interpretation is aided by finding the median (untransformed) number of injuries before and after intervention.[4] Table 18.25 shows a SAS MEANS run for medians. (Note that data are already sorted by `BELT`.)

Notice that the number of injuries predicted by the model after intervention (.87)(3801.5) = 3307.3 is greater than the actual number of injuries observed (2791.5). This is because the model also adjusts the data for trends and persistence of random shocks (moving average components), and thus more accurately reflects the impact of the intervention.

TABLE 18.24 Step Intervention Test of ARIMA (0, 1, 1)(0, 1, 1)$_{12}$ Model of Injury Data (SAS ARIMA Syntax and Selected Output)

```
proc arima data=SASUSER.TIMESER;
  identify var=LOG_INJ (1, 12) nlag=10 crosscor=(BELT(1, 12) MONTH(12));
  estimate noint q=(1)(12) input= (BELT MONTH);
  outlier id=month alpha=.001;
run;
```

```
                    The ARIMA Procedure

                Name of Variable = LOG_INJ

     Period(s) of Differencing                       1,12
     Mean of Working Series                      -0.00055
     Standard Deviation                          0.036572
     Number of Observations                           119
     Observation(s) eliminated by differencing         13
```

[4]Recall from Section 4.1.6 that the median of the raw scores is an appropriate measure of central tendency after data have been transformed.

TABLE 18.24 Continued

Conditional Least Squares Estimation

Parameter	Estimate	Approx Std Error	t Value	Pr > |t|	Lag	Variable	Shift
MA1,1	0.58727	0.07927	7.41	<.0001	1	LOG_INJ	0
MA2,1	0.76890	0.06509	11.81	<.0001	12	LOG_INJ	0
NUM1	-0.06198	0.02119	-2.92	0.0042	0	BELT	0
NUM2	0.00001375	0.00002794	0.49	0.6234	0	MONTH	0

Variance Estimate	0.000719
Std Error Estimate	0.026808
AIC	-519.697
SBC	-508.581
Number of Residuals	119

* AIC and SBC do not include log determinant.

Correlations of Parameter Estimates

Variable	Parameter	LOG_INJ MA1,1	LOG_INJ MA2,1	BELT NUM1	MONTH NUM2
LOG_INJ	MA1,1	1.000	-0.110	-0.238	0.032
LOG_INJ	MA2,1	-0.110	1.000	0.108	0.203
BELT	NUM1	-0.238	0.108	1.000	-0.015
MONTH	NUM2	0.032	0.203	-0.015	1.000

Autocorrelation Check of Residuals

To Lag	Chi-Square	DF	Pr > ChiSq	Autocorrelations					
6	3.50	4	0.4774	0.053	-0.078	0.015	-0.005	-0.137	0.012
12	6.87	10	0.7377	0.067	0.004	-0.008	0.109	0.092	0.025
18	10.66	16	0.8302	-0.031	-0.086	-0.125	-0.040	0.030	0.030
24	14.88	22	0.8671	0.074	-0.015	0.025	-0.061	0.132	-0.032

(continued)

TABLE 18.24 Continued

```
                Model for variable LOG_INJ

           Period(s) of Differencing      1,12

                No mean term in this model

                  Moving Average Factors

             Factor 1:    1 - 0.58727 B**(1)
             Factor 2:    1 - 0.7689 B**(12)

                      Input Number 1

          Input Variable                      BELT
          Period(s) of Differencing           1.12
          Overall Regression Factor       -0.06198

                      Input Number 2

          Input Variable                     MONTH
          Period(s) of Differencing             12
          Overall Regression Factor       0.000014

               Outlier Detection Summary

          Maximum number search                  3
          Number found                           0
          Significant used                   0.001
```

Effect size for the ARIMA intervention model is estimated using Equation 18.15, substituting $s_{\hat{a}}^2$ from the model without intervention as an estimate of SS_{y_t} the ARIMA model without intervention parameters. The `Variance Estimate` for the ARIMA model without intervention parameters (not shown) is 0.000752. Degrees of freedom are $119 - 4 - 1 = 114$, so that $SS_{y_t} = (0.000753)(114) = 0.0858$. The `Variance Estimate` for the full model is 0.000719, as seen in Table 18.24. Degrees of freedom are $119 - 6 - 1 = 112$, so that $SS_{\hat{a}} = (0.000719)(112) = 0.0805$. Thus, effect size using Equation 18.15 is:

$$R^2 = 1 - \frac{0.0805}{0.0858} = .06$$

That is, the intervention accounts for 6% of the variance in the ARIMA $(0, 1, 1)(0, 1, 1)_{12}$ model.

Table 18.26 is a checklist for an ARIMA intervention analysis. An example of a Results section, in journal format, follows.

