CHAPTER 16

Time Series Forecasting and Index Numbers

Learning Objectives

After mastering the material in this chapter, you will be able to:

- **LO16-1** Identify the components of a time series.
- **LO16-2** Use time series regression to forecast time series having linear, quadratic, and certain types of seasonal patterns.
- **LO16-3** Use data transformations to forecast time series having increasing seasonal variation.
- **LO16-4** Use multiplicative decomposition and moving averages to forecast a time series.
- **LO16-5** Use simple exponential smoothing to forecast a time series.
- **LO16-6** Use double exponential smoothing to forecast a time series.
- **LO16-7** Use multiplicative Winters' method to forecast a time series.
- **LO16-8** Appreciate some of the basic concepts of Box–Jenkins forecasting models.
- **LO16-9** Compare time series models by using forecast errors.
- **L16-10** Use index numbers to compare economic data over time.

Chapter Outline

Note: After completing Section 16.2, the reader may study Sections 16.3, 16.4, 16.6, 16.7, and 16.8 in any order without loss of continuity. Section 16.5 requires background from Sections 16.1, 16.2, and 16.4. Section 16.8 may be covered at any time.

A **time series** is a set of observations on a variable of interest that has been collected in **time order.** In this chapter we discuss developing and using **univariate time series models,** which forecast future values of a time series **solely on the basis of past values of the time series.** Often univariate time series models forecast future time series values by extrapolating the **trend** and/or **seasonal patterns** exhibited by the past values of the time series. To illustrate these ideas, we consider several cases in this chapter, including:

C

The Calculator Sales Case: By extrapolating an upward trend in past sales of the Bismark X-12 electronic calculator, Smith's Department Stores, Inc., forecasts future sales of this calculator. The forecasts help the department store chain to better implement its inventory and financial policies.

The Traveler's Rest Case: By extrapolating an upward trend and the seasonal behavior of its past hotel room occupancies, Traveler's Rest, Inc., forecasts future hotel room occupancies. The forecasts help the hotel chain to more effectively hire help and acquire supplies.

16.1 Time Series Components and Models

LO16-1 Identify the components of a time series.

In order to identify patterns in time series data, it is often convenient to think of such data as consisting of several components: **trend, cycle, seasonal variations,** and **irregular fluctuations. Trend** refers to the upward or downward movement that characterizes a time series over time. Thus trend reflects the long-run growth or decline in the time series. Trend movements can represent a variety of factors. For example, long-run movements in the sales of a particular industry might be determined by changes in consumer tastes, increases in total population, and increases in per capita income. **Cycle** refers to recurring up-and-down movements around trend levels. These fluctuations can last from 2 to 10 years or even longer measured from peak to peak or trough to trough. One of the common cyclical fluctuations found in time series data is the *business cycle,* which is represented by fluctuations in the time series caused by recurrent periods of prosperity and recession. **Seasonal variations** are periodic patterns in a time series that complete themselves within a calendar year or less and then are repeated on a regular basis. Often seasonal variations occur yearly. For example, soft drink sales and hotel room occupancies are annually higher in the summer months, while department store sales are annually higher during the winter holiday season. Seasonal variations can also last less than one year. For example, daily restaurant patronage might exhibit within-week seasonal variation, with daily patronage higher on Fridays and Saturdays. **Irregular fluctuations** are erratic time series movements that follow no recognizable or regular pattern. Such movements represent what is "left over" in a time series after trend, cycle, and seasonal variations have been accounted for.

Time series that exhibit trend, seasonal, and cyclical components are illustrated in Figure 16.1. In Figure 16.1(a) a time series of sales observations that has an essentially straight-line or linear trend is plotted. Figure 16.1(b) portrays a time series of sales observations that contains a

FIGURE 16.1 **Time Series Exhibiting Trend, Seasonal, and Cyclical Components**

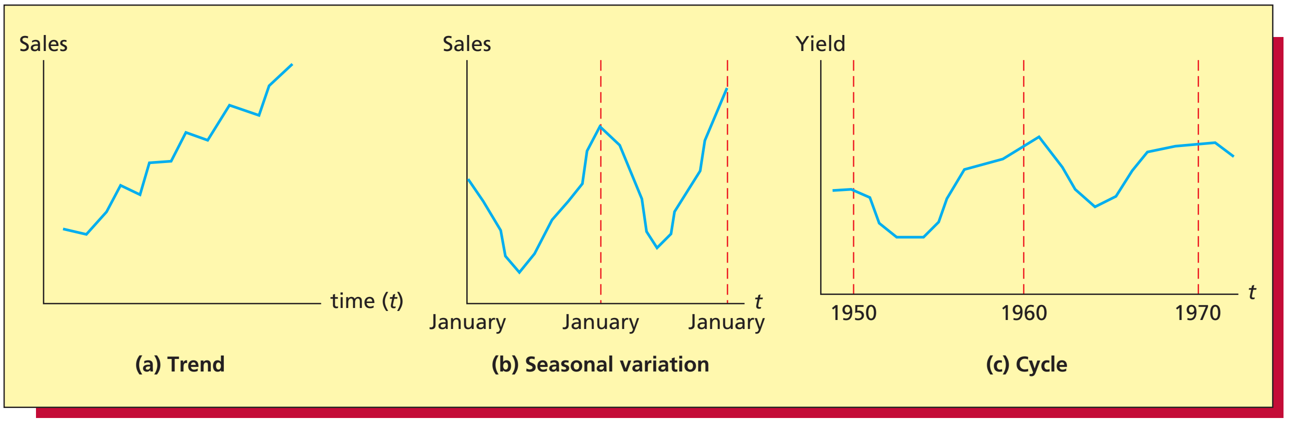

seasonal pattern that repeats annually. Figure 16.1(c) exhibits a time series of agricultural yields that is cyclical, repeating a cycle about once every 10 years.

Time series models attempt to identify significant patterns in the components of a time series. Then, assuming that these patterns will continue into the future, time series models extrapolate these patterns to forecast future time series values. In Section 16.2 we discuss forecasting by **time series regression models,** and in Section 16.3 we discuss forecasting by using an intuitive method called **multiplicative decomposition.** Both of these approaches assume that the time series components remain essentially constant over time. If the time series components might be changing slowly over time, it is appropriate to forecast by using **exponential smoothing.** This approach is discussed in Sections 16.4 and 16.5. If the time series components might be changing fairly quickly over time, it is appropriate to forecast by using the **Box–Jenkins methodology.** This more advanced approach is discussed in (Optional) Section 16.6.

LO16-2 Use time series regression to forecast time series having linear, quadratic, and certain types of seasonal patterns.

16.2 Time Series Regression

Modeling trend components We begin this section with two examples.

EXAMPLE 16.1 The Cod Catch Case: No Trend Regression

C

TABLE 16.1 Cod Catch (in Tons)
DS CodCatch

Month	Year 1	Year 2
Jan.	362	276
Feb.	381	334
Mar.	317	394
Apr.	297	334
May	399	384
June	402	314
July	375	344
Aug.	349	337
Sept.	386	345
Oct.	328	362
Nov.	389	314
Dec.	343	365

The Bay City Seafood Company owns a fleet of fishing trawlers and operates a fish processing plant. In order to forecast its minimum and maximum possible revenues from cod sales and plan the operations of its fish processing plant, the company desires to make both point forecasts and prediction interval forecasts of its monthly cod catch (measured in tons). The company has recorded monthly cod catch for the previous two years (years 1 and 2). The cod history is given in Table 16.1. A time series plot shows that the cod catches appear to randomly fluctuate around a constant average level. (See the plot in Figure 16.2.) Because the company subjectively believes that this data pattern will continue in the future, it seems reasonable to use the **"no trend"** regression model

$$y_t = \beta_0 + \varepsilon_t$$

to forecast cod catch in future months. It can be shown that for the no trend regression model the least squares point estimate b_0 of β_0 is $\bar{y}$, the average of the n observed time series values. Because the average $\bar{y}$ of the $n = 24$ observed cod catches is 351.29, it follows that $\hat{y}_t = b_0 = 351.29$ is the point prediction of the cod catch (y_t) in any future month. Furthermore, it can be shown that a $100(1 - \alpha)$ percent prediction interval for any future y_t value described by the no trend model is $[\hat{y}_t \pm t_{\alpha/2}\, s\sqrt{1 + (1/n)}]$. Here s is the sample standard deviation of the n observed time series values, and $t_{\alpha/2}$ is based on $n - 1$ degrees of freedom. For example, because s can be calculated to be

FIGURE 16.2 Plot of Cod Catch versus Time

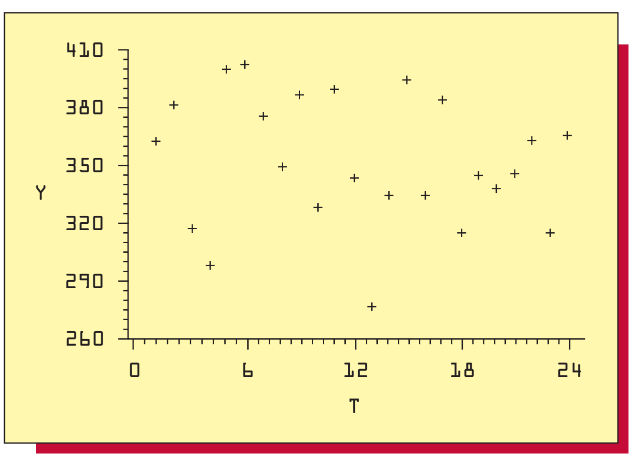

FIGURE 16.3 Plot of Calculator Sales versus Time

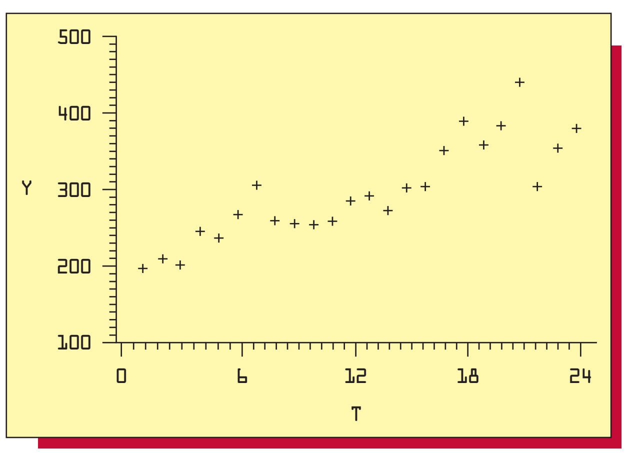

33.82 for the $n = 24$ cod catches, and because $t_{.025}$ based on $n - 1 = 23$ degrees of freedom is 2.069, it follows that a 95 percent prediction interval for the cod catch in any future month is $[351.29 \pm 2.069(33.82)\sqrt{1 + (1/24)}]$, or [279.92, 422.66].

EXAMPLE 16.2 The Calculator Sales Case: Inventory Policy

C

For the last two years Smith's Department Stores, Inc., has carried a new type of electronic calculator called the Bismark X-12. Sales of this calculator have generally increased over these two years. Smith's inventory policy attempts to ensure that stores will have enough Bismark X-12 calculators to meet practically all demand for the Bismark X-12, while at the same time ensuring that Smith's does not needlessly tie up its money by ordering many more calculators than can be sold. In order to implement this inventory policy in future months, Smith's requires both point predictions and prediction intervals for total monthly Bismark X-12 demand.

The monthly calculator demand data for the last two years are given in Table 16.2. A time series plot of the demand data is shown in Figure 16.3. The demands appear to randomly fluctuate around an average level that increases over time in a linear fashion. Furthermore, Smith's believes that this trend will continue for at least the next year. Thus it is reasonable to use the **"linear trend"** regression model

$$y_t = \beta_0 + \beta_1 t + \varepsilon_t$$

to forecast calculator sales in future months. Notice that this model is just a simple linear regression model where the time period t plays the role of the independent variable. The least squares point estimates of β_0 and β_1 can be calculated to be $b_0 = 198.0290$ and $b_1 = 8.0743$. [see Figure 16.4(a).] Therefore, for example, point forecasts of Bismark X-12 demand in January and February of year 3 (time periods 25 and 26) are, respectively,

$$\hat{y}_{25} = 198.0290 + 8.0743(25) = 399.9 \qquad \text{and}$$

$$\hat{y}_{26} = 198.0290 + 8.0743(26) = 408.0$$

Note that Figure 16.4(b) gives these point forecasts. In addition, it can be shown using either the formulas for simple linear regression or a computer software package [see Figure 16.4(c)] that a 95 percent prediction interval for demand in time period 25 is [328.6, 471.2] and that a 95 percent prediction interval for demand in time period 26 is [336.0, 479.9]. These prediction intervals can help Smith's implement its inventory policy. For instance, if Smith's stocks 471 Bismark X-12 calculators in January of year 3, we can be reasonably sure that monthly demand will be met.

BI

TABLE 16.2 Calculator Sales Data DS CalcSale

Month	Year 1	Year 2
Jan.	197	296
Feb.	211	276
Mar.	203	305
Apr.	247	308
May	239	356
June	269	393
July	308	363
Aug.	262	386
Sept.	258	443
Oct.	256	308
Nov.	261	358
Dec.	288	384

FIGURE 16.4 Excel Analysis of the Calculator Sales Data Using the Linear Trend Regression Model

(a) The Excel Output

ANOVA

	df	SS	MS	F	Significance F
Regression	1	74974.3567	74974.3567	74.7481	1.5893E-08
Residual	22	22066.6016	1003.0273		
Total	33	97040.9583			

	Coefficients	Standard Error	t Stat	P-value	Lower 95%	Upper 95%
Intercept	198.0290	13.3444	14.8398	6.0955E-13	170.3543	225.7036
T	8.0743	8.9339	8.6457	1.5893E-08	6.1375	10.0112

(b) Prediction using Excel

	A	B	C	D
24	358	23		
25	384	24		
26	399.8877	25	USING TREND	
27	407.962	26		

(c) Prediction using an Excel add-in (MegaStat)

Predicted values for: Sales

		95% Confidence Intervals		95% Prediction Intervals	
t	Predicted	lower	upper	lower	upper
25	399.9	372.2	427.6	328.6	471.2
26	408.0	378.6	437.3	336.0	479.9

Example 16.1 illustrates that the intercept β_0 can be used to model a lack of trend over time, and Example 16.2 illustrates that the expression $(\beta_0 + \beta_1 t)$ can model a linear trend over time. In addition, as will be illustrated in Exercise 16.42, the expression $(\beta_0 + \beta_1 t + \beta_2 t^2)$ can model a quadratic trend over time.

Modeling seasonal components We next consider how to forecast time series described by trend and seasonal components.

EXAMPLE 16.3 The Bike Sales Case: Inventory Policy

C

Table 16.3 presents quarterly sales of the TRK-50 mountain bike for the previous four years at a bicycle shop in Switzerland. The time series plot in Figure 16.5 shows that the bike sales exhibit a linear trend and a strong seasonal pattern, with bike sales being higher in the spring and summer quarters than in the winter and fall quarters. If we let y_t denote the number of TRK-50 mountain bikes sold in time period t at the Swiss bike shop, then a regression model describing y_t is

$$y_t = \beta_0 + \beta_1 t + \beta_{Q2}Q_2 + \beta_{Q3}Q_3 + \beta_{Q4}Q_4 + \varepsilon_t$$

Here the expression $(\beta_0 + \beta_1 t)$ models the linear trend evident in Figure 16.5. Q_2, Q_3, and Q_4 are dummy variables defined for quarters 2, 3, and 4. Specifically, Q_2 equals 1 if quarterly bike sales were observed in quarter 2 (spring) and 0 otherwise; Q_3 equals 1 if quarterly bike sales were observed in quarter 3 (summer) and 0 otherwise; Q_4 equals 1 if quarterly bike sales were observed in quarter 4 (fall) and 0 otherwise. Note that we have not defined a dummy variable for quarter 1 (winter). It follows that the regression parameters β_{Q2}, β_{Q3}, and β_{Q4} compare quarters 2, 3, and 4 with quarter 1. Intuitively, for example, β_{Q4} is the difference, excluding trend, between the level of the time series (y_t) in quarter 4 (fall) and the level of the time series in quarter 1 (winter). A positive β_{Q4} would imply that, excluding trend, bike sales in the fall can be expected to be higher than bike sales in the winter. A negative β_{Q4} would imply that, excluding trend, bike sales in the fall can be expected to be lower than bike sales in the winter.

Figure 16.6 gives the MINITAB output of a regression analysis of the quarterly bike sales by using the dummy variable model. The MINITAB output tells us that the linear trend and the seasonal dummy variables are significant (every t statistic has a related p-value less than .01). Also, notice that the least squares point estimates of β_{Q2}, β_{Q3}, and β_{Q4} are, respectively, $b_{Q2} = 21$, $b_{Q3} = 33.5$, and $b_{Q4} = 4.5$. It follows that, excluding trend, expected bike sales in quarter 2 (spring), quarter 3 (summer), and quarter 4 (fall) are estimated to be, respectively, 21, 33.5, and 4.5 bikes greater than expected bike sales in quarter 1 (winter). Furthermore, using all of the least squares point estimates in Figure 16.6, we can compute point forecasts of bike sales in quarters

TABLE 16.3 Quarterly Sales of the TRK-50 Mountain Bike DS BikeSales

Year	Quarter	t	Sales, y_t
1	1 (Winter)	1	10
	2 (Spring)	2	31
	3 (Summer)	3	43
	4 (Fall)	4	16
2	1	5	11
	2	6	33
	3	7	45
	4	8	17
3	1	9	13
	2	10	34
	3	11	48
	4	12	19
4	1	13	15
	2	14	37
	3	15	51
	4	16	21

FIGURE 16.5 Time Series Plot of TRK-50 Bike Sales

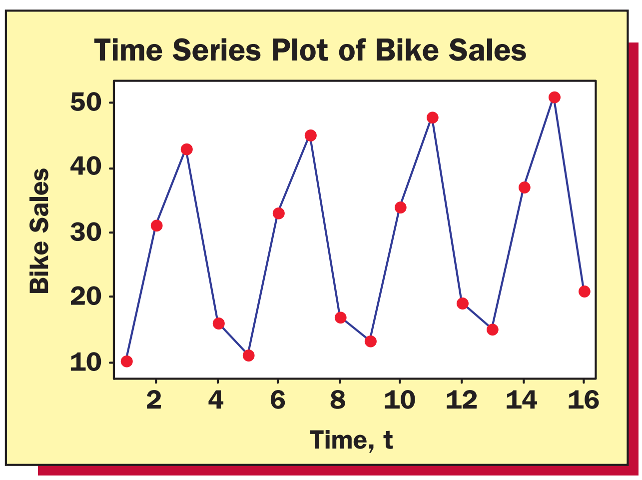

FIGURE 16.6 MINITAB Output of an Analysis of the Quarterly Bike Sales by Using Dummy Variable Regression

```
The regression equation is
BikeSales = 8.75 + 0.500 Time + 21.0 Q2 + 33.5 Q3 + 4.50 Q4

Predictor      Coef  SE Coef      T      P
Constant     8.7500   0.4281  20.44  0.000
Time        0.50000  0.03769  13.27  0.000
Q2          21.0000   0.4782  43.91  0.000
Q3          33.5000   0.4827  69.41  0.000
Q4           4.5000   0.4900   9.18  0.000

S = 0.674200   R-Sq = 99.8%   R-Sq(adj) = 99.8%

Values of Predictors for New Obs   Predicted Values for New Observations
New Obs  Time     Q2     Q3     Q4    New Obs     Fit  SE Fit       95% CI              95% PI
      1  17.0      0      0      0          1  17.250   0.506  (16.137, 18.363)  (15.395, 19.105)
      2  18.0      1      0      0          2  38.750   0.506  (37.637, 39.863)  (36.895, 40.605)
      3  19.0      0      1      0          3  51.750   0.506  (50.637, 52.863)  (49.895, 53.605)
      4  20.0      0      0      1          4  23.250   0.506  (22.137, 24.363)  (21.395, 25.105)
```

1 through 4 of next year (periods 17 through 20) as follows:

$$\hat{y}_{17} = b_0 + b_1(17) + b_{Q2}(0) + b_{Q3}(0) + b_{Q4}(0) = 8.75 + .5(17) = 17.250$$
$$\hat{y}_{18} = b_0 + b_1(18) + b_{Q2}(1) + b_{Q3}(0) + b_{Q4}(0) = 8.75 + .5(18) + 21 = 38.750$$
$$\hat{y}_{19} = b_0 + b_1(19) + b_{Q2}(0) + b_{Q3}(1) + b_{Q4}(0) = 8.75 + .5(19) + 33.5 = 51.750$$
$$\hat{y}_{20} = b_0 + b_1(20) + b_{Q2}(0) + b_{Q3}(0) + b_{Q4}(1) = 8.75 + .5(20) + 4.5 = 23.250$$

These point forecasts are given at the bottom of the MINITAB output, as are 95 percent prediction intervals for y_{17}, y_{18}, y_{19}, and y_{20}. The upper limits of these prediction intervals suggest that the bicycle shop can be reasonably sure that it will meet demand for the TRK-50 mountain bike if the numbers of bikes it stocks in quarters 1 through 4 are, respectively, 19, 41, 54, and 25 bikes.

BI

EXAMPLE 16.4 The Traveler's Rest Case: Predicting Hotel Room Occupancy

C

Table 16.4 presents a time series of hotel room occupancies observed by Traveler's Rest, Inc., a corporation that operates four hotels in a midwestern city. The analysts in the operating division of the corporation were asked to develop a model that could be used to obtain short-term forecasts (up to one year) of the number of occupied rooms in the hotels. These forecasts were needed by various personnel to assist in hiring additional help during the summer months, ordering materials that have long delivery lead times, budgeting of local advertising expenditures, and so on. The available historical data consisted of the number of occupied rooms during each day for the previous 14 years. Because it was desired to obtain monthly forecasts, these data were reduced to monthly averages by dividing each monthly total by the number of days in the month. The monthly room averages for the previous 14 years are the time series values given in Table 16.4. A time series plot of these values in Figure 16.7(a) shows that the monthly room averages follow a strong trend and have a seasonal pattern with one major and several minor peaks during the year. Note that the major peak each year occurs during the high summer travel months of June, July, and August.

Although the quarterly bike sales and monthly hotel room averages both exhibit seasonal variation, they exhibit different kinds of seasonal variation. The quarterly bike sales plotted in Figure 16.5 exhibit *constant seasonal variation.* In general, **constant seasonal variation** is seasonal variation where the magnitude of the seasonal swing does not depend on the level of the time series. On the other hand, **increasing seasonal variation** is seasonal variation where the magnitude of the seasonal swing increases as the level of the time series increases. Figure 16.7(a) shows that the monthly hotel room averages exhibit increasing seasonal variation. If a time series exhibits increasing seasonal variation, one approach is to first use a **fractional power transformation** (see Section 14.9) that produces a transformed time series that exhibits constant seasonal

TABLE 16.4 Monthly Hotel Room Averages DS TravRest

Year	Jan.	Feb.	Mar.	Apr.	May	June	July	Aug.	Sept.	Oct.	Nov.	Dec.
1	501	488	504	578	545	632	728	725	585	542	480	530
2	518	489	528	599	572	659	739	758	602	587	497	558
3	555	523	532	623	598	683	774	780	609	604	531	592
4	578	543	565	648	615	697	785	830	645	643	551	606
5	585	553	576	665	656	720	826	838	652	661	584	644
6	623	553	599	657	680	759	878	881	705	684	577	656
7	645	593	617	686	679	773	906	934	713	710	600	676
8	645	602	601	709	706	817	930	983	745	735	620	698
9	665	626	649	740	729	824	937	994	781	759	643	728
10	691	649	656	735	748	837	995	1040	809	793	692	763
11	723	655	658	761	768	885	1067	1038	812	790	692	782
12	758	709	715	788	794	893	1046	1075	812	822	714	802
13	748	731	748	827	788	937	1076	1125	840	864	717	813
14	811	732	745	844	833	935	1110	1124	868	860	762	877

FIGURE 16.7 Plots for the Monthly Hotel Room Averages in Table 16.4

(a) Plot of the monthly hotel room averages versus time

(b) Plot of the quartic roots of the monthly hotel averages versus time

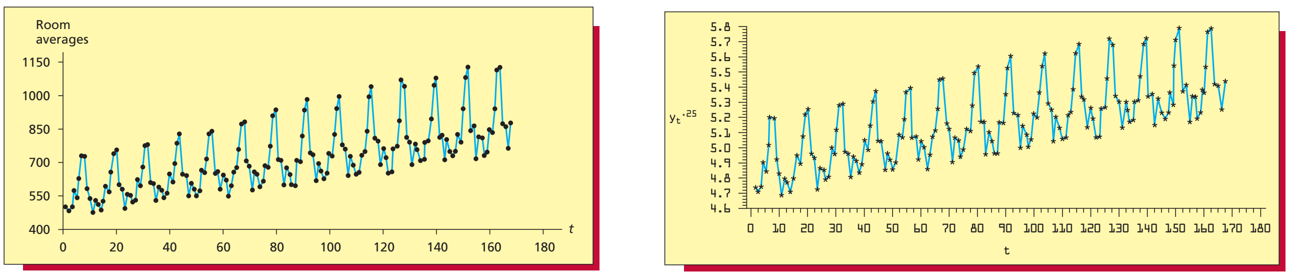

LO16-3 Use data transformations to forecast time series having increasing seasonal variation.

variation. Therefore, consider taking the square roots, quartic roots, and natural logarithms of the monthly hotel room averages in Table 16.4. If we do this and plot the resulting three sets of transformed values versus time, we find that the quartic root transformation best equalizes the seasonal variation. Figure 16.7(b) presents a plot of the quartic roots of the monthly hotel room averages versus time. Letting y_t denote the hotel room average observed in time period t, it follows that a regression model describing the quartic root of y_t is

$$y_t^{.25} = \beta_0 + \beta_1 t + \beta_{M1}M_1 + \beta_{M2}M_2 + \cdots + \beta_{M11}M_{11} + \varepsilon_t$$

The expression $(\beta_0 + \beta_1 t)$ models the linear trend evident in Figure 16.7(b). Furthermore, M_1, $M_2, \ldots, M_{11}$ are dummy variables defined for months January (month 1) through November (month 11). For example, M_1 equals 1 if a monthly room average was observed in January, and 0 otherwise; M_2 equals 1 if a monthly room average was observed in February, and 0 otherwise. Note that we have not defined a dummy variable for December (month 12). It follows that the regression parameters $\beta_{M1}, \beta_{M2}, \ldots, \beta_{M11}$ compare January through November with December. Intuitively, for example, β_{M1} is the difference, excluding trend, between the level of the time series ($y_t^{.25}$) in January and the level of the time series in December. A positive β_{M1} would imply that, excluding trend, the value of the time series in January can be expected to be greater than the value in December. A negative β_{M1} would imply that, excluding trend, the value of the time series in January can be expected to be smaller than the value in December.