TABLE 18.25 Descriptive Statistics for Untransformed Time Series Data (SAS MEANS Syntax and Selected Output)

```
proc means data=SASUSER.TIMESER
   N MEDIAN;
  by BELT;
  var INJURIES;
run;
```

```
----------------------------- BELT=0 -----------------------------
                      The MEANS Procedure
                 Analysis Variable : INJURIES
                         N      Median
                       --------------
                       66     3801.50
                       --------------

----------------------------- BELT=1 -----------------------------
                 Analysis Variable : INJURIES
                         N      Median
                       --------------
                       66     2791.50
                       --------------
```

TABLE 18.26 Checklist for ARIMA Intervention Analysis

1. Issues
   a. Normality of sampling distributions
   b. Homogeneity of variance
   c. Outliers
2. Major analyses
   a. Identification, estimation, and diagnosis of baseline model
   b. Diagnosis of intervention model
   c. Interpretation of intervention model parameters, if significant
   d. Effect size
3. Additional analyses
   a. Tests of alternative intervention models, if nonsignificant
   b. Forecast of future observations
   c. Interpretation of ARIMA model parameters, if significant

Results

A time-series model for incapacitating injuries was developed to examine the effect of a seat belt law, introduced in Illinois in January 1985. Data were collected for 66 months before and 66 months after implementing the law. As seen in Figure 18.9, there are no obvious outliers among the observations, but variance appears to decrease over the time series, particularly after intervention. Therefore, a logarithmic transform was applied to the series, producing Figure 18.9. The 66-month pre-intervention series was used to identify a seasonal ARIMA $(0, 1, 1)(0, 1, 1)_{12}$ model, with differencing required at lags 1 and 12 to achieve stationarity.

The local moving average parameter was 0.58799 and the seasonal parameter 0.55492, both statistically significant, with t = 7.41 and t = 11.81, $p < .05$ respectively. The intervention parameter (-0.06198), was strong and statistically significant, t = -2.92, $p > .05$, $R^2 = .06$. The antilog of the parameter ($10^{-0.062}$) was 0.87, suggesting that the impact of the seat belt law was a 13% reduction in mean number of incapacitating injuries per month. With 3,802 injuries per month on average before the law, the reduction is approximately 500 per month. Thus, the seat belt law is shown to be effective in reducing incapacitating injuries.

18.8 Comparison of Programs

SAS and SPSS each have a single program for ARIMA time-series modeling. SPSS has additional programs for producing time-series graphs and for forecasting seasonally adjusted time series. SAS has a variety of programs for time-series analysis, including a special one for seasonally adjusted time series. SYSTAT also has a time-series program, but it does not support intervention analysis. The programs reviewed in Table 18.27 are those that do ordinary ARIMA modeling and produce time-series plots.

18.8.1 SPSS Package

SPSS has a number of time-series procedures, only a few of which are relevant to the ARIMA models of this chapter. ACF produces autocorrelation and partial autocorrelation plots, CCF produces cross-correlation plots, and TSPLOT produces the time-series plot itself. All of the plots are produced in high-resolution form; ACF and CCF plots also are shown in low-resolution form with numerical values. This is the only program that permits alternative specification of standard errors for autocorrelation plots, with an independence model possible (IND) as well as the usual Bartlett's approximation (MA) used by other programs.

SPSS ARIMA is available for modeling time series with a basic set of features. Several options for iteration are available, as well as two methods of parameter estimation. Initial values for parameters can be specified, but there is no provision for centering the series (although the mean may be subtracted in a separate transformation). There are built-in functions for a log (base 10 or natural) transform of the series. Differencing is awkward for lag greater than 1. A differenced variable is created, and then the differencing parameter is omitted from the model specification (along with the constant). Residuals (and their upper and lower confidence values) and predicted values are automatically written to the existing data file, and tests are available to compare models (log-likelihood, AIC, and SBC).

18.8.2 SAS System

SAS ARIMA also is a full-featured modeling program, with three estimation methods as well as several options for controlling the iteration process. SAS ARIMA has the most options for saving things to data files. Inverse autocorrelation plots and (optionally) autocorrelation plots of residuals are available in addition to the usual ACF, PACF, and cross-correlation plots that are produced by default. However, a plot of the raw time series itself must be requested outside PROC ARIMA.

The autocorrelation checks for white noise and residuals are especially handy for diagnosing a model, and models may be compared using AIC and SBC. Missing data are estimated under some conditions.

18.8.3 SYSTAT System

Except for graphs and built-in transformations, SYSTAT SERIES is a bare bones program. The time-series plot can be edited and annotated; however, ACF, PACF, and CCF plots have no numerical values. Maximum number of iterations and a convergence criterion can be specified. Iteration history is provided, along with a final value of error variance. Other than that, only the parameter estimates, standard errors, and 95% confidence interval are shown. There is no provision for intervention analysis or any other input variables.