Figure 16.8 gives relevant portions of the Excel output of a regression analysis of the hotel room data using the quartic root dummy variable model. The Excel output tells us that the linear trend and the seasonal dummy variables are significant (every t statistic has a related p-value less than .05). In addition, although not shown on the output, $R^2 = .988$. Now consider time period 169, which is January of next year and which therefore implies that $M_1 = 1$ and that all the other

FIGURE 16.8 Excel Output of an Analysis of the Quartic Roots of the Room Averages Using Dummy Variable Regression

(a) The Excel output

	Coefficients	Standard Error	t Stat	P-value
Intercept	4.8073	0.0085	568.0695	4.06E-259
t	0.0035	0.0000	79.0087	3.95E-127
M1	-0.0525	0.0106	-4.9709	1.75E-06
M2	-0.1408	0.0106	-13.3415	1.59E-27
M3	-0.1071	0.0106	-10.1509	7.016E-19
M4	0.0499	0.0105	4.7284	5.05E-06
M5	0.0254	0.0105	2.4096	0.0171
M6	0.1902	0.0105	18.0311	6.85E-40
M7	0.3825	0.0105	36.2663	1.28E-77
M8	0.4134	0.0105	39.2009	2.41E-82
M9	0.0714	0.0105	6.7731	2.47E-10
M10	0.0506	0.0105	4.8029	3.66E-06
M11	-0.1419	0.0105	-13.4626	7.47E-28

(b) Prediction of $y_t^{.25}$ using an Excel add-in (MegaStat)

		95% Prediction Intervals	
t	Predicted	lower	upper
169	5.3489	5.2913	5.4065
170	5.2641	5.2065	5.3217
171	5.3013	5.2437	5.3589
172	5.4618	5.4042	5.5194
173	5.4409	5.3833	5.4984
174	5.6091	5.5515	5.6667
175	5.8049	5.7473	5.8625
176	5.8394	5.7818	5.8969
177	5.5009	5.4433	5.5585
178	5.4837	5.4261	5.5412
179	5.2946	5.2370	5.3522
180	5.4400	5.3825	5.4976

dummy variables equal 0. Using the least squares point estimates in Figure 16.8(a), we compute a point forecast of $y_{169}^{.25}$ to be

$$b_0 + b_1(169) + b_{M1}(1) = 4.8073 + 0.0035(169) + (-.0525)(1) = 5.3489$$

Note that this point forecast is given in Figure 16.8(b) [see time period 169]. It follows that a point forecast of y_{169} is

$$(5.3489)^4 = 818.57$$

Furthermore, Figure 16.8(b) shows that a 95 percent prediction interval for $y_{169}^{.25}$ is [5.2913, 5.4065]. It follows that a 95 percent prediction interval for y_{169} is

$$[(5.2913)^4, (5.4065)^4] = [783.88, 854.41]$$

This interval says that Traveler's Rest, Inc., can be 95 percent confident that the monthly hotel room average in period 169 will be no less than 783.88 rooms per day and no more than 854.41 rooms per day. Lastly, note that Figure 16.8(b) also gives point forecasts of and 95 percent prediction intervals for the quartic roots of the hotel room averages in February through December of next year (time periods 170 through 180).

The validity of the regression methods just illustrated requires that the independence assumption be satisfied. However, when time series data are analyzed, this assumption is often violated. It is quite common for the time-ordered error terms to exhibit **positive or negative autocorrelation.** In Section 14.9 we discussed positive and negative autocorrelation, and we saw that we can use the **Durbin–Watson statistic *d*** to check for such autocorrelation. For example, it can be verified that this statistic shows no evidence of positive or negative first-order autocorrelation in the error terms of the calculator sales model or in the error terms of the bike sales model. However, the Durbin–Watson statistic for the dummy variable regression model describing the quartic roots of the hotel room averages can be calculated to be $d = 1.26$. Because the dummy variable regression model uses $k = 12$ independent variables, and because Tables A.11, A.12, and A.13 (pages 801 and 802) do not give the **Durbin–Watson critical points** corresponding to $k = 12$, we cannot test for autocorrelation using these tables. However, it can be shown that $d = 1.26$ is quite small and indicates **positive autocorrelation** in the error terms. One approach to dealing with autocorrelation in the error terms is to predict a future error term ε_t by using an **autoregressive model** that relates ε_t to past error terms $\varepsilon_{t-1}, \varepsilon_{t-2}, \ldots$. A better way (in the authors' opinion) is to use the Box–Jenkins methodology to forecast the time series. This methodology is presented in Section 16.6.

Exercises for Section 16.2

CONCEPTS

16.1 Discuss how we model no trend and a linear trend.

16.2 Discuss the difference between constant seasonal variation and increasing seasonal variation.

TABLE 16.5 Annual Total U.S. Lumber Production DS LumberProd

35,404	35,733
37,462	35,791
32,901	34,592
33,178	38,902
34,449	37,858
38,044	32,926
36,762	35,697
36,742	34,548
33,385	32,087
34,171	37,515
36,124	38,629
38,658	32,019
32,901	35,710
36,356	36,693
37,166	37,153

TABLE 16.6 Watch Sales Values DS WatchSale

298	356
302	371
301	399
351	392
336	425
361	411
407	455
351	457
357	465
346	481

METHODS AND APPLICATIONS

16.3 THE LUMBER PRODUCTION CASE DS LumberProd

In this exercise we consider annual U.S. lumber production over 30 years. The data were obtained from the U.S. Department of Commerce *Survey of Current Business* and are presented in Table 16.5. (The lumber production values are given in millions of board feet.) Plot the lumber production values versus time and discuss why the plot indicates that the model $y_t = \beta_0 + \varepsilon_t$ might appropriately describe these values.

16.4 THE LUMBER PRODUCTION CASE DS LumberProd

Referring to the situation of Exercise 16.3, the mean and the standard deviation of the lumber production values can be calculated to be $\bar{y} = 35{,}651.9$ and $s = 2{,}037.3599$. Find a point forecast of and a 95 percent prediction interval for any future lumber production value.

16.5 THE WATCH SALES CASE DS WatchSale

The past 20 monthly sales figures for a new type of watch sold at Lambert's Discount Stores are given in Table 16.6.

a Plot the watch sales values versus time and discuss why the plot indicates that the model

$$y_t = \beta_0 + \beta_1 t + \varepsilon_t$$

might appropriately describe these values.

b The least squares point estimates of β_0 and β_1 can be calculated to be $b_0 = 290.0895$ and $b_1 = 8.6677$. Use b_0 and b_1 to show (calculate) that a point forecast of watch sales in period 21 is $\hat{y}_{21} = 472.1$. Use the formulas of simple linear regression analysis or a computer software package to show that a 95 percent prediction interval for watch sales in period 21 is [421.5, 522.7].

16.6 THE AIR CONDITIONER SALES CASE DS ACSales

Bargain Department Stores, Inc., is a chain of department stores in the Midwest. Quarterly sales of the "Bargain 8000-Btu Air Conditioner" over the past three years are as given in the lefthand portion of Table 16.7.

a Plot sales versus time and discuss why the plot indicates that the model

$$y_t = \beta_0 + \beta_1 t + \beta_2 t^2 + \beta_{Q2} Q_2 + \beta_{Q3} Q_3 + \beta_{Q4} Q_4 + \varepsilon_t$$

might appropriately describe the sales values. In this model Q_2, Q_3, and Q_4 are appropriately defined dummy variables for quarters 2, 3, and 4.

TABLE 16.7 Air Conditioner Sales and A Dummy Variable Regression Analysis of the Sales Data DS ACSales

Year	Quarter	Sales
1	1	2,915
	2	8,032
	3	10,411
	4	2,427
2	1	4,381
	2	9,138
	3	11,386
	4	3,382
3	1	5,105
	2	9,894
	3	12,300
	4	4,013

```
The regression equation is
Sales = 2625 + 383 T - 11.4 TSq + 4630 Q2 + 6739 Q3 - 1565 Q4

Predictor      Coef   SE Coef      T      P
Constant     2624.5    100.4   26.15  0.000    S = 92.4244
T            382.82    34.03   11.25  0.000    R-Sq = 100.0%
TSq         -11.354    2.541   -4.47  0.004    R-Sq(adj)= 99.9%
Q2          4629.74    76.08   60.86  0.000
Q3          6738.85    77.38   87.09  0.000
Q4         -1565.32    79.34  -19.73  0.000

Time     Fit   SE Fit         95% CI                95% PI
  13   5682.4   112.6  ( 5406.9,  5957.9)  ( 5325.9,  6038.8)
  14  10388.4   142.8  (10039.0, 10737.8)  ( 9972.2, 10804.6)
  15  12551.0   177.2  (12117.4, 12984.7)  (12061.9, 13040.2)
  16   4277.7   213.9  ( 3754.4,  4801.1)  ( 3707.6,  4847.8)
```

b The right-hand portion of Table 16.7 is the MINITAB output of a regression analysis of the air conditioner sales data using the model in part *a*. **(1)** Define the dummy variables Q_2, Q_3, and Q_4. **(2)** Use the MINITAB output to find, report, and interpret the least squares point estimates of β_{Q2}, β_{Q3}, and β_{Q4}.

c At the bottom of the MINITAB output are point and prediction interval forecasts of air conditioner sales in the four quarters of year 4. Find and report these forecasts and hand calculate the point forecasts.

16.7 Table 16.8 gives the monthly international passenger totals over the last 11 years for an airline company. A plot of these passenger totals reveals an upward trend with increasing seasonal variation, and the natural logarithmic transformation is found to best equalize the seasonal variation [see Figure 16.9(a) and (b)]. Figure 16.9(c) gives the MINITAB output of a regression analysis of the monthly international passenger totals by using the model

$$\ln y_t = \beta_0 + \beta_1 t + \beta_{M1}M_1 + \beta_{M2}M_2 + \cdots + \beta_{M11}M_{11} + \varepsilon_t$$

Here $M_1, M_2, \ldots, M_{11}$ are appropriately defined dummy variables for January (month 1) through November (month 11). Let y_{133} denote the international passenger totals in month 133 (January of next year). The MINITAB output tells us that a point forecast of and a 95 percent prediction interval for $\ln y_{133}$ are, respectively, 6.08610 and [5.96593, 6.20627]. **(1)** Using the least squares point estimates on the MINITAB output, calculate the point forecast. **(2)** By calculating $e^{6.08610}$ and $[e^{5.96593}, e^{6.20627}]$, find a point forecast of and a 95 percent prediction interval for y_{133}. DS AirPass

16.8 Use the Durbin–Watson statistic given at the bottom of the MINITAB output in Figure 16.9(c) to assess whether there is positive autocorrelation.

TABLE 16.8 Monthly International Passenger Totals (Thousands of Passengers) DS AirPass

Year	Jan.	Feb.	Mar.	Apr.	May	June	July	Aug.	Sept.	Oct.	Nov.	Dec.
1	112	118	132	129	121	135	148	148	136	119	104	118
2	115	126	141	135	125	149	170	170	158	133	114	140
3	145	150	178	163	172	178	199	199	184	162	146	166
4	171	180	193	181	183	218	230	242	209	191	172	194
5	196	196	236	235	229	243	264	272	237	211	180	201
6	204	188	235	227	234	264	302	293	259	229	203	229
7	242	233	267	269	270	315	364	347	312	274	237	278
8	284	277	317	313	318	374	413	405	355	306	271	306
9	315	301	356	348	355	422	465	467	404	347	305	336
10	340	318	362	348	363	435	491	505	404	359	310	337
11	360	342	406	396	420	472	548	559	463	407	362	405

Source: *FAA Statistical Handbook of Civil Aviation* (several annual issues). These data were originally presented by Box and Jenkins (1976). We have updated the situation in this exercise to be more modern.

FIGURE 16.9 Analysis of the Monthly International Passenger Totals

(a) Plot of the passenger totals **(b) Plot of the natural logarithms of the passenger totals**

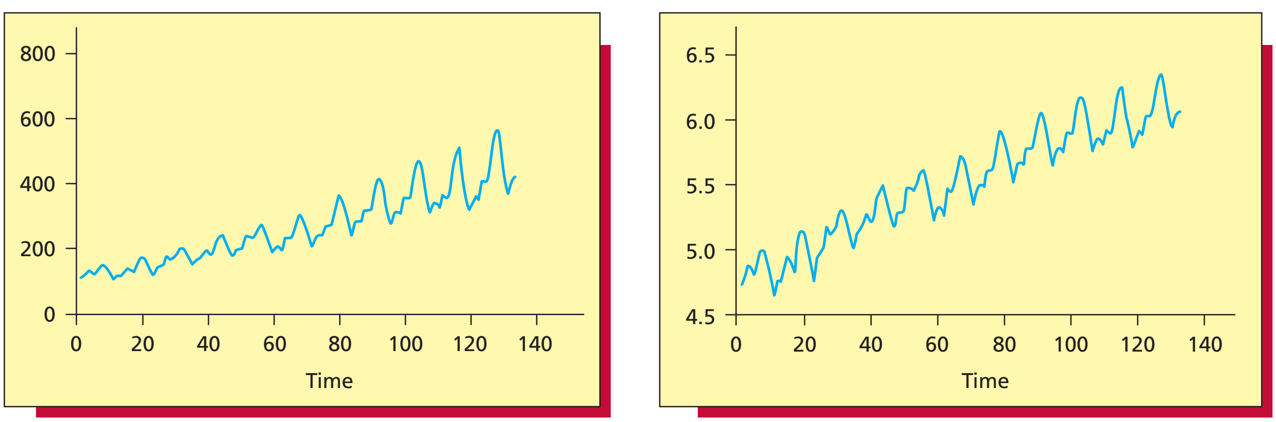

FIGURE 16.9 Analysis of the Monthly International Passenger Totals (*continued*)

(c) MINITAB Output of a Regression Analysis of the Monthly International Passenger Totals Using the Dummy Variable Model

```
Predictor        Coef     SE Coef       T      P    Predicted Values for New Observations
Constant      4.69618     0.01973  238.02  0.000    Time      Fit    SE Fit         95% PI
Time        0.0103075   0.0001316   78.30  0.000     133  6.08610  0.01973  (5.96593, 6.20627)
Jan           0.01903     0.02451    0.78  0.439     134  6.07888  0.01973  (5.95871, 6.19905)
Feb           0.00150     0.02451    0.06  0.951     135  6.22564  0.01973  (6.10547, 6.34581)
March         0.13795     0.02450    5.63  0.000     136  6.19383  0.01973  (6.07366, 6.31400)
April         0.09583     0.02449    3.91  0.000     137  6.20008  0.01973  (6.07991, 6.32025)
May           0.09178     0.02449    3.75  0.000     138  6.33292  0.01973  (6.21276, 6.45309)
June          0.21432     0.02448    8.75  0.000     139  6.44360  0.01973  (6.32343, 6.56377)
July          0.31469     0.02448   12.85  0.000     140  6.44682  0.01973  (6.32665, 6.56699)
Aug           0.30759     0.02448   12.57  0.000     141  6.31605  0.01973  (6.19588, 6.43622)
Sept          0.16652     0.02448    6.80  0.000     142  6.18515  0.01973  (6.06498, 6.30531)
Oct           0.02531     0.02447    1.03  0.303     143  6.05455  0.01973  (5.93438, 6.17472)
Nov          -0.11559     0.02447   -4.72  0.000     144  6.18045  0.01973  (6.06028, 6.30062)

S = 0.0573917   R-Sq = 98.3%   R-Sq(adj) = 98.1%    Durbin-Watson statistic = 0.420944
```

LO16-4 Use multiplicative decomposition and moving averages to forecast a time series.

16.3 Multiplicative Decomposition

When a time series exhibits increasing (or decreasing) seasonal variation, we can use the **multiplicative decomposition method** to decompose the time series into its **trend, seasonal, cyclical,** and **irregular** components. This is illustrated in the following example.

EXAMPLE 16.5 The Tasty Cola Case: Predicting Soft Drink Sales

C

The Discount Soda Shop, Inc., owns and operates 10 drive-in soft drink stores. Discount Soda has been selling Tasty Cola, a soft drink introduced just three years ago and gaining in popularity. Periodically, Discount Soda orders Tasty Cola from the regional distributor. To better implement its inventory policy, Discount Soda needs to forecast monthly Tasty Cola sales (in hundreds of cases).

Discount Soda has recorded monthly Tasty Cola sales for the previous three years. This time series is given in Table 16.9 and plotted in Figure 16.10. Notice that, in addition to having a linear trend, the Tasty Cola sales time series possesses seasonal variation, with sales of the soft drink

TABLE 16.9 Monthly Sales of Tasty Cola (in Hundreds of Cases) DS TastyCola

Year	Month	t	Sales, y_t	Year	Month	t	Sales, y_t
1	1 (Jan.)	1	189	2	7	19	831
	2 (Feb.)	2	229		8	20	960
	3 (Mar.)	3	249		9	21	1,152
	4 (Apr.)	4	289		10	22	759
	5 (May)	5	260		11	23	607
	6 (June)	6	431		12	24	371
	7 (July)	7	660	3	1	25	298
	8 (Aug.)	8	777		2	26	378
	9 (Sept.)	9	915		3	27	373
	10 (Oct.)	10	613		4	28	443
	11 (Nov.)	11	485		5	29	374
	12 (Dec.)	12	277		6	30	660
2	1	13	244		7	31	1,004
	2	14	296		8	32	1,153
	3	15	319		9	33	1,388
	4	16	370		10	34	904
	5	17	313		11	35	715
	6	18	556		12	36	441

FIGURE 16.10 Time Series Plot of the Tasty Cola Sales Data

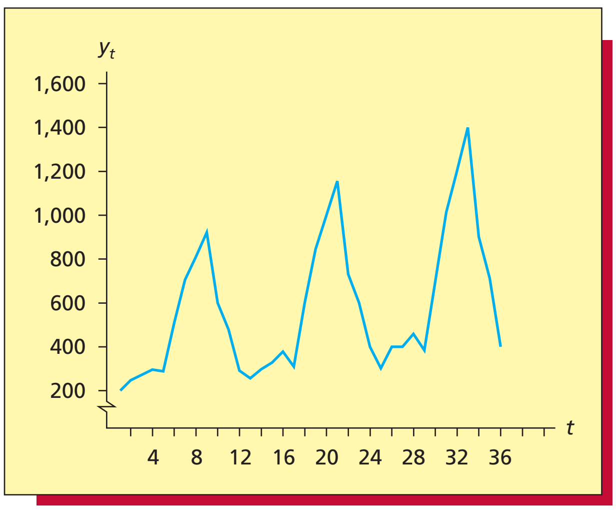

greatest in the summer and early fall months and lowest in the winter months. Because the seasonal variation seems to be increasing, we will see as we progress through this example that it might be reasonable to conclude that y_t, the sales of Tasty Cola in period t, is described by the **multiplicative model**

$$y_t = TR_t \times SN_t \times CL_t \times IR_t$$

Here TR_t, SN_t, CL_t, and IR_t represent the trend, seasonal, cyclical, and irregular components of the time series in time period t.

Table 16.10 summarizes the calculations needed to find estimates—denoted tr_t, sn_t, cl_t, and ir_t—of TR_t, SN_t, CL_t, and IR_t. As shown in the table, we begin by calculating **moving averages** and **centered moving averages.** The purpose behind computing these averages is to eliminate seasonal variations and irregular fluctuations from the data. The first moving average of the first 12 Tasty Cola sales values is

$$\frac{189 + 229 + 249 + 289 + 260 + 431 + 660 + 777 + 915 + 613 + 485 + 277}{12}$$
$$= 447.833$$

TABLE 16.10 Tasty Cola Sales and the Multiplicative Decomposition Method

t Time Period	y_t Tasty Cola Sales	First Step: 12-Period Moving Average	$tr_t \times cl_t$: Centered Moving Average	$sn_t \times ir_t$: $\frac{y_t}{tr_t \times cl_t}$	sn_t: Table 16.11	d_t: $\frac{y_t}{sn_t}$	tr_t: $380.163 + 9.489t$	$tr_t \times sn_t$: Multiply tr_t by sn_t	$cl_t \times ir_t$: $\frac{y_t}{tr_t \times sn_t}$	cl_t: 3-Period Moving Average	ir_t: $\frac{cl_t \times ir_t}{cl_t}$
1 (Jan)	189				.493	383.37	389.652	192.10	.9839		
2	229				.596	384.23	399.141	237.89	.9626	.9902	.9721
3	249				.595	418.49	408.630	243.13	1.0241	1.0010	1.0231
4	289				.680	425	418.119	284.32	1.0165	1.0396	.9778
5	260				.564	460.99	427.608	241.17	1.0781	1.0315	1.0452
6	431	447.833			.986	437.12	437.097	430.98	1.0000	1.0285	.9723
7	660	452.417	450.125	1.466	1.467	449.9	446.586	655.14	1.0074	1.0046	1.0028
8	777	458	455.2085	1.707	1.693	458.95	456.075	772.13	1.0063	1.0004	1.0059
9	915	563.833	460.9165	1.985	1.990	459.79	465.564	926.47	.9876	.9937	.9939
10	613	470.583	467.208	1.312	1.307	469.01	475.053	620.89	.9873	.9825	1.0049
11	485	475	472.7915	1.026	1.029	471.33	489.542	498.59	.9727	.9648	1.0082
12	277	485.417	480.2085	.577	.600	461.67	494.031	296.42	.9345	.9634	.9700
13 (Jan)	244	499.667	492.542	.495	.493	494.97	503.520	248.24	.9829	.9618	1.0219
14	296	514.917	507.292	.583	.596	496.64	513.009	305.75	.9681	.9924	.9755
15	319	534.667	524.792	.608	.595	536.13	522.498	310.89	1.0261	1.0057	1.0203
16	370	546.833	540.75	.684	.680	544.12	531.987	361.75	1.0228	1.0246	.9982
17	313	557	551.9165	.567	.564	554.97	541.476	305.39	1.0249	1.0237	1.0012
18	556	564.833	560.9165	.991	.986	563.89	550.965	543.25	1.0235	1.0197	1.0037
19	831	569.333	567.083	1.465	1.467	566.46	560.454	822.19	1.0107	1.0097	1.0010
20	960	576.167	572.75	1.676	1.693	567.04	569.943	964.91	.9949	1.0016	.9933
21	1,152	580.667	578.417	1.992	1.990	578.89	579.432	1,153.07	.9991	.9934	1.0057
22	759	586.75	583.7085	1.300	1.307	580.72	588.921	769.72	.9861	.9903	.9958
23	607	591.833	589.2915	1.030	1.029	589.89	598.410	615.76	.9858	.9964	.9894
24	371	600.5	596.1665	.622	.600	618.33	607.899	364.74	1.0172	.9940	1.0233
25 (Jan)	298	614.917	607.7085	.490	.493	604.46	617.388	304.37	.9791	1.0027	.9765
26	378	631	622.9585	.607	.596	634.23	626.877	373.62	1.0117	.9920	1.0199
27	373	650.667	640.8335	.582	.595	626.89	636.366	378.64	.9851	1.0018	.9833
28	443	662.75	656.7085	.675	.680	651.47	645.855	439.18	1.0087	1.0030	1.0057
29	374	671.75	667.25	.561	.564	663.12	655.344	369.61	1.0119	1.0091	1.0028
30	660	677.583	674.6665	.978	.986	669.37	664.833	655.53	1.0068	1.0112	.9956
31	1,004				1.467	684.39	674.322	989.23	1.0149	1.0059	1.0089
32	1,153				1.693	681.04	683.811	1,157.69	.9959	1.0053	.9906
33	1,388				1.990	697.49	693.300	1,379.67	1.0060	.9954	1.0106
34	904				1.307	691.66	702.789	918.55	.9842	.9886	.9955
35	715				1.029	694.85	712.278	732.93	.9755	.9927	.9827
36	441				.600	735	721.707	433.06	1.0183		

Here we use a "12-period moving average" because the Tasty Cola time series data are monthly (12 time periods or "seasons" per year). If the data were quarterly, we would compute a "4-period moving average." The second moving average is obtained by dropping the first sales value (y_1) from the average and by including the next sales value (y_{13}) in the average. Thus we obtain

$$\frac{229 + 249 + 289 + 260 + 431 + 660 + 777 + 915 + 613 + 485 + 277 + 244}{12}$$
$$= 452.417$$

The third moving average is obtained by dropping y_2 from the average and by including y_{14} in the average. We obtain

$$\frac{249 + 289 + 260 + 431 + 660 + 777 + 915 + 613 + 485 + 277 + 244 + 296}{12}$$
$$= 458$$

Successive moving averages are computed similarly until we include y_{36} in the last moving average. Note that we use the term *moving average* here because, as we calculate these averages, we move along by dropping the most remote observation in the previous average and by including the "next" observation in the new average.

The first moving average corresponds to a time that is midway between periods 6 and 7, the second moving average corresponds to a time that is midway between periods 7 and 8, and so forth. In order to obtain averages corresponding to time periods in the original Tasty Cola time series, we calculate **centered moving averages.** The centered moving averages are two-period moving averages of the previously computed 12-period moving averages. Thus the first centered moving average is

$$\frac{447.833 + 452.417}{2} = 450.125$$

The second centered moving average is

$$\frac{452.417 + 458}{2} = 455.2085$$

Successive centered moving averages are calculated similarly. The 12-period moving averages and centered moving averages for the Tasty Cola sales time series are given in Table 16.10.

If the original moving averages had been computed using an odd number of time series values, the centering procedure would not have been necessary. For example, if we had three seasons per year, we would compute three-period moving averages. Then, the first moving average would correspond to period 2, the second moving average would correspond to period 3, and so on. However, most seasonal time series are quarterly, monthly, or weekly, so the centering procedure is necessary.

The centered moving average in time period t is considered to equal $tr_t \times cl_t$, the estimate of $TR_t \times CL_t$, because the averaging procedure is assumed to have removed seasonal variations (note that each moving average is computed using exactly one observation from each season) and (short-term) irregular fluctuations. The (longer-term) trend effects and cyclical effects—that is, $tr_t \times cl_t$—remain.

Because the model

$$y_t = TR_t \times SN_t \times CL_t \times IR_t$$

implies that

$$SN_t \times IR_t = \frac{y_t}{TR_t \times CL_t}$$

it follows that the estimate $sn_t \times ir_t$ of $SN_t \times IR_t$ is

$$sn_t \times ir_t = \frac{y_t}{tr_t \times cl_t}$$

TABLE 16.11 Estimation of the Seasonal Factors

		$sn_t \times ir_t = y_t/(tr_t \times cl_t)$ Year 1	Year 2	$\overline{sn}_t$	$sn_t = 1.0008758(\overline{sn}_t)$
1	Jan.	.495	.490	.4925	.493
2	Feb.	.583	.607	.595	.596
3	Mar.	.608	.582	.595	.595
4	Apr.	.684	.675	.6795	.680
5	May	.567	.561	.564	.564
6	June	.991	.978	.9845	.986
7	July	1.466	1.465	1.4655	1.467
8	Aug.	1.707	1.676	1.6915	1.693
9	Sep.	1.985	1.992	1.9885	1.990
10	Oct.	1.312	1.300	1.306	1.307
11	Nov.	1.026	1.030	1.028	1.029
12	Dec.	.577	.622	.5995	.600

FIGURE 16.11 Plot of Tasty Cola Sales and Deseasonalized Sales

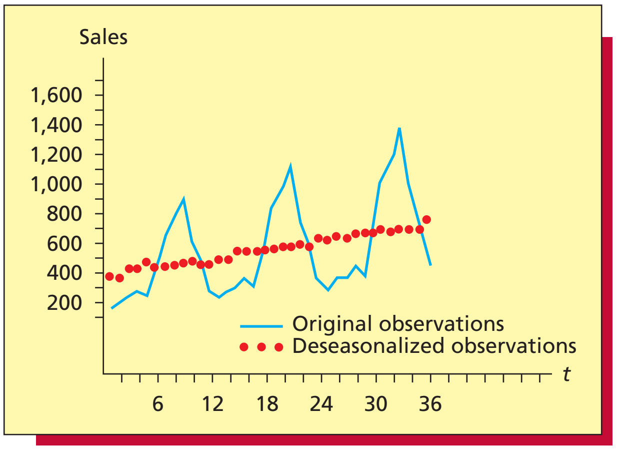

Noting that the values of $sn_t \times ir_t$ are calculated in Table 16.10, we can find sn_t by grouping the values of $sn_t \times ir_t$ by months and calculating an average, $\overline{sn}_t$, for each month. These monthly averages are given for the Tasty Cola data in Table 16.11. The monthly averages are then normalized so that they sum to the number of time periods in a year. Denoting the number of time periods in a year by L (for instance, $L = 4$ for quarterly data, $L = 12$ for monthly data), we accomplish the normalization by multiplying each value of $\overline{sn}_t$ by the quantity

$$\frac{L}{\sum \overline{sn}_t} = \frac{12}{.4925 + .595 + \cdots + .5995}$$

$$= \frac{12}{11.9895} = 1.0008758$$

This normalization process results in the estimate $sn_t = 1.0008758(\overline{sn}_t)$, which is the estimate of SN_t. These calculations are summarized in Table 16.11.

Having calculated the values of sn_t and placed them in Table 16.10, we next define the **deseasonalized observation** in time period t to be

$$d_t = \frac{y_t}{sn_t}$$

Deseasonalized observations are computed to better estimate the trend component TR_t. Dividing y_t by the estimated seasonal factor removes the seasonality from the data and allows us to better understand the nature of the trend. The deseasonalized observations are calculated in Table 16.10 and are plotted in Figure 16.11. Because the deseasonalized observations have a straight-line appearance, it seems reasonable to assume a linear trend

$$TR_t = \beta_0 + \beta_1 t$$

We estimate TR_t by fitting a straight line to the deseasonalized observations. That is, we compute the least squares point estimates of the parameters in the simple linear regression model relating the dependent variable d_t to the independent variable t:

$$d_t = \beta_0 + \beta_1 t + \varepsilon_t$$

We obtain $b_0 = 380.163$ and $b_1 = 9.489$. It follows that the estimate of TR_t is

$$tr_t = b_0 + b_1 t = 380.163 + 9.489t$$

The values of tr_t are calculated in Table 16.10. Note that, for example, although $y_{22} = 759$, Tasty Cola sales in period 22 (October of year 2), are larger than $tr_{22} = 588.921$ (the estimated trend in period 22), $d_{22} = 580.72$ is smaller than $tr_{22} = 588.921$. This implies that, on a deseasonalized basis, Tasty Cola sales were slightly down in October of year 2. This might have been caused by a slightly colder October than usual.

Thus far, we have found estimates sn_t and tr_t of SN_t and TR_t. Because the model

$$y_t = TR_t \times SN_t \times CL_t \times IR_t$$

implies that

$$CL_t \times IR_t = \frac{y_t}{TR_t \times SN_t}$$

it follows that the estimate of $CL_t \times IR_t$ is

$$cl_t \times ir_t = \frac{y_t}{tr_t \times sn_t}$$

Moreover, experience has shown that, when considering either monthly or quarterly data, we can average out ir_t and thus calculate the estimate cl_t of CL_t by computing a three-period moving average of the $cl_t \times ir_t$ values.

Finally, we calculate the estimate ir_t of IR_t by using the equation

$$ir_t = \frac{cl_t \times ir_t}{cl_t}$$

The calculations of the values cl_t and ir_t for the Tasty Cola data are summarized in Table 16.10. Because there are only three years of data, and because most of the values of cl_t are near 1, we cannot discern a well-defined cycle. Furthermore, examining the values of ir_t, we cannot detect a pattern in the estimates of the irregular factors.

Traditionally, the estimates tr_t, sn_t, cl_t, and ir_t obtained by using the multiplicative decomposition method are used to describe the time series. However, we can also use these estimates to forecast future values of the time series. If there is no pattern in the irregular component, we predict IR_t to equal 1. Therefore, the point forecast of y_t is

$$\hat{y}_t = tr_t \times sn_t \times cl_t$$

if a well-defined cycle exists and can be predicted. The point forecast is

$$\hat{y}_t = tr_t \times sn_t$$

if a well-defined cycle does not exist or if CL_t cannot be predicted, as in the Tasty Cola example. Because values of $tr_t \times sn_t$ have been calculated in column 9 of Table 16.10, these values are the point forecasts of the $n = 36$ historical Tasty Cola sales values. Furthermore, we present in Table 16.12 point forecasts of future Tasty Cola sales in the 12 months of year 4. Recalling that

TABLE 16.12 Forecasts of Future Values of Tasty Cola Sales Calculated Using the Multiplicative Decomposition Method

t	sn_t	$tr_t = 380.163 + 9.489t$	Point Prediction, $\hat{y}_t = tr_t \times sn_t$	Approximate 95% Prediction Interval	y_t
37	.493	731.273	360.52	[333.72, 387.32]	352
38	.596	740.762	441.48	[414.56, 468.40]	445
39	.595	750.252	446.40	[419.36, 473.44]	453
40	.680	759.741	516.62	[489.45, 543.79]	541
41	.564	769.231	433.85	[406.55, 461.15]	457
42	.986	778.720	767.82	[740.38, 795.26]	762
43	1.467	788.209	1,156.30	[1,128.71, 1,183.89]	1,194
44	1.693	797.699	1,350.50	[1,322.76, 1,378.24]	1,361
45	1.990	807.188	1,606.30	[1,578.41, 1,634.19]	1,615
46	1.307	816.678	1,067.40	[1,039.35, 1,095.45]	1,059
47	1.029	826.167	850.12	[821.90, 878.34]	824
48	.600	835.657	501.39	[473, 529.78]	495

the estimated trend equation is $tr_t = 380.163 + 9.489t$ and that the estimated seasonal factor for August is 1.693 (see Table 16.11), it follows, for example, that the point forecast of Tasty Cola sales in period 44 (August of year 4) is

$$\begin{aligned}\hat{y}_{44} &= tr_{44} \times sn_{44} \\ &= (380.163 + 9.489(44))(1.693) \\ &= 797.699(1.693) \\ &= 1{,}350.50\end{aligned}$$

Although there is no theoretically correct prediction interval for y_t, a fairly accurate **approximate 100(1 − α) percent prediction interval for y_t** is obtained by computing an interval that is centered at $\hat{y}_t$ and that has a length equal to the length of the 100(1 − α) percent prediction interval for the **deseasonalized observation d_t.** Here the interval for d_t is obtained by using the model

$$\begin{aligned}d_t &= TR_t + \varepsilon_t \\ &= \beta_0 + \beta_1 t + \varepsilon_t\end{aligned}$$

For instance, using MINITAB to predict d_t on the basis of this model, we find that a 95 percent prediction interval for d_{44} is [769.959, 825.439]. Because this interval has a **half-length** equal to (825.439 − 769.959)/2 = 55.48/2 = 27.74, it follows that an approximate 95 percent prediction interval for y_{44} is

$$\begin{aligned}[\hat{y}_{44} \pm 27.74] &= [1{,}350.50 \pm 27.74] \\ &= [1{,}322.76, 1{,}378.24]\end{aligned}$$

In Table 16.12 we give the approximate 95 percent prediction intervals (calculated by the above method) for Tasty Cola sales in the 12 months of year 4.

Next, suppose we actually observe Tasty Cola sales in year 4, and these sales are as given in Table 16.12. In Figure 16.12 we plot the observed and forecasted sales for all 48 sales periods. In practice, the comparison of the observed and forecasted sales in years 1 through 3 would be used by the analyst to determine whether the forecasting equation adequately fits the historical data. An adequate fit (as indicated by Figure 16.12, for example) might prompt an analyst to use this equation to calculate forecasts for future time periods. One reason that the Tasty Cola forecasting equation

$$\begin{aligned}\hat{y}_t &= tr_t \times sn_t \\ &= (380.163 + 9.489t)sn_t\end{aligned}$$

FIGURE 16.12 A Plot of the Observed and Forecasted Tasty Cola Sales Values

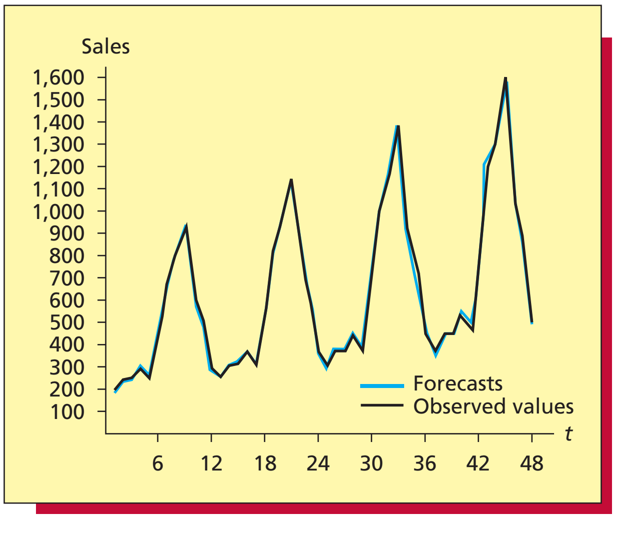

provides reasonable forecasts is that this equation *multiplies* tr_t by sn_t. Therefore, as the average level of the time series (determined by the trend) increases, the seasonal swing of the time series increases, which is consistent with the data plots in Figures 16.10 and 16.12. For example, note from Table 16.11 that the estimated seasonal factor for August is 1.693. The forecasting equation yields a prediction of Tasty Cola sales in August of year 1 equal to

$$\begin{aligned}\hat{y}_8 &= [380.163 + 9.489(8)]1.693\\ &= (456.075)(1.693)\\ &= 772.13\end{aligned}$$

This implies a seasonal swing of $772.13 - 456.075 = 316.055$ (hundreds of cases) above 456.075, the estimated trend level. The forecasting equation yields a prediction of Tasty Cola sales in August of year 2 equal to

$$\begin{aligned}\hat{y}_{20} &= [380.163 + 9.489(20)]1.693\\ &= (569.943)(1.693)\\ &= 964.91\end{aligned}$$

which implies an increased seasonal swing of $964.91 - 569.943 = 394.967$ (hundreds of cases) above 569.943, the estimated trend level. In general, then, the forecasting equation is appropriate for forecasting a time series with a seasonal swing that is proportional to the average level of the time series as determined by the trend—that is, a time series exhibiting increasing seasonal variation.

We next note that the U.S. Bureau of the Census has developed the **Census II method,** which is a sophisticated version of the multiplicative decomposition method discussed in this section. The initial version of Census II was primarily developed by Julius Shiskin in the late 1950s when a computer program was written to perform the rather complex calculations. Several modifications have been made to the first version of the method over the years. Census II continues to be widely used by a variety of businesses and government agencies.

Census II first adjusts the original data for "trading day variations." That is, the data are adjusted to account for the fact that, for example, different months or quarters will consist of different numbers of business days or "trading days." The method then uses an iterative procedure to obtain estimates of the seasonal component (SN_t), the trading day component, the so-called trend-cycle component ($TR_t \times CL_t$), and the irregular component (IR_t). The iterative procedure makes extensive use of moving averages and a method for identifying and replacing extreme values in order to eliminate randomness. For a good explanation of the details involved here and in the Census II method as a whole, see Makridakis, Wheelwright, and McGee (1983). After carrying out a number of tests to check the correctness of the estimates, the method estimates the trend-cycle, seasonal, and irregular components.

MINITAB carries out a modified version of the multiplicative decomposition method discussed in this section. We believe that MINITAB's modified version (at the time of the writing of this book) makes some conceptual errors that can result in biased estimates of the time series components. Therefore, we will not present MINITAB output of multiplicative decomposition. The Excel add-in (MegaStat) estimates the seasonal factors and the trend line exactly as described in this section. MegaStat does not estimate the cyclical and irregular components. However, because it is often reasonable to make forecasts by using estimates of the seasonal factors and trend line, MegaStat can be used to do this. In Appendix 16.2, we show a MegaStat output that estimates the seasonal factors and trend line for the Tasty Cola data.

Exercises for Section 16.3

CONCEPTS

connect™

16.9 Explain how the multiplicative decomposition model estimates seasonal factors.

16.10 Explain how the multiplicative decomposition method estimates the trend effect.

16.11 Discuss how the multiplicative decomposition method makes point forecasts of future time series values.

METHODS AND APPLICATIONS

Exercises 16.12 through 16.16 are based on the following situation: International Machinery, Inc., produces a tractor and wishes to use **quarterly** tractor sales data observed in the last four years to predict quarterly tractor sales next year. The following MegaStat output gives the tractor sales data and the estimates of the seasonal factors and trend line for the data: DS IntMach

t	Year	Quarter	Sales, y	Centered Moving Average	Ratio to CMA	Seasonal Indexes	Sales, y Deseasonalized
1	1	1	293			1.191	245.9
2	1	2	392			1.521	257.7
3	1	3	221	275.125	0.803	0.804	275.0
4	1	4	147	302.000	0.487	0.484	303.9
5	2	1	388	325.250	1.193	1.191	325.7
6	2	2	512	338.125	1.514	1.521	336.6
7	2	3	287	354.125	0.810	0.804	357.1
8	2	4	184	381.500	0.482	0.484	380.4
9	3	1	479	405.000	1.183	1.191	402.0
10	3	2	640	417.375	1.533	1.521	420.7
11	3	3	347	435.000	0.798	0.804	431.8
12	3	4	223	462.125	0.483	0.484	461.0
13	4	1	581	484.375	1.199	1.191	487.7
14	4	2	755	497.625	1.517	1.521	496.3
15	4	3	410			0.804	510.2
16	4	4	266			0.484	549.9

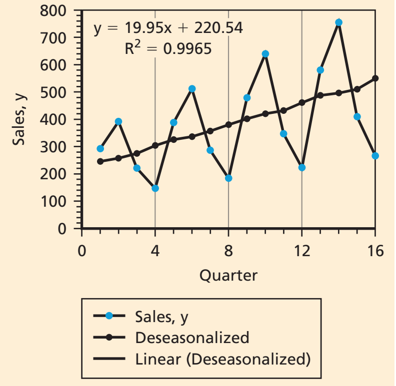

Calculation of Seasonal Indexes

	1	2	3	4	
1			0.803	0.487	
2	1.193	1.514	0.810	0.482	
3	1.183	1.533	0.798	0.483	
4	1.199	1.517			
mean:	1.192	1.522	0.804	0.484	4.001
adjusted:	1.191	1.521	0.804	0.484	4.000

16.12 Find and identify the four seasonal factors for quarters 1, 2, 3, and 4.

16.13 What type of trend is indicated by the plot of the deseasonalized data?

16.14 What is the equation of the estimated trend that has been calculated using the deseasonalized data?

16.15 Compute a point forecast of tractor sales (based on trend and seasonal factors) for each of the quarters next year.

16.16 Compute an approximate 95 percent prediction interval forecast of tractor sales for each of the quarters next year. Use the fact that the half-lengths of 95 percent prediction intervals for the deseasonalized sales values in the four quarters of next year are, respectively, 14, 14.4, 14.6, and 15.

16.17 If we use the multiplicative decomposition method to analyze the quarterly bicycle sales data given in Table 16.3 (page 634), we find that the quarterly seasonal factors are .46, 1.22, 1.68, and .64. Furthermore, if we use a statistical software package to fit a straight line to the deseasonalized sales values, we find that the estimate of the trend is $tr_t = 22.61 + .59t$.

In addition, we find that the half-lengths of 95 percent prediction intervals for the deseasonalized sales values in the four quarters of the next year are, respectively, 2.80, 2.85, 2.92, and 2.98. DS BikeSales

a Calculate point predictions of bicycle sales in the four quarters of the next year.

b Calculate approximate 95 percent prediction intervals for bicycle sales in the four quarters of the next year.

16.4 Simple Exponential Smoothing

LO16-5 Use simple exponential smoothing to forecast a time series.

In ongoing forecasting systems, forecasts of future time series values are made each period for succeeding periods. At the end of each period the estimates of the time series parameters and the forecasting equation need to be updated to account for the most recent observation. This updating accounts for possible changes in the parameters that may occur over time. In addition, such changes may imply that unequal weights should be applied to the time series observations when the estimates of the parameters are updated.

In this section we assume that a time series is appropriately described by the no trend equation

$$y_t = \beta_0 + \varepsilon_t$$

When the parameter β_0 remains constant over time, we have seen that it is reasonable to forecast future values of y_t by using regression analysis (see Example 16.1 on page 632). In such a case, the least squares point estimate of β_0 is

$$b_0 = \bar{y} = \text{the average of the observed time series values}$$

When we compute the point estimate b_0 we are **equally weighting** each of the previously observed time series values $y_1, y_2, \ldots, y_n$.

When the value of the parameter β_0 is slowly changing over time, the equal weighting scheme may not be appropriate. Instead, it may be desirable to weight recent observations more heavily than remote observations. **Simple exponential smoothing** is a forecasting method that applies unequal weights to the time series observations. This unequal weighting is accomplished by using a **smoothing constant** that determines how much weight is attached to each observation. The most recent observation is given the most weight. More distantly past observations are given successively smaller weights. The procedure allows the forecaster to update the estimate of β_0 so that changes in the value of this parameter can be detected and incorporated into the forecasting equation. We illustrate simple exponential smoothing in the following example.

EXAMPLE 16.6 The Cod Catch Case: Simple Exponential Smoothing

C

Consider the cod catch data of Example 16.1, which are given in Table 16.1 (page 632). The plot of these data (in Figure 16.2 on page 632) suggests that the no trend model

$$y_t = \beta_0 + \varepsilon_t$$

may appropriately describe the cod catch series. It is also possible that the parameter β_0 could be slowly changing over time.

We begin the simple exponential smoothing procedure by calculating an initial estimate of the average level β_0 of the series. This estimate is denoted S_0 and is computed by averaging the first six time series values. We obtain

$$S_0 = \frac{\sum_{t=1}^{6} y_t}{6} = \frac{362 + 381 + \cdots + 402}{6} = 359.667$$

Note that, because simple exponential smoothing attempts to track changes over time in the average level β_0 by using newly observed values to update the estimates of β_0, we use only six of the $n = 24$ time series observations to calculate the initial estimate of β_0. If we do this, then 18 observations remain to tell us how β_0 may be changing over time. Experience has shown that, in general, it is reasonable to calculate initial estimates in exponential smoothing procedures by using half of the historical data. However, it can be shown that, in simple exponential smoothing, using six observations is reasonable (it would not, however, be reasonable to use a very small number of observations because doing so might make the initial estimate so different from the true value of β_0 that the exponential smoothing procedure would be adversely affected).

Next, assume that at the end of time period $T - 1$ we have an estimate S_{T-1} of β_0. Then, assuming that in time period T we obtain a new observation y_T, we can update S_{T-1} to S_T, which is an estimate made in period T of β_0. We compute the updated estimate by using the so-called **smoothing equation**

$$S_T = \alpha y_T + (1 - \alpha)S_{T-1}$$

Here α is a smoothing constant between 0 and 1 (α will be discussed in more detail later). The updating equation says that S_T, the estimate made in time period T of β_0, equals a fraction α (for example, .1) of the newly observed time series observation y_T plus a fraction $(1 - \alpha)$ (for example, .9) of S_{T-1}, the estimate made in time period $T - 1$ of β_0. The more the average level of the process is changing, the more a newly observed time series value should influence our estimate, and thus the larger the smoothing constant α should be set. We will soon see how to use historical data to determine an appropriate value of α.

TABLE 16.13 One-Period-Ahead Forecasting of the Historical Cod Catch Time Series Using Simple Exponential Smoothing with $\alpha = .02$

Year	Month	Actual Cod Catch, y_T	Smoothed Estimate, S_T ($S_0 = 359.667$)	Forecast Made Last Period	Forecast Error
1	Jan.	362	359.713	359.667	2.333
	Feb.	381	360.139	359.713	21.287
	Mar.	317	359.276	360.139	−43.139
	Apr.	297	358.031	359.276	−62.276
	May	399	358.850	358.031	40.969
	June	402	359.713	358.850	43.150
	July	375	360.019	359.713	15.287
	Aug.	349	359.799	360.019	−11.019
	Sept.	386	360.323	359.799	26.201
	Oct.	328	359.676	360.323	−32.323
	Nov.	389	360.263	359.676	29.324
	Dec.	343	359.917	360.263	−17.263
2	Jan.	276	358.239	359.917	−83.917
	Feb.	334	357.754	358.239	−24.239
	Mar.	394	358.479	357.754	36.246
	Apr.	334	357.990	358.479	−24.479
	May	384	358.510	357.990	26.010
	June	314	357.620	358.510	−44.510
	July	344	357.347	357.620	−13.620
	Aug.	337	356.940	357.347	−20.347
	Sept.	345	356.701	356.940	−11.940
	Oct.	362	356.807	356.701	5.299
	Nov.	314	355.951	356.807	−42.807
	Dec.	365	356.132	355.951	9.049

We will now begin with the initial estimate $S_0 = 359.667$ and update this initial estimate by applying the smoothing equation to the 24 observed cod catches. To do this, we arbitrarily set α equal to .02, and to judge the appropriateness of this choice of α we calculate "one-period-ahead" forecasts of the historical cod catches as we carry out the smoothing procedure. Because the initial estimate of β_0 is $S_0 = 359.667$, it follows that 359.667 is the forecast made at time 0 for y_1, the value of the time series in period 1. Because we see from Table 16.13 that $y_1 = 362$, we have a forecast error of $362 - 359.667 = 2.333$. Using $y_1 = 362$, we can update S_0 to S_1, an estimate made in period 1 of the average level of the time series, by using the equation

$$\begin{aligned} S_1 &= \alpha y_1 + (1-\alpha)S_0 \\ &= .02(362) + .98(359.667) = 359.713 \end{aligned}$$

Because this implies that 359.713 is the forecast made in period 1 for y_2, and because we see from Table 16.13 that $y_2 = 381$, we have a forecast error of $381 - 359.713 = 21.287$. Using $y_2 = 381$, we can update S_1 to S_2, an estimate made in period 2 of β_0, by using the equation

$$\begin{aligned} S_2 &= \alpha y_2 + (1-\alpha)S_1 \\ &= .02(381) + .98(359.713) = 360.139 \end{aligned}$$

Because this implies that 360.139 is the forecast made in period 2 for y_3, and because we see from Table 16.13 that $y_3 = 317$, we have a forecast error of $317 - 360.139 = -43.139$. This procedure is continued through the entire 24 periods of historical data. The results are summarized in Table 16.13. Using the results in this table, we find that, for $\alpha = .02$, the mean of the squared forecast errors is 1161.14. To find a "good" value of α, we evaluate the mean of the squared forecast errors for values of α ranging from .02 to .30 in increments of .02. (In most exponential smoothing applications, the value of the smoothing constant used is between .01 and .30.) When we do this, we find that $\alpha = .02$ minimizes the mean of the squared forecast errors. Because this minimizing value of α is small, it appears to be best to apply small weights to new observations, which tells us that the level of the time series is not changing very much.

In general, simple exponential smoothing is carried out as follows:

Simple Exponential Smoothing

1 Suppose that the time series $y_1, \ldots, y_n$ is described by the equation

$$y_t = \beta_0 + \varepsilon_t$$

where the average level β_0 of the process may be slowly changing over time. Then the estimate S_T of β_0 made in time period T is given by the **smoothing equation**

$$S_T = \alpha y_T + (1 - \alpha)S_{T-1}$$

where α is a smoothing constant between 0 and 1 and S_{T-1} is the estimate of β_0 made in time period $T-1$.

2 A point forecast made in time period T for any future value of the time series is S_T.

3 If we observe y_{T+1} in time period $T+1$, we can update S_T to S_{T+1} by using the equation

$$S_{T+1} = \alpha y_{T+1} + (1 - \alpha)S_T$$

and a point forecast made in time period $T+1$ for any future value of the time series is S_{T+1}.