SYSTAT does differencing at lags greater than 1 in a manner similar to that of SPSS. Differenced values are produced as a new variable (in a new file) and then the differenced values are used in the model. Thus, forecast values are not in the scale of the original data, but in the scale of the differenced values, making interpretation more difficult. SYSTAT does, however, provide a plot showing values for the known (differenced) series as well as forecast values.

TABLE 18.27 Comparison of SPSS, SAS, and SYSTAT Programs for ARIMA Time-Series Analysis

Feature	SPSS ARIMA	SAS ARIMA	SYSTAT SERIES
Input			
Specify intervention variable	Yes	Yes	No
Specify additional continuous variable(s)	No	Yes	No
Include constant or not	Yes	Yes	Yes
Specify maximum number of iterations	MXITER	MAXIT	Yes
Specify tolerance	PAREPS	SINGULAR	No
Specify change in sum of squares	SSQPCT	No	No
Parameter estimate stopping criterion	No	CONVERGE	Yes
Specify maximum lambda	MXLAMB	No	No
Specify delta	No	DELTA	No
Define maximum number of psi weights	No	No	No
Conditional least squares estimation method	CLS	CLS	Yes
Unconditional estimation method	EXACT	ULS	No
Maximum likelihood estimation method	No	ML	No
Options regarding display of iteration details	Yes	PRINTALL	No
User-specified initial values for AR and MA parameters and constant	Yes	INITVAL	No
Request that mean be subtracted from each observation	No	CENTER	No
Options for forecasting	Yes	Yes	Yes
Specify size of confidence interval	CINPCT	ALPHA	No
Erase current time-series model	No	CLEAR	CLEAR SERIES
Use a previously defined model without respecification	APPLY	No	No
Request ACF plot with options	Yes[a]	Default	ACF
Request PACF plot with options	Yes[a]	Default	PACF
Request cross-correlation plot with options	Yes[b]	Default	CCF
Request inverse autocorrelation plot	No	Default	No
Request residual autocorrelation plots	No	PLOT	No
Request log transforms of series	Yes	No[e]	Yes
Request time series plot with options	Yes[c]	No[f]	TPLOT
Estimate missing data	No[d]	Yes	Yes
Specify analysis by groups	No	BY	No
Specify number of lags in plots	MXAUTO[a]	NLAG	No
Specify method for estimating standard errors of autocorrelations	SERROR[a]	No	No
Specify method for estimating variance	No	NODF	No
Request serach for outliers in the solution	No	Yes	No
Output			
Number of observations or residuals (after differencing)	Yes	Yes	No

TABLE 18.27 Continued

Feature	SPSS ARIMA	SAS ARIMA	SYSTAT SERIES
Output *(continued)*			
Mean of observations (after differencing)	No	Yes	No
Standard deviation	No	Yes	No
t value of mean (against zero)	No	No	No
Numerical values of autocorrelations (and partial autocorrelations) and standard errors	Yes	Yes	No
Probability value for autocorrelations	Yes	No	No
Covariances for autocorrelation plot	No	Yes	No
Autocorrelation check for white noise	No	Yes	No
Box-Ljung statistic	Yes	No	No
Model parameter estimates	B	Estimate	Yes
Type of parameter	Yes	Yes	Yes
Standard error of parameter estimate	SEB	Approx. Std Error	A.S.E.
t ratio for parameter estimate	T-RATIO	T Ratio	No
Significance value for *t* ratio	APPROX. PROB.	$\Pr > \lvert t \rvert$	No
Confidence interval for parameter estimate	No	No	Yes
Residual sum of squares	(adjusted)	No	No
Residual df	Yes	No	No
Residual mean square (variance)	Yes	Yes	MSE
Standard error of residuals	Yes	Yes	No
Autocorrelation check of residuals	No	Yes	No
Correlations among parameter estimates	Yes	Yes	Yes
Covariances among parameter estimates	Yes	No	No
Log-likelihood	Yes	No	No
AIC	Yes	Yes	No
SBC	Yes	Yes	No
Output data set			
Residual time series	Default	RESIDUAL	Yes
Predicted values for existing time series	Default	FORECAST	No
95% upper and lower confidence interval of predicted values	Default	L95, U95	No
Standard errors of predicted values	Default	STD	No
Forecast values with options	Yes	LEAD	Yes
Parameter estimates for a model	No	OUTEST, OUTMODEL	No
Diagnostic statistics for a model	No	OUTSTAT	No

[a]Done through ACF procedure. [b]Done through CCF procedure.
[c]Done through TSPLOT. [d]Done through MVA.
[e]Done through DATA step. [f]Done through PROC GPLOT.