EXAMPLE 16.7 The Cod Catch Case: Forecasting

C

In Example 16.6 we saw that $\alpha = .02$ is a "good" value of the smoothing constant when forecasting the 24 observed cod catches in Table 16.13. Therefore, we will use simple exponential smoothing with $\alpha = .02$ to forecast future monthly cod catches. From Table 16.13 we see that $S_{24} = 356.132$ is the estimate made in month 24 of the average level β_0 of the monthly cod catches. It follows that the point forecast made in month 24 of any future monthly cod catch is 356.132 tons of cod. Now, assuming that we observe a cod catch in January of year 3 of $y_{25} =$ 384, we can update S_{24} to S_{25} by using the equation

$$\begin{aligned} S_{25} &= \alpha y_{25} + (1 - \alpha)S_{24} \\ &= .02(384) + .98(356.132) \\ &= 356.689 \end{aligned}$$

This implies that the point forecast made in month 25 of any future monthly cod catch is 356.689 tons of cod.

By using the smoothing equation

$$S_T = \alpha y_T + (1 - \alpha)S_{T-1}$$

it can be shown that S_T, the estimate made in time period T of the average level β_0 of the time series, can be expressed as

$$\begin{aligned} S_T = {} & \alpha y_T + \alpha(1 - \alpha)y_{T-1} + \alpha(1 - \alpha)^2 y_{T-2} \\ & + \cdots + \alpha(1 - \alpha)^{T-1}y_1 + (1 - \alpha)^T S_0 \end{aligned}$$

The coefficients measuring the contributions of the observations $y_T, y_{T-1}, y_{T-2}, \ldots, y_1$—that is, $\alpha, \alpha(1 - \alpha), \alpha(1 - \alpha)^2, \ldots, \alpha(1 - \alpha)^{T-1}$—decrease *exponentially* with age. For this reason we refer to this procedure as simple *exponential* smoothing.

Because the coefficients measuring the contributions of $y_T, y_{T-1}, y_{T-2}, \ldots, y_1$ are decreasing exponentially, the most recent observation y_T makes the largest contribution to the current estimate of β_0. Older observations make smaller and smaller contributions to this estimate. Thus remote observations are "dampened out" of the current estimate of β_0 as time advances. The rate at which remote observations are dampened out depends on the smoothing constant α. For values of α near 1, remote observations are dampened out quickly. For example, if $\alpha = .9$ we obtain coefficients .9, .09, .009, .0009, ... For values of α near 0, remote observations are dampened out more slowly (if $\alpha = .1$, we obtain coefficients .1, .09, .081, .0729, . . .). The choice of a smoothing constant α is usually made by simulated forecasting of a historical data set as illustrated in Example 16.6.

Computer software packages can be used to implement exponential smoothing. These packages choose the smoothing constant (or constants) in different ways and also compute approximate

FIGURE 16.13 MINITAB Output of Using Simple Exponential Smoothing to Forecast the Cod Catches

(a) The graphical forecasts

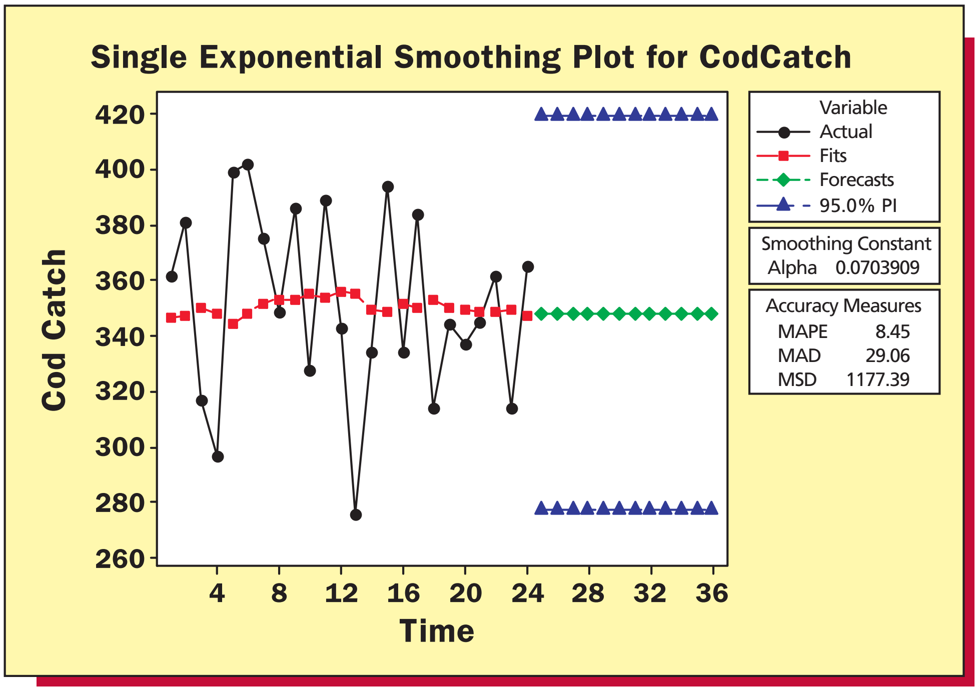

(b) The numerical forecasts of the cod catch in month 25 (and any other future month)

Forecasts

Period	Forecast	Lower	Upper
25	348.168	276.976	419.360

prediction intervals in different ways. Optimally, the user should carefully investigate how the computer software package implements exponential smoothing. At a minimum, the user should not trust the forecasts given by the software package if they seem illogical.

Figure 16.13 gives the MINITAB output of using simple exponential smoothing to forecast in month 24 the cod catches in future months. Note that MINITAB has selected the smoothing constant $\alpha = .0703909$ and tells us that the point forecast and the 95 percent prediction interval forecast of the cod catch in any future month are, respectively, 348.168 and [276.976, 419.360]. Looking at Figure 16.13(a), these forecasts seem intuitively reasonable. An Excel add-in (MegaStat) output of simple exponential smoothing for the cod catch data is given in Appendix 16.1.

Exercises for Section 16.4

CONCEPTS

connect

16.18 In general, when is it appropriate to use exponential smoothing?

16.19 What is the purpose of a smoothing constant in exponential smoothing?

METHODS AND APPLICATIONS

16.20 THE COD CATCH CASE DS CodCatch

Consider Table 16.13 (page 649). Verify (calculate) that S_3, an estimate made in period 3 of β_0, is 359.276. Also verify (calculate) that the one-period-ahead forecast error for period 4 is -62.276, as shown in Table 16.13.

FIGURE 16.14 MINITAB Output of Using Simple Exponential Smoothing to Forecast Lumber Production

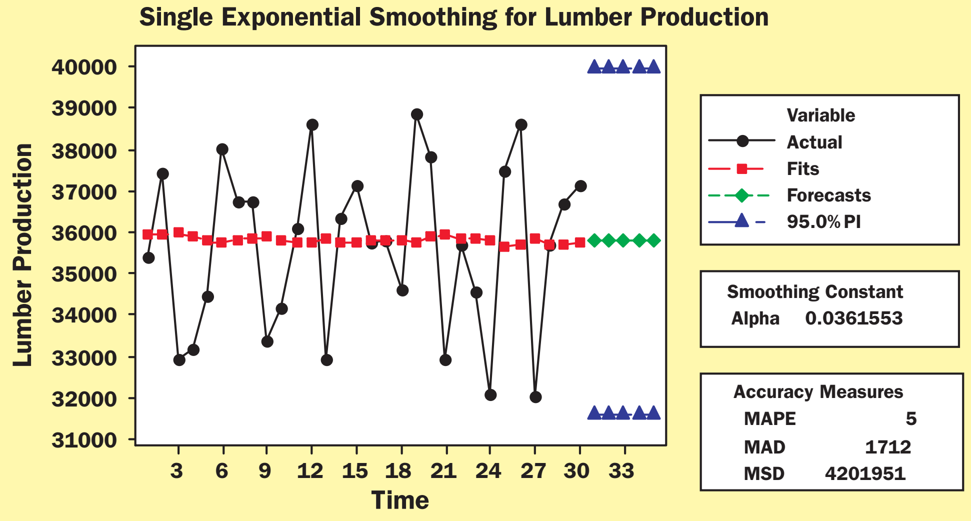

Period	Forecast	Lower	Upper
31	35782.6	31588.9	39976.3

16.21 THE COD CATCH CASE DS CodCatch

Consider Example 16.7 (page 650). Suppose that we observe a cod catch in February of year 3 of $y_{26} = 328$. Update $S_{25} = 356.689$ to S_{26}, a point forecast made in month 26 of any future monthly cod catch. Use $\alpha = .02$ as in Example 16.7.

16.22 THE LUMBER PRODUCTION CASE DS LumberProd

Figure 16.14 gives the MINITAB output of using simple exponential smoothing to forecast yearly U.S. lumber production. Here MINITAB has estimated the smoothing constant alpha to be .0361553. Use the MINITAB output to find and report the point prediction of and the 95 percent prediction interval for the total U.S. lumber production in a future year.

LO16-6 Use double exponential smoothing to forecast a time series.

16.5 Holt–Winters' Models

Holt–Winters' double exponential smoothing Various extensions of simple exponential smoothing can be used to forecast time series that are described by models that are different from the model $y_t = \beta_0 + \varepsilon_t$. For example, **Holt–Winters' double exponential smoothing** can be used to forecast time series that are described by the linear trend model

$$y_t = \beta_0 + \beta_1 t + \varepsilon_t$$

Here we assume that β_0 and β_1 (and thus the linear trend) may be changing slowly over time. To implement Holt–Winters' double exponential smoothing, we let ℓ_{T-1} denote the estimate of the level $\beta_0 + \beta_1(T - 1)$ of the time series in time period $T - 1$, and we let b_{T-1} denote the estimate of the slope β_1 of the time series in time period $T - 1$. Then, if we observe a new time series value y_T in time period T, the estimate of the level $\beta_0 + \beta_1 T$ of the time series in time period T uses the **smoothing constant α** and is

$$\ell_T = \alpha y_T + (1 - \alpha)\,[\ell_{T-1} + b_{T-1}]$$

This equation says that ℓ_T equals a fraction α of the newly observed time series value y_T plus a fraction $(1 - \alpha)$ of $[\ell_{T-1} + b_{T-1}]$, which is the estimate of the level of the time series in time period T, as calculated using the estimates ℓ_{T-1} and b_{T-1} computed in time period $T - 1$. Furthermore, the estimate of the slope β_1 of the time series in time period T uses the **smoothing constant γ** and is

$$b_T = \gamma[\ell_T - \ell_{T-1}] + (1 - \gamma)b_{T-1}$$

This equation says that b_T equals a fraction γ of $[\ell_T - \ell_{T-1}]$, which is an estimate of the difference between the levels of the time series in periods T and $T - 1$, plus a fraction $(1 - \gamma)$ of b_{T-1}, the estimate of the slope made in time period $T - 1$.

To use the updating equations, we first obtain initial estimates ℓ_0 and b_0 of the level and the slope of the time series in time period 0. One way to do this is to fit a least squares trend line to part (say, one-half) of the historical data and let the y-intercept and slope of the trend line be ℓ_0 and b_0. For example, consider the 24 observed calculator sales values in Table 16.2 (page 633). If we fit a least squares trend line to the first 12 of those values, we obtain

$$\hat{y}_t = 204.803 + 6.9406t$$

This would imply that $\ell_0 = 204.803$ and $b_0 = 6.9406$. MINITAB uses a more complicated method to find initial estimates and obtains $\ell_0 = 198.0290$ and $b_0 = 8.0743$. Starting with the MINITAB initial estimates ℓ_0 and b_0, we calculate a point forecast of y_1 from time origin 0 to be

$$\hat{y}_1(0) = \ell_0 + b_0 = 198.0290 + 8.0743 = 206.103$$

This point forecast is shown on the MINITAB output of Figure 16.15(a) [it is the first number under the column headed $\hat{y}_T(T - 1)$]. Also shown on the output are the actual calculator sales value $y_1 = 197$ and the forecast error, which is

$$y_1 - \hat{y}_1(0) = 197 - 206.103 = -9.103$$

FIGURE 16.15 The MINITAB Output of Double Exponential Smoothing for the Calculator Sales Data

(a) The updated level and slope estimates when $\alpha = .2$ and $\gamma = .2$

$\ell_0 = 198.0290$ $b_0 = 8.0743$

Time T	Sales y_T	Level ℓ_T	Slope b_T	Forecast $\hat{y}_T(T-1)$	Error $y_T - \hat{y}_T(T-1)$
1	197	204.283	7.7102	206.103	-9.1033
2	211	211.794	7.6705	211.993	-0.9929
3	203	216.172	7.0119	219.465	-16.4648
4	247	227.947	7.9646	223.184	23.8162
5	239	236.529	8.0881	235.912	3.0884
6	269	249.494	9.0634	244.617	24.3827
7	308	268.446	11.0411	258.557	49.4427
8	262	275.990	10.3416	279.487	-17.4869
9	258	280.665	9.2084	286.331	-28.3312
10	256	283.099	7.8535	289.873	-33.8733
11	261	284.962	6.6554	290.952	-29.9521
12	288	290.894	6.5107	291.617	-3.6171
13	296	297.123	6.4545	297.404	-1.4043
14	276	298.062	5.3514	303.578	-27.5780
15	305	303.731	5.4148	303.414	1.5862
16	308	308.917	5.3690	309.146	-1.1459
17	356	322.629	7.0376	314.286	41.7143
18	393	342.333	9.5709	329.666	63.3339
19	363	354.123	10.0148	351.904	11.0962
20	386	368.510	10.8893	364.138	21.8621
21	443	392.120	13.4333	379.400	63.6004
22	308	386.042	9.5312	405.553	-97.5529
23	358	388.059	8.0282	395.574	-37.5735
24	384	393.670	7.5447	396.087	-12.0870

(b) Point and 95 percent prediction interval forecasts when $\alpha = .2$ and $\gamma = .2$

Period	Forecast	Lower	Upper
25	401.214	337.813	464.616
26	408.759	344.037	473.482
27	416.304	350.159	482.449
28	423.849	356.186	491.512
29	431.393	362.124	500.663
30	438.938	367.979	509.898

(c) Point and 95 percent prediction interval forecasts when $\alpha = .496$ and $\gamma = .142$

Period	Forecast	Lower	Upper
25	383.677	319.135	448.220
26	389.121	316.066	462.177
27	394.565	312.109	477.022
28	400.010	307.534	492.486
29	405.454	302.521	508.386
30	410.898	297.191	524.604

We next choose values of the smoothing constants α and γ. A reasonable choice (and the default option of MINITAB) is to let each of α and γ be .2. Then, using $y_1 = 197$ and the equation for ℓ_T, it follows that the estimate of the level of the time series in time period 1 is

$$\begin{aligned}\ell_1 &= \alpha y_1 + (1 - \alpha)[\ell_0 + b_0] \\ &= .2(197) + .8[198.0290 + 8.0743] \\ &= 204.283\end{aligned}$$

Furthermore, using the equation for b_T, the estimate of the slope of the time series in time period 1 is

$$\begin{aligned}b_1 &= \gamma[\ell_1 - \ell_0] + (1 - \gamma)b_0 \\ &= .2[204.283 - 198.0290] + .8(8.0743) \\ &= 7.7102\end{aligned}$$

It follows that a point forecast made in time period 1 of y_2 is

$$\hat{y}_2(1) = \ell_1 + b_1 = 204.283 + 7.7102 = 211.993$$

Because the actual calculator sales value in period 2 is $y_2 = 211$, the forecast error is

$$y_2 - \hat{y}_2(1) = 211 - 211.933 = -.993$$

The MINITAB output in Figure 16.15(a) on the previous page shows the entire process of using the double exponential smoothing updating equations to find new period-by-period estimates of the level and slope of the time series. The output also shows the one-period-ahead forecasts and forecast errors, which are utilized to evaluate the effectiveness of the double exponential smoothing procedure. At the end of the updating process, MINITAB uses $\ell_{24} = 393.670$ and $b_{24} = 7.5447$ to calculate point forecasts of future calculator sales values. For example, point forecasts of y_{25} and y_{26} made from time origin 24 are

$$\hat{y}_{25}(24) = \ell_{24} + b_{24} = 393.670 + 7.5447 = 401.214$$

and

$$\hat{y}_{26}(24) = \ell_{24} + 2b_{24} = 393.670 + 2(7.5447) = 408.759$$

These point forecasts, as well as point forecasts of y_{27} through y_{30}, are shown on the MINITAB output in Figure 16.15(b). Also shown are 95 percent prediction interval forecasts of y_{25} through y_{30}.

Figure 16.16 shows a MINITAB output that graphically illustrates the forecasts when $\alpha = .2$ and $\gamma = .2$. Generally speaking, choosing $\alpha = .2$ and $\gamma = .2$ gives reasonable results, but MINITAB will choose its own values of α and γ. If we have MINITAB do this, it chooses $\alpha = .496$ and $\gamma = .142$. The forecasts given by this choice of α and γ are given in Figure 16.15(c) and graphically illustrated in Figure 16.17. To evaluate the choice of a particular set of values for α and γ, MINITAB gives the **mean of the absolute forecast errors (the MAD)** and the **mean of the squared forecast errors (the MSD)** for the 24 historical calculator sales values. Comparing Figures 16.16 and 16.17, we see that $\alpha = .2$ and $\gamma = .2$ give a smaller MAD and MSD than do $\alpha = .496$ and $\gamma = .142$. Therefore, we might conclude that we should use the forecasts of y_{25} through y_{30} based on $\alpha = .2$ and $\gamma = .2$. On the other hand, we might believe that the lower sales values at the end of the observed data signal that the sales values will not continue to increase as fast as they have been increasing. In this case, we might use the lower forecasts given by $\alpha = .496$ and $\gamma = .142$ (see Figure 16.17).

LO16-7 Use multiplicative Winters' method to forecast a time series.

Multiplicative Winters' method (requires material from Section 16.3) The **Multiplicative Winters' method** can be used to forecast time series that are described by the model

$$y_t = (\beta_0 + \beta_1 t) \times SN_t + \varepsilon_t$$

Here we assume that β_0 and β_1 (and thus the linear trend) and SN_t (which represents the seasonal pattern) may be changing slowly over time. To implement the multiplicative Winters' method, we let ℓ_{T-1} denote the estimate of the deseasonalized level $\beta_0 + \beta_1(T - 1)$ of the time series in time period $T - 1$, and we let b_{T-1} denote the estimate of the slope β_1 of the time series in time period $T - 1$. Then, suppose that we observe a new time series value y_T in time period T, and let sn_{T-L} denote the "most recent" estimate of the seasonal factor for the season corresponding to time period T. Here L denotes the number of seasons in a year ($L = 12$ for monthly data, and $L = 4$ for quarterly data),

FIGURE 16.16 The MINITAB Graphical Forecasts of Calculator Sales When α = .2 and γ = .2

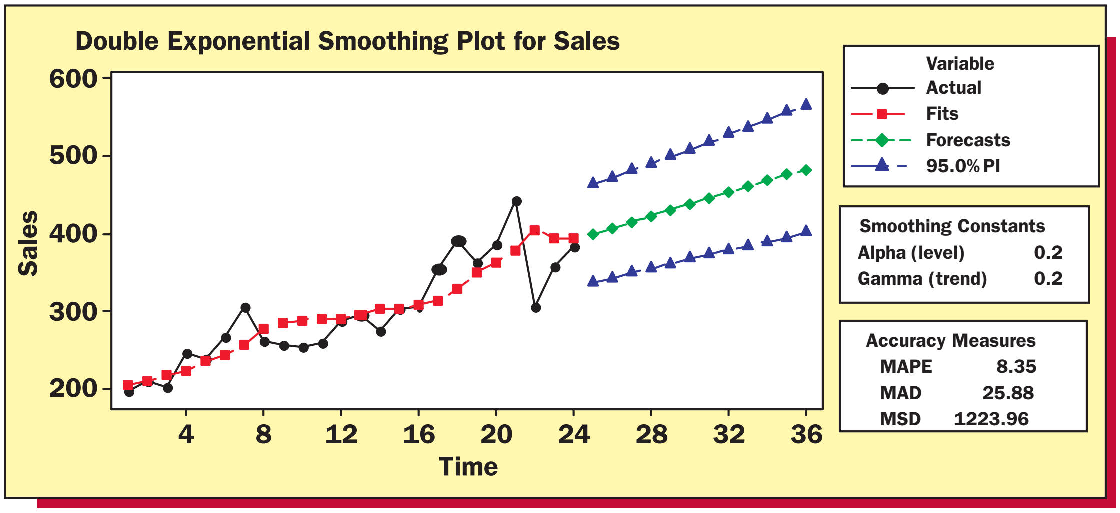

FIGURE 16.17 The MINITAB Graphical Forecasts When α = .496 and γ = .142

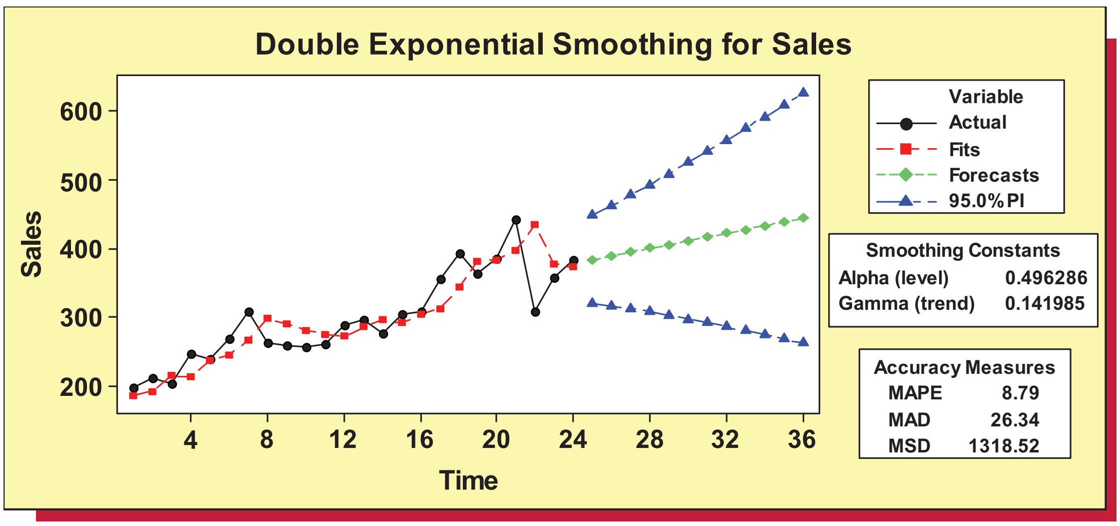

and thus $T - L$ denotes the time period occurring one year prior to time period T. Furthermore, the subscript $T - L$ of sn_{T-L} denotes the fact that the time series value observed in time period $T - L$ was the most recent time series value observed in the season being analyzed and thus was the most recent time series value used to help find sn_{T-L}. Then, the estimate of the deseasonalized level $\beta_0 + \beta_1 T$ of the time series in time period T uses the smoothing constant α and is

$$\ell_T = \alpha \frac{y_T}{sn_{T-L}} + (1 - \alpha)[\ell_{T-1} + b_{T-1}]$$

where y_T/sn_{T-L} is the deseasonalized observation in time period T. The estimate of the slope β_1 of the time series in time period T uses the smoothing constant γ and is

$$b_T = \gamma[\ell_T - \ell_{T-1}] + (1 - \gamma)b_{T-1}$$

The new estimate of the seasonal factor SN_T in time period T uses the smoothing constant δ and is

$$sn_T = \delta \frac{y_T}{\ell_T} + (1 - \delta)sn_{T-L}$$

where y_T/ℓ_T is an estimate of the newly observed seasonal variation.

To use the updating equations, we first obtain initial estimates ℓ_0, b_0, and sn_0 of the deseasonalized level, slope, and seasonal factors of the time series in time period 0. One way to do this is to use the multiplicative decomposition method (see Section 16.3 on page 640) to analyze part (say, one-half) of the historical data. Here, if there are less than five years of historical data, it is probably best to base the initial estimates on all of the historical data. Then, we regard the y-intercept and slope of the trend line fit to the deseasonalized data as the initial estimates ℓ_0 and b_0. Furthermore, we regard the multiplicative decomposition method's seasonal factors as the initial estimates of the seasonal factors in time period 0. For example, consider the 36 Tasty Cola sales values in Table 16.9 (page 640). Using the multiplicative decomposition method results summarized in Tables 16.10 (page 641) and 16.11 (page 643), we obtain the initial estimates $\ell_0 = 380.163$ and $b_0 = 9.489$ and the initial seasonal factor estimates given in the page margin. Starting with these initial estimates, we calculate a point forecast of y_1 from time origin 0 to be

Initial Seasonal Factor Estimates

Month	sn_0
Jan.	.493
Feb.	.596
Mar.	.595
Apr.	.680
May	.564
June	.986
July	1.467
Aug.	1.693
Sept.	1.990
Oct.	1.307
Nov.	1.029
Dec.	.600

$$\begin{aligned}\hat{y}_1(0) &= (\ell_0 + b_0)sn_0 \\ &= (380.163 + 9.489)(.493) \\ &= 192.098\end{aligned}$$

Here we have used the initial January seasonal factor estimate $sn_0 = .493$ because y_1 is Tasty Cola sales in January of year 1. The actual value of y_1 is 189, so the forecast error is

$$y_1 - \hat{y}_1(0) = 189 - 192.098 = -3.098$$

We next choose values of the smoothing constants α, γ, and δ. A reasonable choice (and the default option of MINITAB) is to let each of α, γ, and δ be .2. Then, using $y_1 = 189$ and the equation for ℓ_T, it follows that the estimate of the deseasonalized level of the time series in time period 1 is

$$\begin{aligned}\ell_1 &= \alpha \frac{y_1}{sn_0} + (1 - \alpha)[\ell_0 + b_0] \\ &= .2\left[\frac{189}{.493}\right] + .8[380.163 + 9.489] \\ &= 388.395\end{aligned}$$

Here we have used the initial January seasonal factor estimate $sn_0 = .493$ as the most recent Winters' method estimate of the January seasonal factor. Using the equation for b_T, the estimate of the slope of the time series in time period 1 is

$$\begin{aligned}b_1 &= \gamma[\ell_1 - \ell_0] + (1 - \gamma)b_0 \\ &= .2[388.395 - 380.163] + .8(9.489) \\ &= 9.238\end{aligned}$$

Using the equation for sn_T, the new estimate of the January seasonal factor in time period 1 is

$$\begin{aligned}sn_1 &= \delta \frac{y_1}{\ell_1} + (1 - \delta)sn_0 \\ &= .2\left[\frac{189}{388.395}\right] + .8(.493) \\ &= .492\end{aligned}$$

It follows that a point forecast made in period 1 of y_2 is

$$\begin{aligned}\hat{y}_2(1) &= (\ell_1 + b_1)sn_0 \\ &= (388.395 + 9.238)(.596) \\ &= 236.989\end{aligned}$$

Here we have used the initial February seasonal factor estimate $sn_0 = .596$ because y_2 is the Tasty Cola sales in February of year 1. The actual value of y_2 is 229, so the forecast error is

$$y_2 - \hat{y}_2(1) = 229 - 236.989 = -7.989$$

The MINITAB output in Figure 16.18(a) on the next page shows the entire process of using the Winters' method updating equations to find new period-by-period estimates of the level, slope, and seasonal factors of the time series. The output also shows the one-period-ahead forecasts and forecast errors, which are utilized to evaluate the effectiveness of the Winters' method procedure. MINITAB does not find initial estimates by using the multiplicative decomposition method. We will not discuss how MINITAB obtains initial estimates, but note from Figure 16.18(a) that the values of ℓ_1 and b_1 obtained by MINITAB ($\ell_1 = 278.768$ and $b_1 = 44.9736$) are very different from the values that we obtained by hand calculation ($\ell_1 = 388.395$ and $b_1 = 9.238$). In addition, the one-period-ahead forecast errors obtained by MINITAB are generally quite large in periods 1 through 12 but then become reasonably small for periods 13 through 36. To further illustrate the Winters' method updating equations, note from Figure 16.18(a) that $\ell_{35} = 725.603$ and $b_{35} = 8.9026$. Because the most recent estimate of the December seasonal factor is $sn_{24} = .60767$, the point forecast made in period 35 of y_{36} (sales in December of year 3) is

$$\begin{aligned}\hat{y}_{36}(35) &= (\ell_{35} + b_{35})sn_{24} \\ &= (725.603 + 8.9026)(.60767) \\ &= 446.34\end{aligned}$$

The actual sales value in period 36 is $y_{36} = 441$, so the forecast error is

$$y_{36} - \hat{y}_{36}(35) = 441 - 446.34 = -5.34$$

The updated estimates ℓ_{36}, b_{36}, and sn_{36} are calculated as follows:

$$\begin{aligned}\ell_{36} &= \alpha \frac{y_{36}}{sn_{24}} + (1 - \alpha)[\ell_{35} + b_{35}] \\ &= .2\left[\frac{441}{.60767}\right] + .8[725.603 + 8.9026] \\ &= 732.75 \\ b_{36} &= \gamma[\ell_{36} - \ell_{35}] + (1 - \gamma)b_{35} \\ &= .2[732.75 - 725.603] + .8(8.9026) \\ &= 8.5514\end{aligned}$$

and

$$\begin{aligned}sn_{36} &= \delta \frac{y_{36}}{\ell_{36}} + (1 - \delta)sn_{24} \\ &= .2\left[\frac{441}{732.75}\right] + .8(.60767) \\ &= .6065\end{aligned}$$

We are now at the end of the historical data, so we can forecast future Tasty Cola sales values. Figure 16.18(b) gives the point and 95 percent prediction interval forecasts of future sales values in periods 37 through 48, and Figure 16.19 graphically portrays the forecasts. To see how the point forecasts are calculated, note that, for example, the most recent estimates of the January and July seasonal factors are $sn_{25} = .48019$ and $sn_{31} = 1.42891$. Therefore, point forecasts made in period 36 of Tasty Cola sales in periods 37 and 43 (January and July of year 4) are

$$\hat{y}_{37}(36) = (\ell_{36} + b_{36})sn_{25} = (732.75 + 8.5514)(.48019) = 355.96$$

and

$$\hat{y}_{43}(36) = (\ell_{36} + 7b_{36})sn_{31} = [732.75 + 7(8.5514)](1.42891) = 1{,}132.57$$

FIGURE 16.18 The MINITAB Output of Winters' Method for the Tasty Cola Sales Data, When $\alpha = .2$, $\gamma = .2$, and $\delta = .2$

(a) The updated level, slope, and seasonal factor estimates

Time T	Sales y_T	Level ℓ_T	Slope b_T	Seasonal sn_T	Forecast $\hat{y}_T(T-1)$	Error $y_T - \hat{y}_T(T-1)$
1	189	278.768	44.9736	0.48896	106.67	82.334
2	229	343.270	48.8794	0.56818	175.93	53.065
3	249	401.836	50.8167	0.57606	221.63	27.371
4	289	449.774	50.2409	0.65605	298.49	-9.492
5	260	492.009	48.6398	0.55787	282.62	-22.624
6	431	520.567	44.6235	0.94880	529.30	-98.301
7	660	541.448	39.8750	1.42638	835.48	-175.485
8	777	556.089	34.8280	1.64516	992.39	-215.395
9	915	562.722	29.1891	1.95207	1201.68	-286.680
10	613	565.315	23.8699	1.28544	790.62	-177.623
11	485	563.116	18.6561	1.01787	622.78	-137.777
12	277	552.752	12.8521	0.60770	369.04	-92.044
13	244	552.287	10.1887	0.47953	276.56	-32.557
14	296	554.174	8.5282	0.56137	319.59	-23.586
15	319	560.914	8.1706	0.57459	324.15	-5.151
16	370	568.063	7.9664	0.65511	373.35	-3.349
17	313	573.035	7.3675	0.55554	321.35	-8.352
18	556	581.523	7.5916	0.95026	550.68	5.315
19	831	587.811	7.3308	1.42385	840.30	-9.301
20	960	592.820	6.8664	1.64000	979.10	-19.101
21	1152	597.777	6.4846	1.94709	1170.63	-18.631
22	759	601.501	5.9325	1.28072	776.74	-17.742
23	607	605.216	5.4890	1.01488	618.29	-11.287
24	371	610.664	5.4807	0.60767	371.13	-0.126
25	298	617.205	5.6927	0.48019	295.46	2.542
26	378	632.989	7.7111	0.56853	349.67	28.326
27	373	642.391	8.0493	0.57580	368.14	4.859
28	443	655.597	9.0806	0.65923	426.11	16.891
29	374	666.385	9.4222	0.55668	369.26	4.743
30	660	679.556	10.1717	0.95445	642.19	17.807
31	1004	692.808	10.7879	1.42891	982.07	21.934
32	1153	703.487	10.7660	1.63980	1153.90	-0.898
33	1388	713.974	10.7103	1.94648	1390.71	-2.712
34	904	720.918	9.9571	1.27537	928.12	-24.118
35	715	725.603	8.9026	1.00898	741.75	-26.753
36	441	732.750	8.5514	0.60650	446.34	-5.336

(b) Point and 95 percent prediction interval forecasts

Period	Forecast	Lower	Upper
37	355.96	240.98	470.95
38	426.31	309.52	543.10
39	436.69	317.90	555.49
40	505.60	384.60	626.60
41	431.71	308.32	555.10
42	748.35	622.39	874.31
43	1132.57	1003.88	1261.26
44	1313.74	1182.16	1445.32
45	1576.09	1441.48	1710.70
46	1043.59	905.81	1181.37
47	834.24	693.17	975.32
48	506.65	362.17	651.14

FIGURE 16.19 MINITAB Output of Using Winters' Method to Forecast Tasty Cola Sales

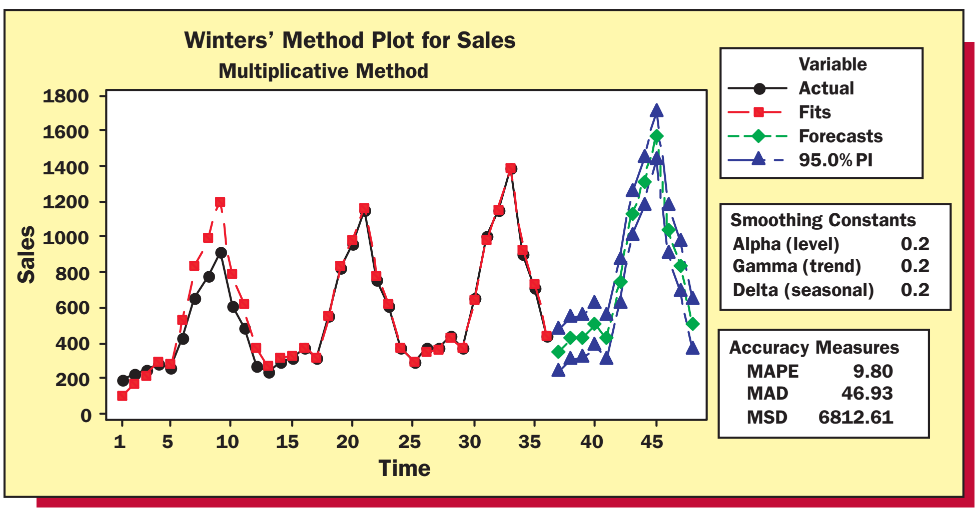

The reason that the 95 percent prediction intervals are so wide is that they can be shown to be functions of the historical forecast errors, which are very large in periods 1 through 12. The mean absolute forecast error in periods 13 through 36 can be calculated to be 12.98 and is more representative of the Winters' method's accuracy than is the mean absolute forecast error in all 36 periods, which is 46.93 (see Figure 16.19). Therefore, to obtain more reasonable prediction intervals, we might multiply the lengths of the prediction intervals by $12.98/46.93 \approx .28$. For example, Figure 16.18(b) tells us that the 95 percent prediction interval for y_{37} is [240.98, 470.95], which has length $470.95 - 240.98 = 229.97$. Multiplying this length by .28, we obtain $(229.97)(.28) = 64.39$. Surrounding the point forecast 355.96 by a new half-length of $64.39/2 = 32.2$, we obtain a new 95 percent prediction interval of $[355.96 \pm 32.2] = [323.76, 388.16]$. The other 95 percent prediction intervals can be modified similarly.

The wide prediction intervals in Figure 16.18(b) result from a combination of a short historical series (36 sales values) and MINITAB obtaining inaccurate initial estimates of the level, slope, and seasonal factors. When the historical series is long (for example, see Exercise 16.28, page 660), MINITAB usually obtains reasonable prediction intervals. Finally, note that MINITAB will not choose its own values of α, γ, and δ. However, the user can simply experiment with different combinations of values of these smoothing constants until a combination is found that produces the "best" results.

Exercises for Section 16.5

CONCEPTS

16.23 When do we use double exponential smoothing?

16.24 When do we use the multiplicative Winters' method?

METHODS AND APPLICATIONS

16.25 Consider Figure 16.15 on page 653. Calculate ℓ_2 and b_2 from ℓ_1, b_1, and y_2. Also, calculate $\hat{y}_{27}(24)$ from ℓ_{24} and b_{24}.

16.26 Consider Figure 16.18 on page 658. Calculate ℓ_{35}, b_{35}, and sn_{35} from ℓ_{34}, b_{34}, y_{35}, and sn_{23}. Also, calculate $\hat{y}_{38}(36)$ from ℓ_{36}, b_{36}, and sn_{26}.

16.27 THE WATCH SALES CASE DS WatchSale

Figure 16.20 gives the MINITAB output of using double exponential smoothing in month 20 to forecast watch sales in months 21 through 26. Here we have used MINITAB's default option that sets each of the smoothing constants alpha and gamma equal to .2. Find and report the point prediction of and a 95 percent prediction interval for watch sales in month 21.

FIGURE 16.20 MINITAB Output of Using Double Exponential Smoothing to Forecast Watch Sales

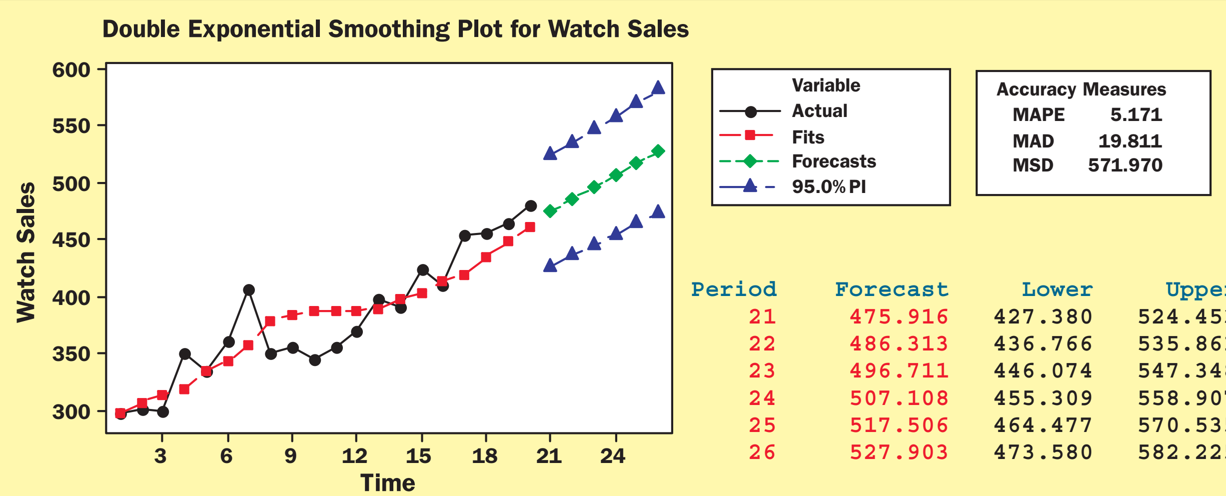

Period	Forecast	Lower	Upper
21	475.916	427.380	524.453
22	486.313	436.766	535.861
23	496.711	446.074	547.348
24	507.108	455.309	558.907
25	517.506	464.477	570.535
26	527.903	473.580	582.225

FIGURE 16.21 MINITAB Output of Using Winters' Method to Forecast the Monthly Hotel Room Averages

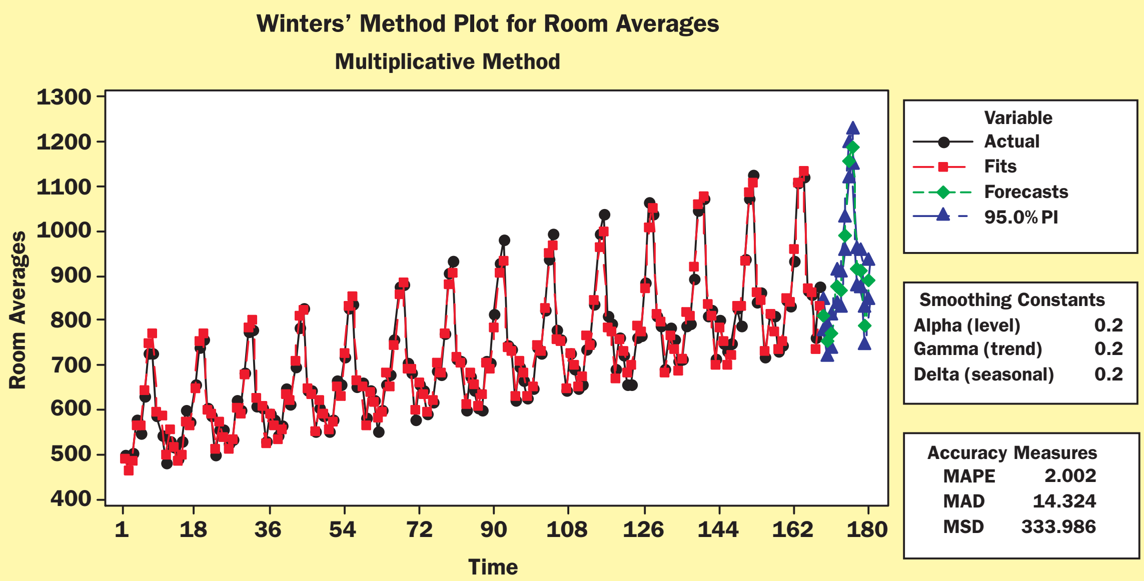

Period	Forecast	Lower	Upper	Period	Forecast	Lower	Upper
169	810.04	774.95	845.13	175	1158.99	1119.72	1198.27
170	754.55	718.91	790.20	176	1189.52	1149.37	1229.68
171	772.31	736.06	808.57	177	917.33	876.25	958.41
172	874.62	837.69	911.55	178	912.83	870.78	954.88
173	867.63	829.97	905.28	179	787.09	744.04	830.15
174	992.02	953.58	1030.47	180	887.66	843.56	931.75

16.28 THE TRAVELER'S REST CASE DS TravRest

Figure 16.21 gives the MINITAB output of using multiplicative Winters' method in month 168 to forecast the monthly hotel room averages in months 169 through 180. Here we have used MINITAB's default option that sets each of the smoothing constants alpha, gamma, and delta equal to .2. Use the MINITAB output to find and report the point prediction of and a 95 percent prediction interval for the monthly hotel room average in period 169.

LO16-8 Appreciate some of the basic concepts of Box–Jenkins forecasting models.

16.6 A Brief Introduction to Box–Jenkins Models (Optional Advanced Section)

In this section we discuss the **Box–Jenkins methodology.** This methodology, developed by G. E. P. Box and G. M. Jenkins (1976), uses an approach to describe the trend and seasonal effects in time series data that is quite different from the approach taken by regression (or exponential smoothing). The Box–Jenkins methodology begins by determining if the time series under consideration is *stationary*. Intuitively, a time series is **stationary** if the statistical properties (for example, the level and the variance) of the time series are essentially constant through time. If we have observed n values $y_1, y_2, \ldots, y_n$ of a time series, we can use a plot of these values (against time) to help us determine whether the time series is stationary. If the n values seem to fluctuate with constant variation around a constant level, then it is reasonable to believe that the time series is stationary. If the n values do not fluctuate around a constant level or do not fluctuate with constant variation, then it is reasonable to believe that the time series is nonstationary. In this case, the Box–Jenkins methodology tells us to *transform* the nonstationary time series values into stationary time series values.

For example, recall that Table 16.8 on page 639 gives monthly international passenger totals for the past 11 years for an airline. The partial MINITAB output in Figure 16.22(a) gives the first 15 and the last 2 of these monthly passenger totals (see Pass). The plot of these passenger totals is

FIGURE 16.22 The Passenger Totals, Natural Logarithms, and Differences Data, and Time Series Plots of the Passenger Totals and Natural Logarithms

(a) Partial MINITAB data output

t	Pass	LnPass	Reg	Seas	RegSeas
1	112	4.71850	*	*	*
2	118	4.77068	0.052186	*	*
3	132	4.88280	0.112117	*	*
4	129	4.85981	−0.022990	*	*
5	121	4.79579	−0.064022	*	*
6	135	4.90527	0.109484	*	*
7	148	4.99721	0.091937	*	*
8	148	4.99721	0.000000	*	*
9	136	4.91265	−0.084557	*	*
10	119	4.77912	−0.133531	*	*
11	104	4.64439	−0.134733	*	*
12	118	4.77068	0.126294	*	*
13	115	4.74493	−0.025752	0.026433	*
14	126	4.83628	0.091350	0.065597	0.039164
15	141	4.94876	0.112478	0.065958	0.000361
131	362	5.89164	−0.117169	0.155072	0.029581
132	405	6.00389	0.112243	0.183804	0.028732

(b) Passenger totals plot

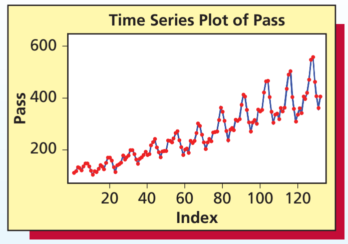

(c) Natural logarithm plot

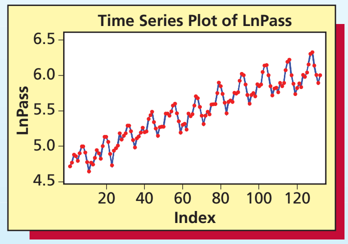

shown in Figure 16.22(b) and indicates that the passenger totals exhibit an upward linear trend with seasonal variation that increases with the level of the totals. In terms of Box–Jenkins modeling, the **increasing seasonal variation** implies that the time series is **nonstationary with respect to its variance** and we should use a variance stabilizing transformation. Figure 16.22(a) gives the natural logarithms (see LnPass) of the monthly passenger totals for the time periods shown in the figure, and Figure 16.22(c) shows that the natural logarithms exhibit an upward linear trend with constant (variance-stabilized) seasonal variation. That is, letting y_t^* denote the variance-stabilized passenger total in month t, it follows that $y_t^* = \ln y_t$. The linear trend and seasonal variation in the y_t^* values indicate that, although these values are stationary with respect to their variance, they are **nonstationary with respect to their level.** To transform these values into stationary values with respect to their level, the Box–Jenkins methodology uses one of three types of *differencing: regular differencing*, *seasonal differencing*, or *combined regular and seasonal differencing*.

The **regular difference** of the passenger totals' natural logarithms in time period t, denoted z_t, is the difference between y_t^*, the natural logarithm in time period t, and y_{t-1}^*, the natural logarithm in time period $t-1$. That is, $z_t = y_t^* - y_{t-1}^*$. For example, Figure 16.22(a) tells us that the natural logarithms in months 1, 2, 3, 131, and 132 are $y_1^* = 4.71850$, $y_2^* = 4.77068$, $y_3^* = 4.88280$, $y_{131}^* = 5.89164$, and $y_{132}^* = 6.00389$. Therefore, the regular differences of the natural logarithms are:

$$z_2 = y_2^* - y_1^* = 4.77068 - 4.71850 = .05218$$
$$z_3 = y_3^* - y_2^* = 4.88280 - 4.77068 = .11212$$
$$\vdots$$
$$z_{132} = y_{132}^* - y_{131}^* = 6.00389 - 5.89164 = .11225$$

Figure 16.22(a) gives the regular differences (see Reg) of the natural logarithms for the time periods shown in the figure. Note that, because MINITAB does calculations using more precise decimal place accuracy than it sometimes shows, the results in Figure 16.22(a) are slightly different from the results that we hand calculate. To calculate the **seasonal differences** of the natural logarithms in time period t, we use the equation $z_t = y_t^* - y_{t-L}^*$, where L is the number of seasons in a year. Because the passenger totals data is monthly, L equals 12, and thus we use the equation $z_t = y_t^* - y_{t-12}^*$. Figure 16.22(a) tells us that $y_1^* = 4.71850$, $y_2^* = 4.77068$, $y_{13}^* = 4.74493$, and $y_{14}^* = 4.83628$. It follows that the first two seasonal differences are:

$$z_{13} = y_{13}^* - y_1^* = 4.74493 - 4.71850 = .02643 \quad \text{and} \quad z_{14} = y_{14}^* - y_2^* = 4.83628 - 4.77068 = .0656$$

To carry out **combined regular and seasonal differencing,** we take the regular differences of the seasonal differences. For example, the regular difference of the just calculated two seasonal

FIGURE 16.23 **Time Series Plots for Analyzing the Passenger Totals' Natural Logarithms**

(a) Regular difference plot **(b) Seasonal difference plot** **(c) Combined difference plot**

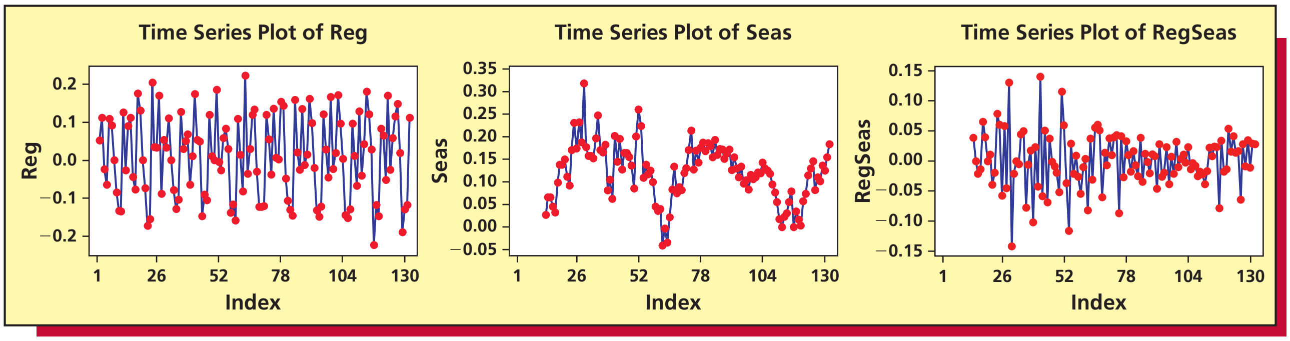

differences is .0656 − .02643 = .03917. Figure 16.22(a) gives the seasonal differences (see Seas) and the combined regular and seasonal differences (see RegSeas) of the natural logarithms for the time periods shown in the figure. Moreover, Figure 16.23 shows time series plots of the regular differences, the seasonal differences, and the combined regular and seasonal differences of the natural logarithms. Examining the plot in Figure 16.23(a), it might seem at first glance that the regular differences are fluctuating with constant variation around a constant level. However, the periodic "dips" in the plot indicate that the regular differences exhibit a seasonal pattern and thus are nonstationary. The plot in Figure 16.23(b) shows that the seasonal differences do not exhibit any long term upward or downward trend movements or any seasonal pattern. However, because these seasonal differences exhibit repeated, short term upward and downward trend movements, we say that they exhibit **changing** (or **stochastic**) **trend effects.** These stochastic trend effects imply that the seasonal differences are nonstationary. Finally, Figure 16.23(c) shows that the combined regular and seasonal differences do not seem to exhibit either trend effects or a seasonal pattern. Because these combined differences seem to fluctuate with (approximately) constant variation around a constant level, we will regard them as being stationary.

A Box–Jenkins model Because the combined regular and seasonal differences are stationary, we will find a Box–Jenkins model describing them. The combined regular and seasonal difference in time period t, z_t, is the regular difference of the seasonal difference $y_t^* - y_{t-12}^*$ and thus can be expressed as

$$z_t = y_t^* - y_{t-12}^* - (y_{t-1}^* - y_{t-13}^*)$$

The simplest Box–Jenkins model describing z_t would set z_t equal to an error term ε_t. This would give

$$y_t^* - y_{t-12}^* - (y_{t-1}^* - y_{t-13}^*) = \varepsilon_t$$

Solving for y_t^*, we would have the following Box–Jenkins model describing y_t^*:

$$y_t^* = y_{t-12}^* + (y_{t-1}^* - y_{t-13}^*) + \varepsilon_t$$

If we temporarily ignore the error term ε_t, this model says that y_t^*, the passenger totals' natural logarithm in time period t, equals y_{t-12}^*, the passenger totals' natural logarithm one year ago, plus $(y_{t-1}^* - y_{t-13}^*)$, the difference between the passenger totals' natural logarithm last month and the passenger totals' natural logarithm one year ago last month. Of course, this and all Box–Jenkins models describe the time series value y_t^* by using ε_t, the error term in time period t, and in fact some Box–Jenkins models also describe y_t^* by using $\varepsilon_{t-1}, \varepsilon_{t-2}, \ldots$, the error terms in previous time periods. Such previous error terms are called *moving average terms.* Later in this section we

FIGURE 16.24 Analysis of the Passenger Totals' Box–Jenkins Model

(a) Estimation and diagnostic checking

Type	Coef	SE Coef	T	P
MA 1	0.3407	0.0868	3.93	0.000
SMA 12	0.6299	0.0766	8.23	0.000

Modified Box-Pierce (Ljung-Box) Chi-Square statistic

Lag	12	24	36	48
Chi-Square	7.5	19.6	30.5	38.7
DF	10	22	34	46
P-Value	0.679	0.607	0.638	0.770

(b) Residuals

t	LnPass	FITS	RESI
120	5.82008	5.86184	−0.041754
121	5.88610	5.85600	0.030104
122	5.83481	5.83182	0.002987
123	6.00635	5.98014	0.026218
124	5.98141	5.97049	0.010927
125	6.04025	6.00109	0.039167
126	6.15698	6.18944	−0.032462
127	6.30628	6.27991	0.026361
128	6.32615	6.30195	0.024195
129	6.13773	6.15349	−0.015763
130	6.00881	6.01104	−0.002232
131	5.89164	5.87364	0.018005
132	6.00389	5.99133	0.012558

(c) Forecasting

		95% Limits	
Period	Forecast	Lower	Upper
133	6.03771	5.96718	6.10823
134	5.99099	5.90652	6.07546
135	6.14666	6.05023	6.24308
136	6.12046	6.01341	6.22751
137	6.15698	6.04026	6.27369
138	6.30256	6.17692	6.42819
139	6.42828	6.29432	6.56224
140	6.43857	6.29677	6.58037
141	6.26527	6.11604	6.41450
142	6.13438	5.97807	6.29069
143	6.00539	5.84231	6.16846
144	6.11358	5.94401	6.28316

will see that the appropriate Box–Jenkins model describing the logged passenger total y_t^* uses the moving average terms ε_{t-1}, ε_{t-2}, and ε_{t-13} and is

$$y_t^* = y_{t-12}^* + (y_{t-1}^* - y_{t-13}^*) + \varepsilon_t - \theta_1\varepsilon_{t-1} - \theta_{12}\varepsilon_{t-12} + \theta_1\theta_{12}\varepsilon_{t-13}$$

Here, θ_1 and θ_{12} are unknown parameters that must be estimated from sample data, and the minus signs are used because of the theory behind the Box–Jenkins methodology. The procedure for finding such a model is quite complex. For this reason we will first give the reader an intuitive understanding of the model by using it to forecast future time series values.

The MINITAB output in Figure 16.24(a) tells us that the least squares point estimates of θ_1 and θ_{12} are $\hat{\theta}_1 = .3407$ and $\hat{\theta}_{12} = .6299$. Moreover, because the *p*-values for testing $H_0: \theta_1 = 0$ and $H_0: \theta_{12} = 0$ are each less than .001, we conclude that each of θ_1 and θ_{12} is important in the model. (The lower half of Figure 16.24(a) provides a *Chi-square analysis* that is used to *diagnostically check* the adequacy of the model. This will be explained later.) To compute a point forecast of y_{133}^*, next month's logged passenger total, we insert $t = 133$, as well as the least squares point estimates $\hat{\theta}_1 = .3407$ and $\hat{\theta}_{12} = .6299$, into the model. This gives the point forecast

$$\hat{y}_{133}^*(132) = y_{121}^* + (y_{132}^* - y_{120}^*) + \hat{\varepsilon}_{133} - .3407\,\hat{\varepsilon}_{132} - .6299\,\hat{\varepsilon}_{121} + (.3407)(.6299)\,\hat{\varepsilon}_{120}$$

To compute this point forecast, we use the natural logarithms (see LnPass in Figure 16.24(b)) $y_{120}^* = 5.82008$, $y_{121}^* = 5.88610$, and $y_{132}^* = 6.00389$ of the passenger totals in months 120, 121, and 132. In addition, we predict the future error term ε_{133} to be $\hat{\varepsilon}_{133} = 0$, and we predict the past error terms ε_{132}, ε_{121}, and ε_{120} by using MINITAB computed residuals for time periods 132, 121, and 120. To obtain the residuals, MINITAB uses an advanced technique and the Box–Jenkins model that it has fit to the 132 logged passenger totals to calculate a point prediction of each logged passenger total. The residual for a particular month is then the difference between the observed logged passenger total and the predicted logged passenger total for the month. Figure 16.24(b) shows the residuals for months 120 through 132. For month 120, for example, the observed logged passenger total is $y_{120}^* = 5.82008$ (see LnPass), the predicted logged passenger total is $\hat{y}_{120}^* = 5.86184$ (see FITS), and the residual $\hat{\varepsilon}_{120}$ is $y_{120}^* - \hat{y}_{120}^* = 5.82008 - 5.86184 = -.041754$ (see RESI). Figure 16.24(b) tells us that the residuals for time periods 132, 121, and 120 are $\hat{\varepsilon}_{132} = .012558$, $\hat{\varepsilon}_{121} = .030104$, and $\hat{\varepsilon}_{120} = -.041754$. Inserting all of the needed quantities into the above equation for the point forecast $\hat{y}_{133}^*(132)$, we find that

$$\begin{aligned}\hat{y}_{133}^*(132) &= 5.88610 + (6.00389 - 5.82008) + 0 \\ &\quad - .3407(.012558) - .6299(.030104) \\ &\quad + (.3407)(.6299)(-.041754) \\ &= 6.03771\end{aligned}$$

This point forecast and a 95 percent prediction interval for y_{133}^* are given in Figure 16.24(c). Exponentiating the point forecast and the ends of the prediction interval, we obtain $e^{6.03771} = 418,933$ and $[e^{5.96718}, e^{6.10823}] = [390,403, 449,542]$. Therefore, the airline forecasts that it will fly 418,933 customers in month 133 (January of year 12). The airline is 95 percent confident it will fly between 390,403 and 449,542 customers in month 133.

FIGURE 16.25 The SAC and SPAC of the Combined Regular and Seasonal Differences

(a) The SAC (b) The SPAC

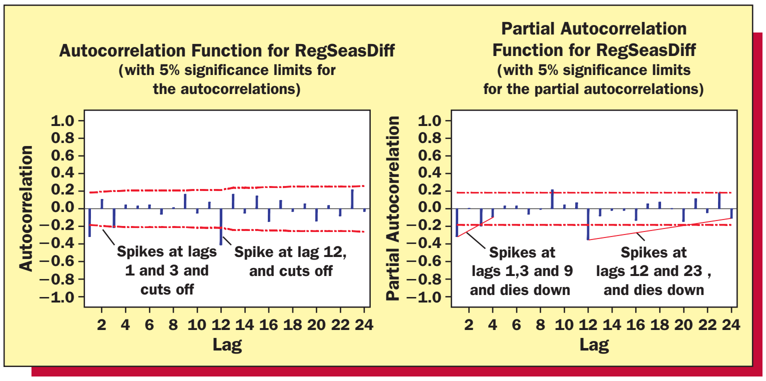

FIGURE 16.26 SAC of the Passenger Totals' Natural Logarithms

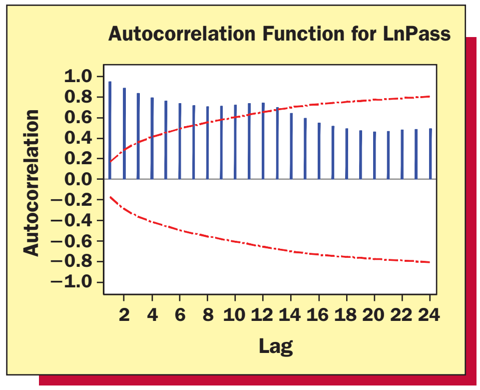

The sample autocorrelation function In order to help find a Box–Jenkins model describing a time series of z_t values, we use the *sample autocorrelation function (SAC)* and the *sample partial autocorrelation (SPAC)* of the z_t values. To explain the sample autocorrelation function, we first define the **sample autocorrelation at lag *k*,** denoted r_k, to be the simple correlation coefficient between z_t values separated by k time units. For example, r_1, r_2, and r_3 are the simple correlation coefficients between z_t values separated by, respectively, one, two, and three time units. To illustrate calculating (for instance) r_3, consider the combined regular and seasonal differences of the passenger totals' natural logarithms. In Figure 16.22(a) we have computed these combined differences to be $z_{14} = .039164$, $z_{15} = .000361, \ldots, z_{131} = .029581$, and $z_{132} = .028732$. The mean of these combined differences can be shown to equal $\bar{z} = .001322$. Inserting these z_t values and $\bar{z}$ into the following equation

$$r_3 = \frac{(z_{14} - \bar{z})(z_{17} - \bar{z}) + (z_{15} - \bar{z})(z_{18} - \bar{z}) + \cdots + (z_{129} - \bar{z})(z_{132} - \bar{z})}{(z_{14} - \bar{z})^2 + (z_{15} - \bar{z})^2 + \cdots + (z_{132} - \bar{z})^2}$$

it follows that r_3, the simple correlation coefficient between z_t values separated by three time units, can be calculated to be $-.216$. In general, we define the **sample autocorrelation function (SAC)** to be a listing, or graph, of the sample autocorrelations at lags $k = 1, 2, \ldots$. Figure 16.25(a) presents the MINITAB output of the SAC of the combined differences of the passenger totals' natural logarithms. The vertical lines on the output represent the sample autocorrelations. For example, the vertical lines at lags 1, 2, and 3 represent r_1, r_2, and r_3, which can be computed to be $-.317$, $.109$, and (as just demonstrated) $-.216$. In order to begin to interpret the meaning of these and the other sample autocorrelations in the SAC, we identify the existence of **spikes** in the SAC. We say that **a spike exists at lag *k* in the SAC** if the sample autocorrelation r_k is large enough in magnitude to conclude that ρ_k, the population autocorrelation of all possible z_t values separated by k time units, does not equal 0. One frequently used convention is to conclude that a spike exists at lag k in the SAC if r_k is at least as large in magnitude as twice its estimated standard deviation. If we examine the MINITAB output of the SAC in Figure 16.25(a), we see that the "center line" on the plot of the r_k values is positioned at 0. Furthermore, for any lag k the dashed line above the center line is two estimated standard deviations greater than 0, and the dashed line below the center line is two estimated standard deviations less than 0. If the vertical line representing r_k extends at least as far up or down as the two standard deviation dashed lines corresponding to lag k, then we conclude that a spike exists at lag k in the SAC. Examining Figure 16.25(a), we conclude that spikes exist at lags 1, 3, and 12 in the SAC.

The sample partial autocorrelation function The **sample partial autocorrelation function (SPAC)** of a time series of z_t values is a listing, or graph, of the **sample partial autocorrelations** at lags $k = 1, 2, \ldots$. It is beyond the scope of this text to give a precise definition of the **sample partial autocorrelation at lag *k*.** However, this quantity, which is denoted r_{kk}, may intuitively be thought of as the sample autocorrelation of z_t values separated by k time units **with the effects of the intervening z_t values eliminated.** Figure 16.25(b) gives the MINITAB output of the SPAC of the combined differences of the passenger totals' natural logarithms. Examining this output, we see that, for example, the vertical lines at lags 1 and 2 represent r_{11} and r_{22}, the sample partial autocorrelations of the combined differences separated by, respectively, one and two time units. In general, if r_{kk} is at least as large in magnitude as twice its estimated standard deviation, we say that **a spike exists at lag k in the SPAC.** This implies that ρ_{kk}, the population partial autocorrelation of all possible z_t values separated by k time units, does not equal 0. The r_{kk} values are plotted in the MINITAB output of Figure 16.25(b). We conclude that a spike exists at lag k in the SPAC if the vertical line representing r_{kk} extends at least as far up or down as the two standard deviation dashed lines corresponding to lag k. Because this occurs at lags 1, 3, 9, 12, and 23, we conclude that spikes occur at these lags in the SPAC. We will now see that the spikes in the SAC and SPAC, as well as the overall behavior of the SAC and SPAC, help us to **tentatively identify** a Box–Jenkins model.

Moving average terms We have previously seen that some Box–Jenkins models describe the **future time series value z_t** by using **past error terms $\varepsilon_{t-1}, \varepsilon_{t-2}, \ldots,$** which are called **moving average terms.** The following guidelines tell us when to use *nonseasonal moving average terms* and *seasonal moving average terms.* [See Bowerman, O'Connell, and Koehler (2005) for the theory behind the guidelines.] To use the guidelines, we define (for monthly seasonal data) the **nonseasonal level** of the SAC and the SPAC to be lags 1 through 9 and the **seasonal level** of the SAC and the SPAC to be lags 12 and 24.

Guideline 1: Using nonseasonal moving average terms Suppose that, at the nonseasonal level, the SAC has **a spike at one or more lags and then abruptly cuts off** (that is, has no spikes) after a certain lag, and the **SPAC dies down fairly quickly** (that is, has sample partial autocorrelations that steadily decrease fairly quickly). In this case, we should describe the time series value z_t by using **past error terms (nonseasonal moving average terms) that correspond to the lags at which spikes exist in the SAC.** For example, we have seen that Figure 16.25 gives the SAC and SPAC of the combined differences of the passenger totals' natural logarithms. At the nonseasonal level, the SAC has spikes at lags 1 and 3 and then cuts off after lag 3, and the SPAC dies down fairly quickly. Here, in concluding that the SPAC dies down fairly quickly, we have ignored a spike at lag 9 in the SPAC. In general, it is a common practice to ignore spikes at higher nonseasonal lags in the SAC and SPAC. Because the SAC has spikes at lags 1 and 3, this guideline says that we should describe the combined difference z_t by using ε_{t-1} and ε_{t-3}. However, because the spike at lag 3 is barely a spike, and because the Box–Jenkins methodology seeks to find simple (*parsimonious*) models, we will begin by attempting to describe z_t by using ε_{t-1} (and not ε_{t-3}). Then, we will *diagnostically check* the model to see if we made the right decision.

Guideline 2: Using a seasonal moving average term Suppose that, at the **seasonal level,** the **SAC has a spike at lag 12 and cuts off after lag 12** (that is, has a small sample autocorrelation at lag 24), and the **SPAC dies down fairly quickly (at lags 12 and 24).** In this case, we should describe the time series value z_t by using the **past error term ε_{t-12}** (which is called a **seasonal moving average term**). For example, consider Figure 16.25. At the seasonal level, the SAC has a spike at lag 12 and cuts off after lag 12, and the SPAC dies down fairly quickly. Therefore, we should describe z_t by using the seasonal moving average term ε_{t-12}. Because we have also concluded in Guideline 1 that we should describe z_t by using ε_{t-1}, it follows that a tentative Box–Jenkins model describing z_t is

$$z_t = \varepsilon_t - \theta_1\varepsilon_{t-1} - \theta_{12}\varepsilon_{t-12} + \theta_1\theta_{12}\varepsilon_{t-13}$$

Here, theory behind the Box–Jenkins methodology tells us to use the **multiplicative component** $\theta_1\theta_{12}\varepsilon_{t-13}$. Moreover, if we insert the expression for the combined difference, $z_t = y_t^* - y_{t-12}^* - (y_{t-1}^* - y_{t-13}^*)$, into this model and solve for y_t^*, we obtain

$$y_t^* = y_{t-12}^* + (y_{t-1}^* - y_{t-13}^*) + \varepsilon_t - \theta_1\varepsilon_{t-1} - \theta_{12}\varepsilon_{t-12} + \theta_1\theta_{12}\varepsilon_{t-13}$$

Before using this tentative model to calculate forecasts (as we have done on page 663), we should check to see if the model is *adequate.* In general, it can be shown that if a tentative model is **adequate,** then the sample autocorrelations of the tentative model's residuals will not be unusually large, and thus the model's **Chi-square statistic *p*-values** will be greater than .05. Because the Chi-square statistic *p*-values in Figure 16.24(a) for the tentative passenger totals model are greater than .05, we conclude that this model is adequate.

We have seen that if a time series fluctuates with constant variation around a constant level, then the time series should be considered to be stationary. It can also be shown that if the SAC cuts off or dies down fairly quickly at both the nonseasonal and seasonal levels, then the time series should be considered to be stationary. However, if the SAC has large sample autocorrelations that die down very slowly at either level, or at both levels (for example, see Figure 16.26 on page 664), then the time series should be considered to be nonstationary.

Autoregressive Terms Some Box–Jenkins models describe the **future time series value z_t,** by using **past time series values $z_{t-1}, z_{t-2}, \ldots$,** which are called **autoregressive terms.** The following guidelines tell us when to use *nonseasonal autoregressive terms* and *seasonal autoregressive terms.*

Guideline 3: Using nonseasonal autoregressive terms Suppose that, at the nonseasonal level, the **SAC dies down fairly quickly** (that is, has sample autocorrelations that steadily decrease fairly quickly), and the **SPAC has a spike at one or more lags and then abruptly cuts off** (that is, has no spikes) after a certain lag. In this case we should describe the time series value z_t by using **past time series values (nonseasonal autoregressive terms) that correspond to the lags at which spikes exist in the SPAC.** For example, consider the seasonal differences ($z_t = y_t^* - y_{t-12}^*$) of the quartic roots ($y_t^* = y_t^{.25}$) of the hotel room averages in Table 16.4 on page 636. It can be verified that these seasonal differences fluctuate with constant variation around a constant level and are, therefore, stationary. Figure 16.27 gives the SAC and SPAC of these seasonal differences. At the nonseasonal level, the SAC dies down fairly quickly, and the SPAC has spikes at lags 1 and 3 and (with the exception of a smaller spike at lag 5) cuts off after lag 3. Because the SPAC has spikes at lags 1 and 3, this guideline says that we should describe the seasonal difference z_t by using z_{t-1} and z_{t-3}. Moreover, MINITAB requires us to use the **intervening term z_{t-2}** because we use z_{t-1} and z_{t-3}.

Guideline 4: Using a seasonal autoregressive term Suppose that, at the **seasonal level**, the **SAC dies down fairly quickly (at lags 12 and 24),** and the **SPAC has a spike at lag 12 and cuts off after lag 12** (that is, has a small sample partial autocorrelation at lag 24). In this case, we should describe the time series value z_t by using the **past time series value z_{t-12}** (which is called a **seasonal autoregressive term**). In practice, the behavior described in this guideline does not occur frequently.

Examining Figure 16.27, we see that at the seasonal level, the SAC has a spike at lag 12 and cuts off after lag 12, and the SPAC dies down fairly quickly. This is the behavior described in Guideline 2, and thus Guideline 2 says that we should describe z_t by using the seasonal moving average term z_{t-12}. Because Guideline 3 says that we should also describe z_t by using z_{t-1}, z_{t-2}, and z_{t-3}, a tentative Box–Jenkins model describing z_t is

$$z_t = \delta + \phi_1 z_{t-1} + \phi_2 z_{t-2} + \phi_3 z_{t-3} - \theta_{12}\varepsilon_{t-12} + \varepsilon_t$$

Here, we have included a **constant term** δ (to be explained momentarily) in the model, and δ, as well as ϕ_1, ϕ_2, ϕ_3, and θ_{12} are unknown parameters that must be estimated from sample data. Inserting the seasonal difference $z_t = y_t^* - y_{t-12}^*$ into this model and solving for y_t^*, we find that a tentative Box–Jenkins model describing y_t^* is

$$y_t^* = \delta + y_{t-12}^* + \phi_1(y_{t-1}^* - y_{t-13}^*) + \phi_2(y_{t-2}^* - y_{t-14}^*) + \phi_3(y_{t-3}^* - y_{t-15}^*) - \theta_{12}\varepsilon_{t-12} + \varepsilon_t$$

The constant term δ in the tentative model says that there is a **deterministic trend** of δ as we move from y_{t-12}^*, the quartic root of the hotel room average one year ago, to y_t^*, the quartic root of the current hotel room average. In the exercises, the reader will assess the appropriateness of the tentative model (δ will be found to be important) and will use the model to forecast. (It can be verified that a constant term δ is not needed in the passenger totals model.)

FIGURE 16.27 The SAC and SPAC of the Hotel Room Seasonal Differences

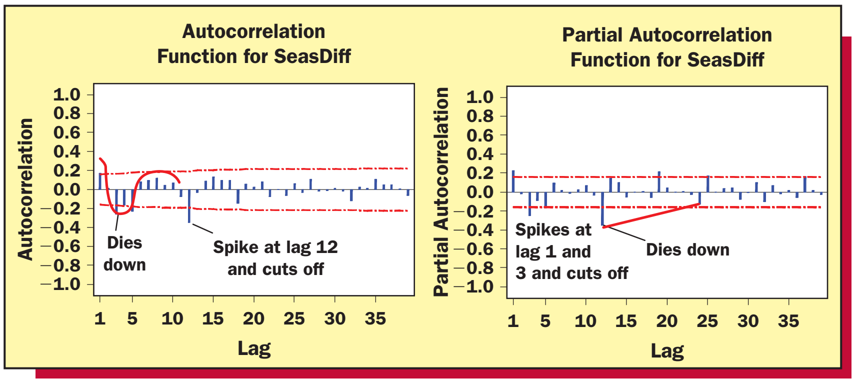

FIGURE 16.28 Forecasting (for Exercise 16.31)

Period	Forecast	95% Limits Lower	95% Limits Upper
169	5.38091	5.33070	5.43111
170	5.26784	5.21622	5.31946
171	5.27552	5.22303	5.32801
172	5.42917	5.37584	5.48249
173	5.40940	5.35592	5.46289
174	5.59285	5.53915	5.64654
175	5.82033	5.76663	5.87404
176	5.85258	5.79888	5.90629
177	5.47779	5.42406	5.53152
178	5.47843	5.42470	5.53215
179	5.28889	5.23516	5.34262
180	5.46090	5.40717	5.51463

To conclude this section, note that the reason we did not use any multiplicative terms in the hotel room average model is that the nonseasonal components $\phi_1 z_{t-1}$, $\phi_2 z_{t-2}$, and $\phi_3 z_{t-3}$ are autoregressive components and the seasonal component $-\theta_{12}\,\varepsilon_{t-12}$ is a moving average component. If the SAC at the seasonal level had told us to use the seasonal autoregressive term z_{t-12} (Guideline 4), then $-\theta_{12}\,\varepsilon_{t-12}$ would be replaced by $\phi_{12}\, z_{t-12} - \phi_1 \phi_{12}\, z_{t-13} - \phi_2 \phi_{12}\, z_{t-14} - \phi_3 \phi_{12}\, z_{t-15}$. This model uses three multiplicative terms, which are appropriate when *both* the nonseasonal and seasonal components of the model are autoregressive.

Exercises for Section 16.6

CONCEPTS

16.29 Discuss the purpose of differencing.

16.30 Explain how we use the SAC and SPAC.

METHODS AND APPLICATIONS

16.31 The least squares point estimates of the parameters of the tentative hotel room average model are $\hat{\delta} = .038056\ (.000)$, $\hat{\phi}_1 = .2392\ (.003)$, $\hat{\phi}_2 = .1322\ (.104)$, $\hat{\phi}_3 = -.2642\ (.001)$, and $\hat{\theta}_{12} = .5271\ (.000)$, where the MINITAB calculated p-values for the importance of the model parameters are given in parentheses. The chi-square statistic p-values for the diagnostic checking of model adequacy are .242 (lag 12), .186 (lag 24), .276 (lag 36), and .461 (lag 48). (1) Assess the appropriateness of the tentative model. (2) Use the model to write out an expression for y^*_{169}. (3) Hand calculate the point forecast of y^*_{169} to be $\hat{y}^*_{169}\,(168) = 5.38091$ (see Figure 16.28). Hint: Recalling that $y^*_t = y_t^{.25}$, use the facts that $y^*_{157} = 5.33648$, $y^*_{168} - y^*_{156} = .10212$, $y^*_{167} - y^*_{155} = .07934$, $y^*_{166} - y^*_{154} = -.00628$, and the residual for period 157 is .057297.

16.32 Find a point forecast and a 95 percent prediction interval for y_{169}.

Note: For an exercise on tentative identification, see Supplemental Exercise 16.44 on page 675.

16.7 Forecast Error Comparisons

LO16-9 Compare time series models by using forecast errors.

Consider comparing the forecasting accuracies of two particular forecasting methods—Method 1 and Method 2. To do this, we use each method to compute point forecasts $(\hat{y}_t)$ of the future values (y_t) of a time series. We then wait for the future values to occur, record them, and compute the forecast errors $(y_t - \hat{y}_t)$. Specifically, suppose that the results we obtain when we use Method 1 and Method 2 to forecast 12 future values of a time series are as shown in Tables 16.14 and 16.15 on the next page. Three criteria by which to compare Method 1 and Method 2 are the **mean absolute deviation (MAD),** the **mean squared deviation (MSD),** and the **mean absolute percentage error (MAPE).**

TABLE 16.14 Method 1 Forecast Errors

y_t	$\hat{y}_t$	$y_t - \hat{y}_t$
352	360.52	−8.52
445	441.48	3.52
453	446.40	6.6
541	516.62	24.38
457	433.85	23.15
762	767.82	−5.82
1,194	1,156.30	37.7
1,361	1,350.50	10.5
1,615	1,606.30	8.7
1,059	1,067.40	−8.4
824	850.12	−26.12
495	501.39	−6.39

TABLE 16.15 Method 2 Forecast Errors

y_t	$\hat{y}_t$	$y_t - \hat{y}_t$
352	355.96	−3.96
445	426.31	18.69
453	436.69	16.31
541	505.60	35.4
457	431.71	25.29
762	748.35	13.65
1,194	1,132.57	61.43
1,361	1,313.74	47.26
1,615	1,576.09	38.91
1,059	1,043.59	15.41
824	834.24	−10.24
495	506.65	−11.65

To calculate the MAD, we find the absolute value of each forecast error and then average the resulting absolute values. For example, if we find the absolute value of each of the 12 forecast errors given by Method 1 in Table 16.14, sum the 12 absolute values, and divide the sum by 12, we find that the MAD is 14.15. By contrast, if we calculate the MAD of the Method 2 forecast errors in Table 16.15, we find that the MAD is 25.6.

To calculate the MSD, we find the squared value of each forecast error and then average the resulting squared values. For example, if we find the squared value of each of the 12 forecast errors given by Method 1 in Table 16.14, sum the 12 squared values, and divide the sum by 12, we find that the MSD is 307.80. By contrast, if we calculate the MSD of the Method 2 forecast errors in Table 16.15, we find that the MSD is 892.44.

To calculate the MAPE, we find the percentage error for each forecast, $[(y_t - \hat{y}_t)/y_t] \times 100\%$, and we then average the absolute values of the percentage errors. For example, the percentage error for the first forecast given by Method 1 is $[(y_1 - \hat{y}_1)/y_1] \times 100\% = [(352 - 360.52)/352] \times 100\% = -2.42\%$. If we average the absolute values of the 12 percentage errors given by Method 1, we find that the MAPE is 2.06%. By contrast, if we calculate the MAPE of the Method 2 forecast errors, we find that the MAPE is 3.24%.

In general, we want a forecasting method that gives small values of the MAD, MSD, and MAPE. Note, however, that the MSD is the average of the *squared forecast errors.* It follows that the MSD, unlike the MAD and MAPE, penalizes a forecasting method much more for large forecast errors than for small forecast errors. Therefore, the forecasting method that gives the smallest MSD may not be the forecasting method that gives the smallest MAD, or MAPE. Furthermore, the forecaster who uses the MSD to choose a forecasting method would prefer several smaller forecast errors to one large error. In our example, Method 1 gives smaller values of the MAD, MSD, and MAPE than does Method 2. Moreover, for those who have studied Sections 16.3 and 16.5, the point forecasts in Tables 16.14 and 16.15 are point forecasts given by the multiplicative decomposition method (Method 1) and multiplicative Winters' method (Method 2). These point forecasts are point forecasts of future Tasty Cola sales in months 37 through 48. Although the multiplicative decomposition method does better in this situation, the multiplicative Winters' method would do better in other situations.

Exercises for Section 16.7

CONCEPTS

16.33 What is the MAD? What is the MSD? What is the MAPE? How do we use these quantities?

16.34 Why does the MSD penalize a forecasting method much more for large forecast errors than for small forecast errors?

METHODS AND APPLICATIONS

Exercises 16.35 and 16.36 compare two forecasting methods—method A and method B. Suppose that method A gives the point forecasts 57, 61, and 70 of three future time series values. Method B gives the point forecasts 59, 65, and 73 of these three future values. The three future values turn out to be 60, 64, and 67.

16.35 Calculate the MAD, MSD, and MAPE for method A. Calculate the MAD, MSD, and MAPE for method B.

16.36 Which method—method A or method B—gives the smallest MAD? The smallest MSD? The smallest MAPE?

16.8 Index Numbers

LO16-10 Use index numbers to compare economic data over time.

We often wish to compare a value of a time series relative to another value of the time series. For instance, according to the U.S. Bureau of Labor Statistics, energy prices increased by 4.7 percent from 1995 to 1996, while apparel prices decreased by .2 percent from 1995 to 1996. In order to make such comparisons, we must describe the time series. We have seen (in Section 16.3) that time series decomposition can be employed to describe a time series. Another way to describe time-related data is to use *index numbers.*

When we compare time series values to the same previous value, we say that the previous value is in the **base time period,** and successive comparisons of time series values to the value in the base period form a sequence of **index numbers.** More formally, a **simple index number** (or **simple index**) is defined as follows:

A **simple index** is obtained by dividing the current value of a time series by the value of the time series in the base time period and by multiplying this ratio by 100. That is, if y_t denotes the current value and if y_0 denotes the value in the base time period, then the **simple index number** is

$$\frac{y_t}{y_0} \times 100$$

The time series values used to construct an index are often *quantities* or *prices.* For instance, in Table 16.16 we give the price of a gallon of regular gasoline in the United States (in dollars) for the years 2003 through 2011. If we consider 2003 to be the base year, we compute an index for each succeeding year by dividing the price per gallon for each year by 1.59 (the price per gallon for the base year 2003) and by multiplying by 100. For example, for 2007 the simple index is $(2.80/1.59) \times 100 = 176.10$, while the simple index for 2008 is $(3.27/1.59) \times 100 = 205.66$. Table 16.16 gives the remaining index values for 2003 through 2011. Notice that (by definition) the index for the base year will always equal 100.0 (as it does here).

Although the simple index is not written with a percentage sign, comparisons of the index with the base year are percentage comparisons. For instance, the index of 205.66 for 2008 tells us that the price per gallon in 2008 was up 105.66 percent compared to the 2003 base year. In general, *if we are comparing the index to the base year,* the difference between the index and 100 gives the percentage change from the base year. It is important to point out that other period-to-period percentage comparisons cannot be made by subtracting indexes. For instance, the percentage difference between the prices per gallon in 2007 and 2008 is *not* $205.66 - 176.10 = 29.56$ percent. Rather, the percentage difference is

$$\frac{205.66 - 176.10}{176.10} \times 100 = 16.79$$

This says that the price per gallon in 2008 was up 16.79 percent relative to 2007.

A simple index is computed by using the values of one time series. Often, however, we compute an index based on the accumulated values of more than one time series. Such an index is called an **aggregate index.** As an example, food prices are often compared with an aggregate index based on a "market basket" of commonly bought grocery items. For instance, consider a market basket consisting of six items—a gallon of whole milk, a one pound jar of peanut butter, a pound of red delicious apples, a dozen eggs, a pound loaf of white bread, and a pound of ground beef. Table 16.17 on the next page lists average city prices for these items in June 2006 and in June 2011 according to the Consumer Price Index Detailed Report.

One way to compare prices would be to compute a simple index for each individual item in the market basket. However, we can create an aggregate price index by totaling the prices for each year and by then computing a simple index of the yearly price totals. Using the data in

TABLE 16.16 Price of a Gallon of Regular Gasoline (in Dollars): 2003 to 2011
DS PriceGas1

Year	2003	2004	2005	2006	2007	2008	2009	2010	2011
Price per Gallon	1.59	1.88	2.30	2.59	2.80	3.27	2.35	2.78	3.63
Index (Base Year = 2003)	100.0	118.24	144.65	162.89	176.10	205.66	147.80	174.84	228.30

Source: U.S. Energy Information Administration.

TABLE 16.17 2006 and 2011 Prices for a Market Basket of Grocery Items DS MkBskt

Grocery Item	2006 Price	2011 Price
1 gal. of whole milk	\$3.00	\$3.62
1 lb. jar of peanut butter	\$1.74	\$1.96
1 lb. of red delicious apples	\$1.05	\$1.32
1 dozen eggs	\$1.24	\$1.68
1 lb. loaf of white bread	\$1.07	\$1.49
1 lb. of ground beef	\$2.24	\$2.77
Totals	\$10.34	\$12.84

Source: http://www.bls.gov/cpi/cpid1106.pdf

Table 16.17, we obtain $(12.84/10.34) \times 100 = 124.18$. This index tells us that prices of the market basket grocery items in 2011 have increased by 24.18 percent over the prices of these items in the base year 2006. Notice that this percentage increase does not necessarily apply to each individual grocery item, nor does this index necessarily apply to any of the individual grocery items. It applies only to the aggregate of grocery items in the market basket.

In general, we compute an aggregate price index as follows:

An **aggregate price index** is

$$\left(\frac{\sum p_t}{\sum p_0}\right) \times 100$$

where Σp_t is the sum of the prices in the current time period and Σp_0 is the sum of the prices in the base year.

A disadvantage of this aggregate price index is that it does not take into account the fact that some items in the market basket are purchased more frequently than others. To remedy this deficiency, we can weight each price by the quantity of that item purchased in a given period (say yearly). Then we can total the weighted prices for each year and compute a simple index of the weighted price totals. To illustrate, Table 16.18 gives the 2006 and 2011 prices of the market basket items and also gives estimates of the quantity of each item purchased in a year by a typical family. The table also gives the price multiplied by the quantity for each item, which is simply the total yearly cost of purchasing the item. These costs are totaled for each year. Looking at Table 16.18, we see that a typical family in 2006 spent \$824.09 purchasing the market basket items during the year, while the family spent \$1049.68 purchasing the market basket items during 2011. We now compute a simple index of the total costs, which is $(1049.68/824.09) \times 100 = 127.37$.

This type of index is called a **weighted aggregate price index.** Two versions of this kind of index are commonly used. The first version is called a **Laspeyres index.** Here the quantities that are specified for the base year are also employed for all succeeding time periods. This is the assumption we have made in Table 16.18. Notice that the quantities for 2011 are the same as those specified for 2006. In general,

A **Laspeyres index** is

$$\frac{\sum p_t q_0}{\sum p_0 q_0} \times 100$$

where p_0 represents a base period price, q_0 represents a base period quantity, and p_t represents a current period price.

Because the Laspeyres index employs the base period quantities in all succeeding time periods, this index allows for ready comparison of prices for identical quantities of goods purchased. Such an index is useful as long as the base quantities provide a reasonable representation of consumption patterns in succeeding time periods. However, sometimes purchasing patterns can change drastically as consumer preferences change or as dramatic price changes occur. If

TABLE 16.18 2006 and 2011 Prices and Quantities for a Market Basket of Grocery Items DS MkBskt

	2006 (Base Year)			2011		
Grocery Item	**Price, p_0**	**Quantity, q**	**$p_0 \times q$ = cost**	**Price, p_t**	**Quantity, q**	**$p_t \times q$ = cost**
1 gal. of whole milk	\$3.00	52	\$156.00	\$3.62	52	\$188.24
1 lb. jar of peanut butter	\$1.74	13	\$22.62	\$1.96	13	\$25.48
1 lb. of red delicious apples	\$1.05	55	\$57.75	\$1.32	55	\$72.60
1 dozen eggs	\$1.24	72	\$89.28	\$1.68	72	\$120.96
1 lb. loaf of white bread	\$1.07	156	\$166.92	\$1.49	156	\$232.44
1 lb. of ground beef	\$2.24	148	\$331.52	\$2.77	148	\$409.96
Totals	\$10.34		\$824.09	\$12.84		\$1049.68

TABLE 16.19 2006 and 2011 Prices and 2011 Quantities for a Market Basket of Grocery Items DS MkBsktR

Grocery Item	**2006 Price, p_0**	**2011 Quantity, q_t**	**$p_0 \times q_t$ = cost**	**2011 Price, p_t**	**2011 Quantity, q_t**	**$p_t \times q_t$ = cost**
1 gal. of whole milk	\$3.00	40	\$120.00	\$3.62	40	\$144.80
1 lb. jar of peanut butter	\$1.74	21	\$36.54	\$1.96	21	\$41.16
1 lb. of red delicious apples	\$1.05	82	\$86.10	\$1.32	82	\$108.24
1 dozen eggs	\$1.24	60	\$74.40	\$1.68	60	\$100.80
1 lb. loaf of white bread	\$1.07	175	\$187.25	\$1.49	175	\$260.75
1 lb. of ground beef	\$2.24	148	\$331.52	\$2.77	148	\$409.96
Totals	\$10.34	Total	\$835.81	\$12.84	Total	\$1065.71

consumption patterns in the current period are very different from the quantities specified in the base period, then a Laspeyres index can be misleading because it relates to quantities of goods that few people would purchase.

A second version of the weighted aggregate price index is called a **Paasche index.** Here we update the quantities so that they reflect consumption patterns in the current time period.

A **Paasche index** is

$$\frac{\sum p_t q_t}{\sum p_0 q_t} \times 100$$

where p_0 represents a base period price, p_t represents a current period price, and q_t represents a current period quantity.

As an example, Table 16.19 presents revised quantities for the grocery items in our previously discussed market basket. These quantities reflect increased consumption of apples, peanut butter, and bread, and decreased consumption of milk and eggs. We calculate a 2006 cost of \$835.81 for the items in the market basket and a 2011 cost of \$1065.71 for the items in the market basket. Therefore, the Paasche index is $(1065.71/835.81) \times 100 = 127.51$.

Because the Paasche index uses quantities from the current period, it reflects current buying habits. However, quantity data for the current period can be difficult to obtain. Furthermore, although each period is compared to the base period, it is difficult to compare the index at other points in time. This is because different quantities are used in different periods, and thus changes in the index are affected by changes in both prices and quantities.

Economic indexes Several commonly quoted economic indexes are compiled monthly by the U.S. Bureau of Labor Statistics. Two important indexes are the **Consumer Price Index** (the **CPI**) and the **Producer Price Index** (the **PPI**). These are both *Laspeyres indexes.*

The CPI monitors the price of a market basket of goods and services that would be purchased by typical nonfarm consumers. Actually, there are two Consumer Price Indexes. The CPI-U, the

FIGURE 16.29 Plot of the Monthly CPI-U from February 1947 until April 2011

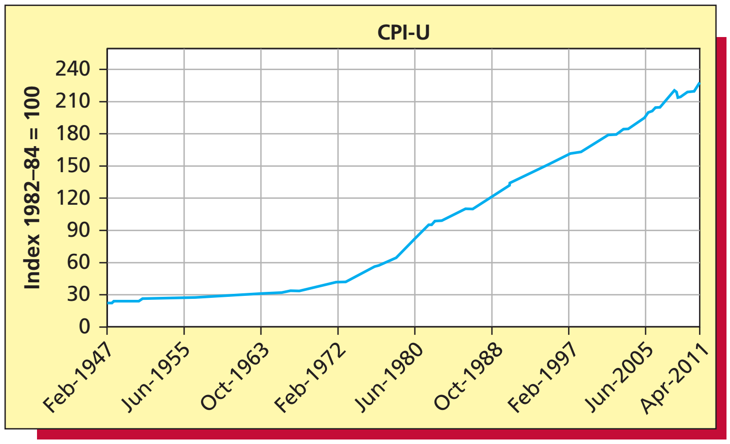

Source: www.data360.org/pdf_print_group.aspx?Print_Group_Id=169.

Consumer Price Index for all Urban Workers, is often reported by the press as an indicator of price changes. Figure 16.29, which gives a plot of the monthly CPI-U from February 1947 until April 2011, shows the general increasing trend in prices over this period. The U.S. Bureau of the Census periodically changes the base period for the CPI. The plot in Figure 16.29 uses a base period of 1982–1984. Here, assume that the base period cost used to compute the CPI-U is the average of the 36 monthly market basket costs over the 1982–1984 base period. The April 2011 CPI-U, computed with a 1982–1984 base index of 100, was 224.91. This says that the cost of purchasing the market basket of goods and services was 124.91 percent higher in April 2011 than it was in the 1982–1984 base period. A second CPI, the CPI-W (Consumer Price Index for Urban Wage Earners and Clerical Workers) is often used to determine wage increases that are written into labor contracts.

The PPI tracks the prices of goods sold by wholesalers. An increase in the PPI is often regarded as an indication that retail prices will soon rise.

Exercises for Section 16.8

connect

CONCEPTS

16.37 Explain the difference between a simple index and an aggregate index.

16.38 Explain the difference between a Laspeyres index and a Paasche index.

METHODS AND APPLICATIONS

16.39 Referring to the discussion of Figure 16.29, the CPI-U for March 2011 was 223.47, and the CPI-U for October 1963 was 30. Interpret these CPI-U's.

16.40 Recall that Table 16.16 (page 669) gives the price of a gallon of regular gasoline in the United States (in dollars) for the years 2003 through 2011. Table 16.20 gives the price of a gallon of regular gasoline in the United States (in dollars) for the years 1990 through 2002. DS PriceGas2

a By using 1990 as the base year, construct a simple index for the prices of a gallon of regular gasoline in Tables 16.20 and 16.16.

b Plot the indices you have constructed versus time (1990 through 2011). Describe the pattern that you see.

16.41 Suppose that Table 16.21 gives the yearly automobile operating expenses in a particular region of the United States for the years 1990, 2000, and 2011. DS AutoExp

a Using 1990 as the base year, construct the Laspeyres index for these operating expenses. Describe how the operating expenses have changed over time.

TABLE 16.20 Gas Prices
DS PriceGas2

Year	Price per Gallon ($)
1990	1.16
1991	1.14
1992	1.13
1993	1.11
1994	1.11
1995	1.15
1996	1.23
1997	1.23
1998	1.06
1999	1.17
2000	1.51
2001	1.46
2002	1.36

TABLE 16.21 Yearly Automobile Operating Expenses in 1990, 2000, and 2011 DS Autoexp

	1990		2000		2011	
Item	**Price**	**Quantity**	**Price**	**Quantity**	**Price**	**Quantity**
Gallon of gasoline	\$1.16	1000	\$1.51	882	\$3.63	682
Quart of oil	\$2.00	15	\$3.50	15	\$5.75	15
Tire	\$125.00	2	\$132.00	2	\$142.00	2
Insurance	\$750.00	1	\$900.00	1	\$1,100.00	1

Note: The quantities are based on 15,000 miles driven per year and a 30,000 mile tire lifetime.

b Using 1990 as the base year, construct the Paasche index for the operating expenses.

c Another index, the **value index,** is given by the formula $(\Sigma p_t\, q_t / \Sigma p_0\, q_0) \times 100$. This index measures the changes in both the prices and quantities involved. Using 1990 as the base year, construct the value index for the operating expenses. Compare the three indices you have constructed.

Chapter Summary

In this chapter we have discussed using **univariate time series** models to forecast future time series values. We began by seeing that it can be useful to think of a time series as consisting of **trend, seasonal, cyclical, and irregular components.** If these components remain *constant* over time, then it is appropriate to describe and forecast the time series by using a **time series regression model.** We discussed using such models to describe **no trend,** a **linear trend,** and **constant seasonal variation** (by utilizing dummy variables). We also considered various transformations that transform **increasing seasonal variation** into constant seasonal variation. As an alternative to using a transformation and dummy variables to model increasing seasonal variation, we discussed using the **multiplicative decomposition method.** We then turned to a consideration of **exponential smoothing,** which is appropriate to use if the components of a time series may be *changing slowly* over time. Specifically, we discussed **simple exponential smoothing, Holt–Winters' double exponential smoothing, and multiplicative Winters' method.** We next considered how to use the Box–Jenkins methodology to forecast a time series having components that may be *changing fairly quickly* over time. We also explained how to compare forecasting methods by using the **mean absolute deviation (MAD),** the **mean squared deviation (MSD),** and the **mean absolute percentage error (MAPE).** We concluded this chapter by showing how to use **index numbers** to describe time-related data. Our discussion included the construction of a **simple index**, an **aggregate index**, a **Laspeyres index**, and a **Paasche index.**

Glossary of Terms

autoregressive terms: Past time series values used in Box–Jenkins models. (page 664)

cyclical variation: Recurring up-and-down movements of a time series around trend levels that last more than one calendar year (often 2 to 10 years) from peak to peak or trough to trough. (page 631)

deseasonalized time series: A time series that has had the effect of seasonal variation removed. (page 645)

double exponential smoothing: An exponential smoothing procedure that can be used to forecast a time series described by a linear trend model with parameters that may be slowly changing over time. (pages 652–654)

exponential smoothing: A forecasting method that weights recent observations more heavily than distantly past observations. (page 648)

index number: A number that compares a value of a time series relative to another value of the time series. (pages 668–672)

irregular component: What is "left over" in a time series after trend, cycle, and seasonal variations have been accounted for. (page 631)

MAD: The mean of the absolute values of a set of forecast errors. (pages 667–668)

MAPE: The mean of the absolute values of the percentage errors of a set of forecasts. (pages 667–668)

MSD: The mean of the squares of a set of forecast errors. (pages 667–668)

moving averages: Averages of successive groups of time series observations. (page 641)

moving average terms: Past error terms used in Box–Jenkins models. (page 665)

multiplicative Winters' method: An exponential smoothing procedure that can be used to forecast a time series described by a linear trend and increasing (or decreasing) seasonal variation with parameters that may be slowly changing over time. (pages 654–659)

seasonal variation: Periodic patterns in a time series that repeat themselves within a calendar year and are then repeated yearly. (page 631)

simple exponential smoothing: An exponential smoothing procedure that can be used to forecast a time series described by a no trend model with an average level that may be slowly changing over time. (pages 647–651)

smoothing constant: A number that determines how much weight is attached to each observation when using exponential smoothing. (page 648)

time series: A set of observations that has been collected in time order. (page 631)

trend: The long-run upward or downward movement that characterizes a time series over a period of time. (page 631)

univariate time series model: A model that predicts future values of a time series solely on the basis of past values of the time series. (page 631)

Important Formulas and Tests

No trend: page 632
Linear trend: page 633
Modeling constant seasonal variation by using dummy variables: pages 634–637
The multiplicative decomposition model: pages 640–646
Simple exponential smoothing: page 650
Double exponential smoothing: pages 652–654
Multiplicative Winters' method: pages 654–659
Mean absolute deviation (MAD): pages 654, 667–668
Mean absolute percentage error (MAPE): pages 667–668
Mean squared deviation (MSD): pages 654, 667–668
Simple index: page 669
Aggregate price index: page 670
Laspeyres index: page 670
Paasche index: page 671

Supplementary Exercises

connect

16.42 The State University Credit Union, a savings institution open to the faculty and staff of State University, handles savings accounts and makes loans to members. In order to plan its investment strategies, the credit union requires both point and prediction interval forecasts of monthly loan requests (in thousands of dollars) to be made by the faculty and staff in future months. The credit union has recorded monthly loan requests for its past two years of operation. These loan requests are as follows: DS Loans

Year 1	297	249	340	406	464	481	549	553	556	642	670	712
Year 2	808	809	867	855	965	921	956	990	1019	1021	1033	1127

If we use MINITAB to fit the **quadratic trend model**

$$y_t = \beta_0 + \beta_1 t + \beta_2 t^2 + \varepsilon_t$$

to these data, we obtain the following partial MINITAB output.

```
The regression equation is
Y = 200 + 50.9 Time - 0.568 TimeSQ

Predictor        Coef      SE Coef         T          P
Constant       199.62        20.85      9.58      0.000
Time           50.937        3.842     13.26      0.000
TimeSQ        -0.5677       0.1492     -3.80      0.001

S = 31.2469   R-Sq = 98.7%   R-Sq(adj) = 98.6%

Predicted Values for New Observations

New Obs  Time  TimeSQ      Fit  SE Fit       95% CI              95% PI
   1     25.0     625  1118.21   20.85  (1074.85, 1161.56)  (1040.09, 1196.32)
   2     26.0     676  1140.19   24.44  (1089.37, 1191.01)  (1057.70, 1222.68)
```

a Does the quadratic term t^2 seem important in the model? Justify your answer.

b At the bottom of the MINITAB output are point and prediction interval forecasts of loan requests in months 25 and 26. **(1)** Find and report these forecasts. **(2)** Then (using the least squares point estimates of β_0, β_1, and β_2 on the computer output) calculate the point forecasts.

16.43 The Olympia Paper Company, Inc., makes Absorbent Paper Towels. The company would like to develop a prediction model that can be used to give point forecasts and prediction interval forecasts of weekly sales over 100,000 rolls, in units of 10,000 rolls, of Absorbent Paper Towels. With a reliable model, Olympia Paper can more effectively plan its production schedule, plan its budget, and estimate requirements for producing and storing this product. For the past 120 weeks the company has recorded weekly sales of Absorbent Paper Towels. The 120 sales figures, $y_1, y_2, \ldots y_{120}$, are given in Table 16.22, and the MINITAB output resulting from using simple exponential smoothing to forecast future values of the sales figures is given in Figure 16.30, on the next page. Note that MINITAB chooses a smoothing constant of 1.35442, which is not (as is the usual case) between 0 and 1. In general, if the trend effects in a time series are changing fairly quickly (as the plot of the sales values in Figure 16.30 indicates is true) MINITAB will often choose a smoothing constant greater than 1. Use the MINITAB output to find a point forecast of and a 95 percent prediction interval for the number of rolls of Absorbent Paper Towels that will be sold in a future week.

TABLE 16.22 Absorbent Paper Towels Sales DS Towels

y_t	y_t	y_t
15.0000	9.2835	11.4986
14.4064	7.7219	13.2778
14.9383	6.8300	13.5910
16.0374	8.2046	13.4297
15.6320	8.5289	13.3125
14.3975	8.8733	12.7445
13.8959	8.7948	11.7979
14.0765	8.1577	11.7319
16.3750	7.9128	11.6523
16.5342	8.7978	11.3718
16.3839	9.0775	10.5502
17.1006	9.3234	11.4741
17.7876	10.4739	11.5568
17.7354	10.6943	11.7986
17.0010	9.8367	11.8867
17.7485	8.1803	11.2951
18.1888	7.2509	12.7847
18.5997	5.0814	13.9435
17.5859	1.8313	13.6859
15.7389	−0.9127	14.1136
13.6971	−1.3173	13.8949
15.0059	−0.6021	14.2853
16.2574	0.1400	16.3867
14.3506	1.4030	17.0884
11.9515	1.9280	15.8861
12.0328	3.5626	14.8227
11.2142	1.9615	15.9479
11.7023	4.8463	15.0982
12.5905	6.5454	13.8770
12.1991	8.0141	14.2746
10.7752	7.9746	15.1682
10.1129	8.4959	15.3818
9.9330	8.4539	14.1863
11.7435	8.7114	13.9996
12.2590	7.3780	15.2463
12.5009	8.1905	17.0179
11.5378	9.9720	17.2929
9.6649	9.6930	16.6366
10.1043	9.4506	15.3410
10.3452	11.2088	15.6453

FIGURE 16.30 MINITAB Output of Simple Exponential Smoothing for the Absorbent Paper Towels Sales Time Series DS Towels

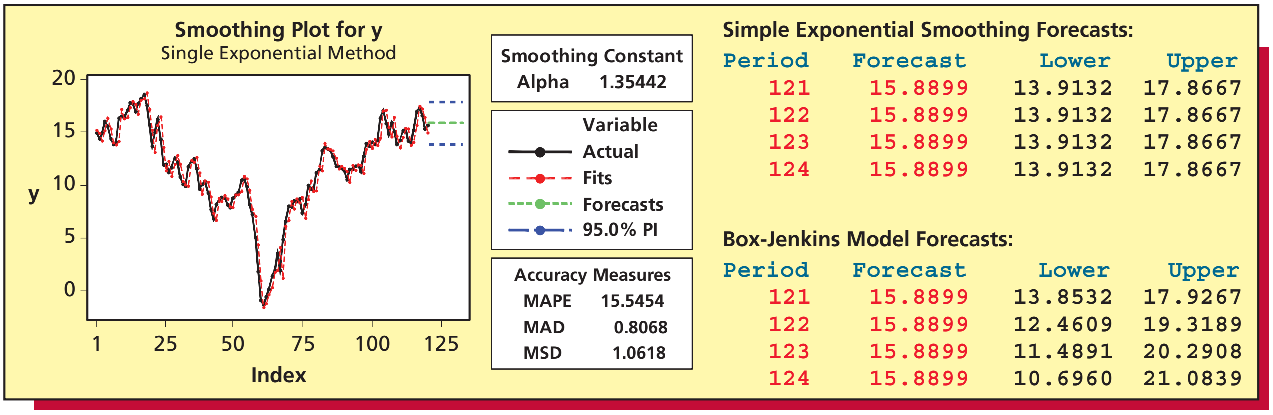

16.44 Figure 16.30 shows that the sales values in Table 16.22 are nonstationary and nonseasonal. At the nonseasonal level (lags 1 through 9), the SAC of the regular differences of the sales values has a spike at lag 1 and cuts off after lag 1, and the SPAC of these regular differences dies down fairly quickly. Discuss why the guidelines in Section 16.6 say that a tentative Box–Jenkins model describing the regular differences $z_t = y_t - y_{t-1}$ is $z_t = \varepsilon_t - \theta_1 \varepsilon_{t-1}$ and thus that a tentative model describing y_t is $y_t = y_{t-1} + \varepsilon_t - \theta_1 \varepsilon_{t-1}$. Estimation shows that θ_1 is important in the tentative model and that a constant term δ does not belong in the model. Diagnostic checking shows that the tentative model is adequate. Because the adequate Box–Jenkins model can be shown to be equivalent to simple exponential smoothing, it gives the same point forecasts. However, MINITAB calculates the Box–Jenkins model prediction intervals differently and obtains increasingly wide intervals (see Figure 16.30).

Appendix 16.1 ■ Time Series Analysis Using Excel

Point forecasts from a linear trend line for the calculator sales data in Table 16.2 on page 633 (data file: CalcSale.xlsx):

- Enter the calculator sales data from Table 16.2 with the label "Sales" in cell A1 and the values of sales in cells A2 through A25.
- Enter the label "Month" in cell B1 and the values 1 to 28 in cells B2 through B29.
- Click on cell A26.
- Click the Insert Function button f_x on the Excel ribbon.
- In the Insert Function dialog box, select Statistical from the "Or select a category:" menu and select TREND from the "Select a function:" menu. Then click OK in the Insert Function dialog box.
- In the "TREND Function Arguments" dialog box, enter $A$2:$A$25 into the "Known_y's" window. Don't forget the dollar signs—this must be an absolute cell reference.
- Enter $B$2:$B$25 into the "Known_x's" window. Again, don't forget the dollar signs.
- Enter B26 into the "New_x's" window.
- Enter the value 1 into the Const window.

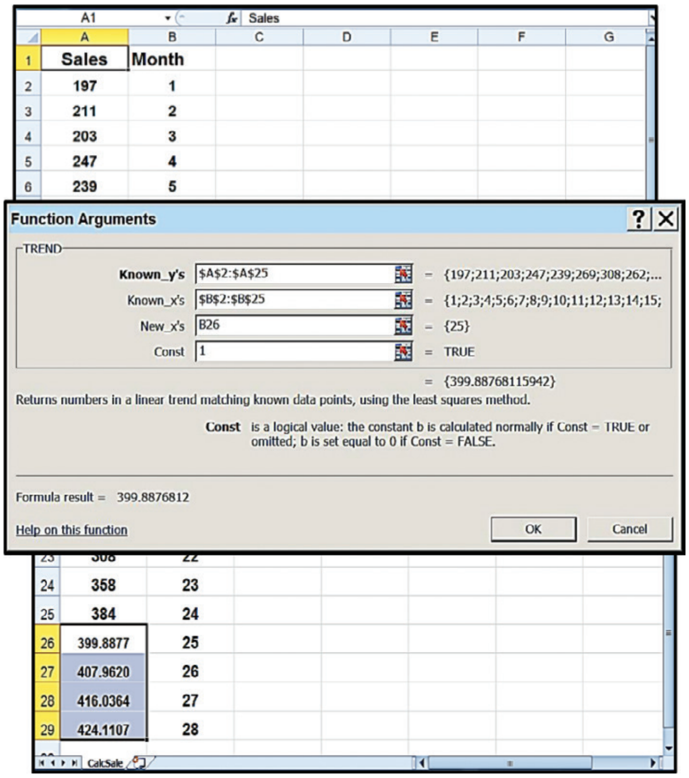

(*Continues on the next page.*)

- Click OK in the "TREND Function Arguments" dialog box to produce the point forecast for time period 25.
- Double-click on the drag handle in cell A26 to extend the forecasts through time period 28.

Appendix 16.2 ■ Time Series Analysis Using MegaStat

Calculation of seasonal factors and deseasonalization similar to Table 16.10, Table 16.11, and Figure 16.11 on pages 641 and 643 (data file: TastyCola.xlsx):

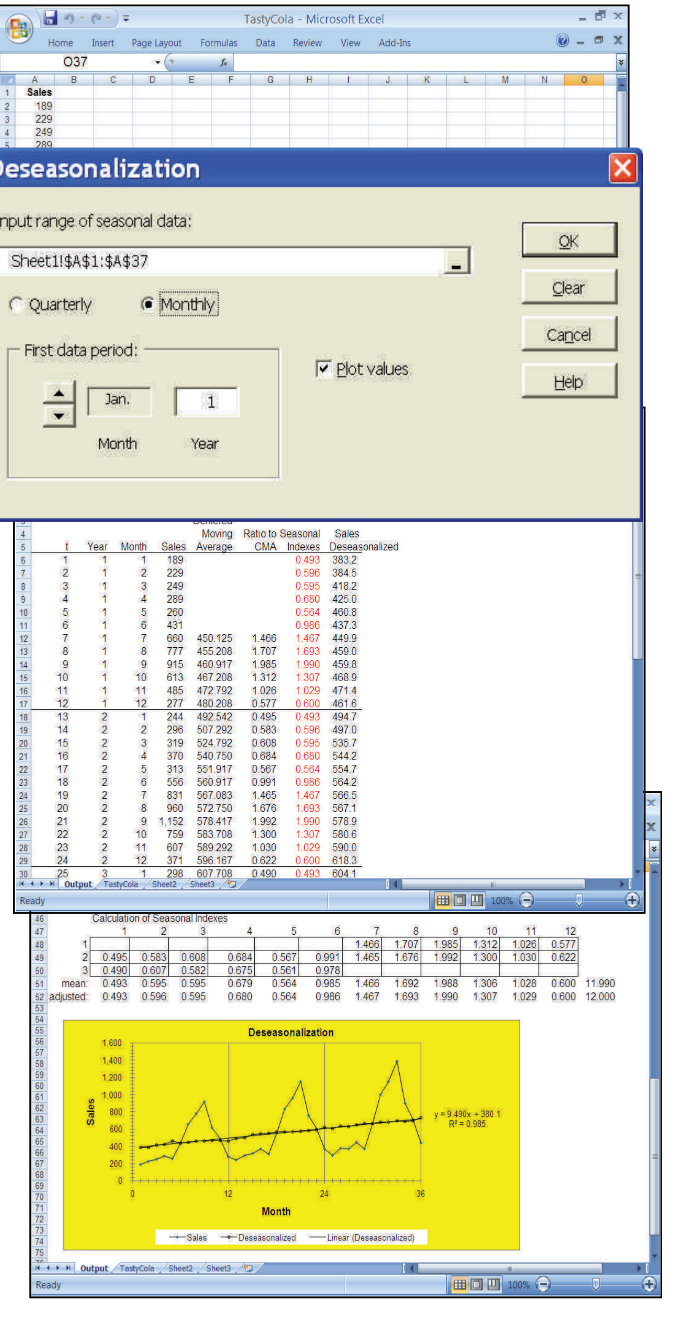

- Enter the Tasty Cola data in Table 16.9 (page 640) into column A with label Sales. Only the sales values in Table 16.9 need to be entered—the year, month, and time period need not be entered.
- Select **Add-Ins : MegaStat : Time Series/ Forecasting : Deseasonalization.**
- In the Deseasonalization dialog box, enter the range A1:A37 into the "Input Range of Seasonal Data" window. This range can be entered by dragging with the mouse—the AutoExpand feature cannot be used in this dialog box.
- Select the type of seasonal data—"quarterly" or "monthly"—by clicking. Here we have selected "monthly" because the Tasty Cola data consists of monthly sales values.
- In the "First data period" box, specify the month (in this case, January) in which the first time series value was observed by using the up or down arrow buttons.
- In the "First data period" box, enter the year in which the first time series value was observed (here equal to 1) into the Year box.
- Check the Plot Values checkbox to obtain plots of the seasonal observations, the deseasonalized data, and a trend line fit to the deseasonalized data.
- Click OK in the Deseasonalization dialog box.
- The seasonal factors are displayed in the "Seasonal Indexes" column of the "Centered Moving Average and Deseasonalization" table in the output worksheet. They are also given in the "adjusted" row at the bottom of the "Calculation of Seasonal Indexes" table in the output worksheet.

Simple exponential smoothing similar to Table 16.13 on page 649 (data file: CodCatch.xlsx):

- Enter the cod catch data in Table 16.1 (page 632) into column A with label CodCatch.
- Select **Add-Ins : MegaStat : Time Series/ Forecasting : Exponential Smoothing : Simple Exponential Smoothing.**
- In the Simple Exponential Smoothing dialog box, enter the range A1:A25 into the "Input Range for Data" window. Enter this range by dragging with the mouse—the autoexpand feature cannot be used in this dialog box.
- Type the value of the smoothing constant (here equal to .02) into the Alpha window.
- Leave the Initial Value window blank if you wish to use an initial value equal to the average of the first six time series observations. If another initial value is desired, type it into the Initial Value window.
- Click OK in the Simple Exponential Smoothing dialog box.
- The forecast for a future value of the time series is found at the bottom of the "Forecast" column in the output worksheet.

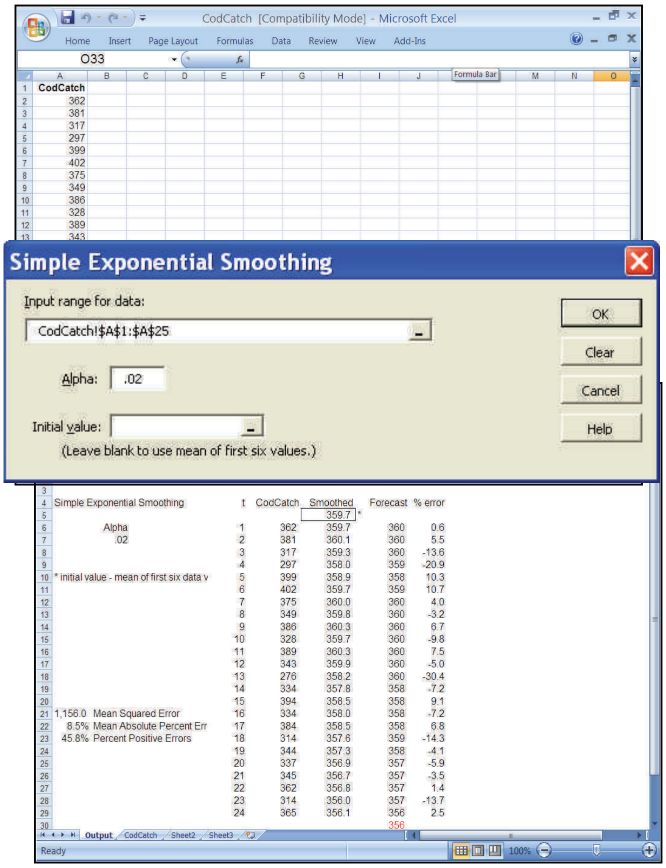

Double exponential smoothing similar to Figure 16.15 on page 653 (data file: CalcSale.xlsx):

- Enter the calculator sales data in Table 16.2 on page 633 into column A with label Sales.
- Select **Add-Ins : MegaStat : Time Series/ Forecasting : Exponential Smoothing : Two-factor Exponential Smoothing.**
- In the Two-Factor Exponential Smoothing dialog box, enter the range A1:A25 into the "Input Range for Data" window. Enter this range by dragging with the mouse—the AutoExpand feature cannot be used in this dialog box.
- Type the desired values of the smoothing constants (here both are set equal to .20) into the Alpha and Beta boxes.
- Leave the "Initial Value" and "Initial Trend" boxes blank if you wish to use initial values that are estimated by the computer using the first six time series observations. If you wish to supply initial values, type an initial value of the intercept into the "Initial Value" box and type an initial value of the slope into the "Initial Trend" box. Here we have supplied the values 198.0 and 8.1.
- Click OK in the Two-Factor Exponential Smoothing dialog box.
- The forecast for the next time series value is found at the bottom of the Forecast column in the output worksheet.

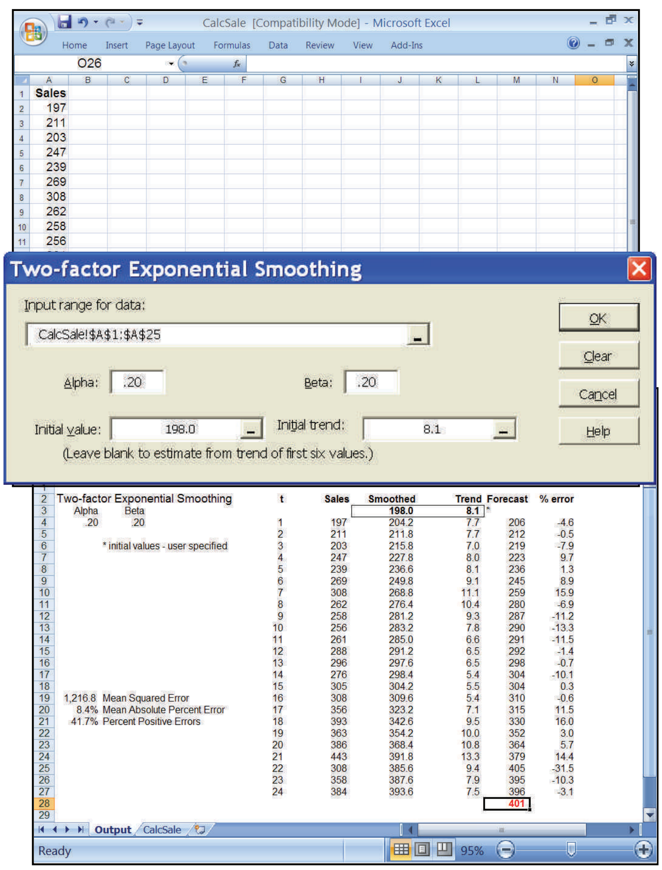

Appendix 16.3 ■ Time Series Analysis Using MINITAB

Simple exponential smoothing in Figure 16.13 on page 651 (data file: CodCatch.MTW):

- In the Data window, enter the cod catch data from Table 16.1 on page 632 into column C1 with variable name CodCatch.
- Select **Stat : Time Series : Single Exp Smoothing.**
- In the Single Exponential Smoothing dialog box, enter CodCatch in the Variable window.
- To request that MINITAB select the smoothing constant, select the "Optimal ARIMA" option under "Weight to Use in Smoothing." To choose your own smoothing constant, select the "Use" option and enter the desired smoothing constant in the window.
- Place a checkmark in the "Generate forecasts" checkbox.
- Enter 12 in the "Number of forecasts" window and enter 24 in the "Starting from origin" window.
- Click OK in the Single Exponential Smoothing dialog box to see the forecast results in the Session window and a graphical summary in a high-resolution graphics window.

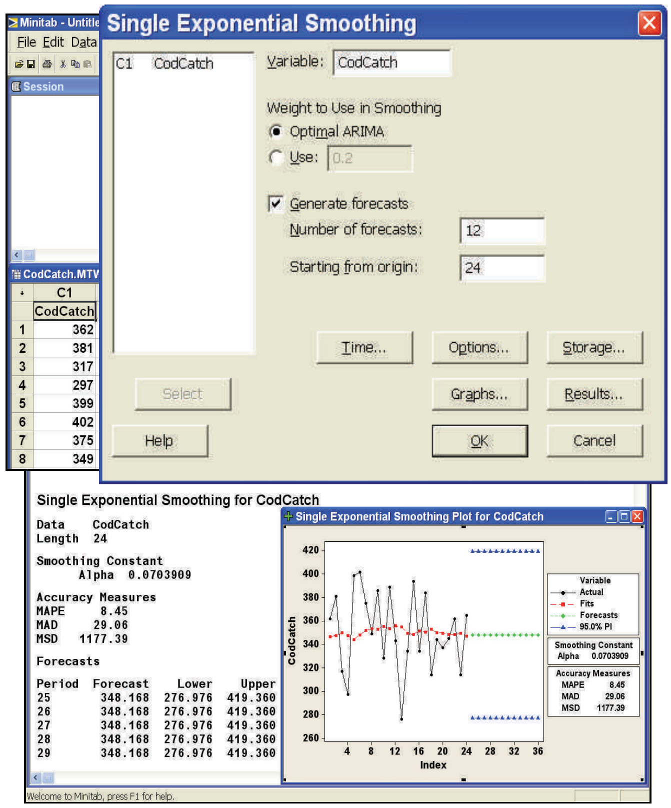

Double exponential smoothing can be performed by choosing **Double Exp Smoothing** from the Time Series menu and by following the remainder of the preceding steps.

Multiplicative Winters' method in Figure 16.19 on page 658 (data file: TastyCola.MTW):

- In the Data window, enter the Tasty Cola data from Table 16.9 (page 640) into column C1 with variable name Sales.
- Select **Stat : Time Series : Winters' Method.**
- In the Winters' Method dialog box, enter Sales into the Variable window.
- Enter 12 in the "Seasonal length" window.
- Click the Multiplicative option under Method Type.
- Use the default values for "Weights to Use in Smoothing" (0.2 in each of the Level, Trend, and Seasonal windows).
- Click the "Generate forecasts" checkbox.
- Enter 12 in the "Number of forecasts" window and enter 36 in the "Starting from origin" window.
- Click OK in the Winters' Method dialog box to obtain the forecast results in the Session window and a graphical summary in a high-resolution graphics window.

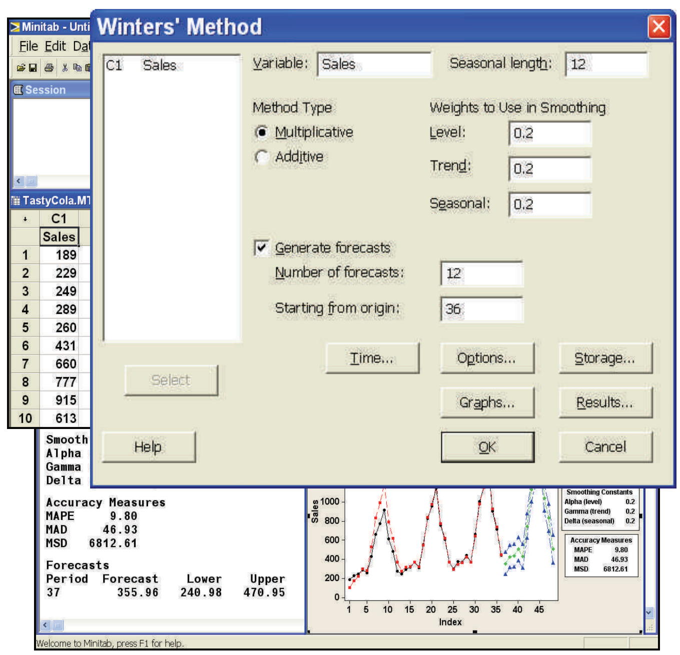

Computing the SAC and SPAC of the combined regular and seasonal differences of the logged passenger totals in Figure 16.25 on page 664 (data file: AirPass.MTW):

- Select **Calc : Calculator** and calculate the natural logarithms of the passenger totals (with variable name LnPassengers) as shown in the calculator dialog box to the right.
- Select **Stat : Time Series : Differences**
- In the Differences dialog box, enter LnPassengers into the Series window.
- Enter 'Seasonal Diff' into the "Store differences in" window.
- Enter 12 into the Lag window.
- Click OK to calculate the first order seasonal differences of the logged passenger totals.
- Select **Stat : Time Series : Differences**
- In the Differences dialog box, enter 'Seasonal Diff' into the series window.
- Enter 'RegSeasonal Diff' into the "Store differences in" window.
- Enter 1 into the Lag window.
- Click OK to calculate the first order regular differences of the first order seasonal differences of the logged passenger totals.
- Select **Stat : Time Series : Autocorrelation**
- In the Autocorrelation Function dialog box, enter 'RegSeasonal Diff' into the series window.
- Enter 24 into the "Number of lags" window.
- Click OK in the Autocorrelation Function dialog box to obtain a graph of the sample autocorrelation function of the combined regular and seasonal differences of the logged passenger totals.
- In a similar fashion, obtain a graph of the sample partial autocorrelation function by selecting **Stat : Time Series : Partial Autocorrelation**

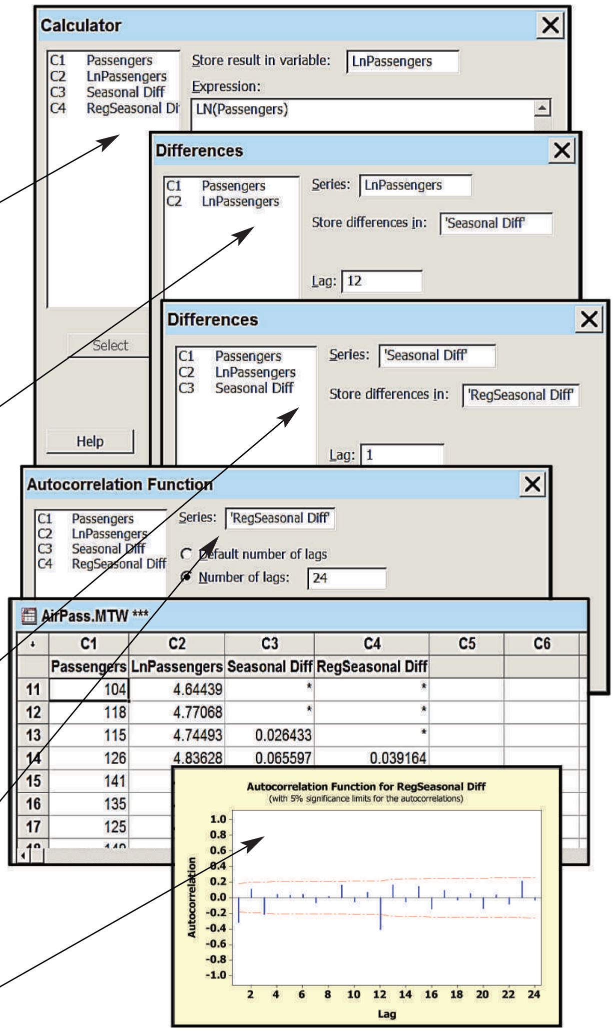

Performing estimation, diagnostic checking, and forecasting using the passenger totals Box–Jenkins model as in Figure 16.24 on page 663 (data file: AirPass.MTW):

- Select **Stat : Time Series : ARIMA**
- Enter "LnPassengers" into the Series window.
- Place a checkmark in the "Fit seasonal model" box and enter 12 into the Period window to specify monthly data.
- Enter 1 into each of the Nonseasonal and Seasonal Difference windows (first order regular and first order seasonal differencing).
- Enter 0 into each of the Nonseasonal Autoregressive and Seasonal Autoregressive windows (the model has no autoregressive terms).
- Enter 1 into each of the Nonseasonal Moving average and Seasonal Moving average windows. (These specify the first-order nonseasonal and first-order seasonal moving average terms.)
- Be sure that the "Include constant term in model" checkbox is not checked (no constant in model).

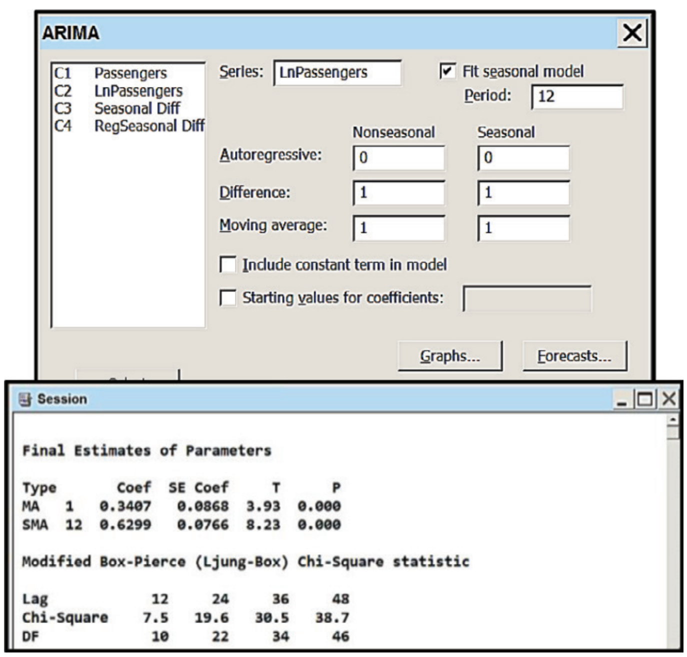

- Click on the Forecasts... button and enter 12 into the "Lead" window to forecast 12 months into the future.
- Click OK in the ARIMA dialog box.