

27

CHAPTER

Forecasting

How much will the economy grow over the next year? Where is the stock market headed? What about interest rates? How will consumer tastes be changing? What will be the hot new products?

Forecasters have answers to all these questions. Unfortunately, these answers will more than likely be wrong. Nobody can accurately predict the future every time.

Nevertheless, the future success of any business depends heavily on how savvy its management is in spotting trends and developing appropriate strategies. The leaders of the best companies often seem to have a sixth sense for when to change direction to stay a step ahead of the competition. These companies seldom get into trouble by badly misestimating what the demand will be for their products. Many other companies do. The ability to forecast well makes the difference.

The preceding chapter has presented a considerable number of models for the management of inventories. All these models are based on a forecast of future demand for a product, or at least a probability distribution for that demand. Therefore, the missing ingredient for successfully implementing these inventory models is an approach for forecasting demand.

Fortunately, when historical sales data are available, some proven **statistical forecasting methods** have been developed for using these data to forecast future demand. Such a method assumes that historical trends will continue, so management then needs to make any adjustments to reflect current changes in the marketplace.

Several **judgmental forecasting methods** that solely use expert judgment also are available. These methods are especially valuable when little or no historical sales data are available or when major changes in the marketplace make these data unreliable for forecasting purposes.

Forecasting product demand is just one important application of the various forecasting methods. A variety of applications are surveyed in the first section. The second section outlines the main judgmental forecasting methods. Section 27.3 then describes *time series*, which form the basis for the statistical forecasting methods presented in the subsequent six sections (including a software package presented in Sec. 27.7 that incorporates these methods). Section 27.10 turns to another important type of statistical forecasting method, *regression analysis*, where the variable to be forecasted is expressed as a mathematical function of one or more other variables whose values will be known at the time of the forecast. The chapter then concludes by surveying forecasting practices in U.S. corporations.

27.1 SOME APPLICATIONS OF FORECASTING

We now will discuss some main areas in which forecasting is widely used. In each case, we will illustrate this use by mentioning one or more actual applications that have been described in published articles. A summary table at the end of the section will tell you where these articles can be found in case you want to read further.

Sales Forecasting

Any company engaged in selling goods needs to forecast the demand for those goods. Manufacturers need to know how much to produce. Wholesalers and retailers need to know how much to stock. Substantially underestimating demand is likely to lead to many lost sales, unhappy customers, and perhaps allowing the competition to gain the upper hand in the marketplace. On the other hand, significantly overestimating demand also is very costly due to (1) excessive inventory costs, (2) forced price reductions, (3) unneeded production or storage capacity, and (4) lost opportunities to market more profitable goods. Successful marketing and production managers understand very well the importance of obtaining good sales forecasts.

The *Merit Brass Company* is a family-owned company that supplies several thousand products to the pipe, valve, and fittings industry. In 1990, Merit Brass embarked on a modernization program that emphasized installing OR methodologies in statistical sales forecasting and finished-goods inventory management (two activities that go hand in glove). This program led to major improvements in customer service (as measured by product availability) while simultaneously achieving substantial cost reductions.

A major Spanish electric utility, *Hidroeléctrica Español*, has developed and implemented a hierarchy of OR models to assist in managing its system of reservoirs used for generating hydroelectric power. All these models are driven by forecasts of both energy demand (this company's sales) and reservoir inflows. A sophisticated statistical forecasting method is used to forecast energy demand on both a short-term and long-term basis. A hydrological forecasting model generates the forecasts of reservoir inflows.

Airline companies now depend heavily on the high fares paid by business people traveling on short notice while providing discount fares to others to help fill the seats. The decision on how to allocate seats to the different fare classes is a crucial one for maximizing revenue. *American Airlines*, for example, uses statistical forecasting of the demand at each fare to make this decision.

Forecasting the Need for Spare Parts

Although effective sales forecasting is a key for virtually any company, some organizations must rely on other types of forecasts as well. A prime example involves forecasts of the need for spare parts.

Many companies need to maintain an inventory of spare parts to enable them to quickly repair either their own equipment or their products sold or leased to customers. In some cases, this inventory is huge. For example, IBM's spare-parts inventory is valued in the billions of dollars and includes many thousand different parts.

Just as for a finished-goods inventory ready for sale, effective management of a spare-parts inventory depends upon obtaining a reliable forecast of the demand for that inventory. Although the types of costs incurred by misestimating demand are somewhat different, the consequences may be no less severe for spare parts. For example, the consequence for an airline not having a spare part available on location when needed to continue flying an airplane probably is at least one canceled flight.

To support its operation of several hundred aircraft, *American Airlines* maintains an extensive inventory of spare parts. Included are over 5,000 different types of rotatable parts (e.g., landing gear and wing flaps) with an average value of \$5,000 per item. When a rotatable part on an airplane is found to be defective, it is immediately replaced by a corresponding part in inventory so the airplane can depart. However, the replaced part then is repaired and placed back into inventory for subsequent use as a replacement part.

American Airlines uses a PC-based forecasting system called the Rotatables Allocation and Planning System (RAPS) to forecast demand for the rotatable parts and to help allocate these parts to the various airports. The statistical forecast uses an 18-month history of parts usage and flying hours for the fleet, and then projects ahead based on planned flying hours.

Forecasting Production Yields

The yield of a production process refers to the percentage of the completed items that meet quality standards (perhaps after rework) and so do not need to be discarded. Particularly with high-technology products, the yield frequently is well under 100 percent.

If the forecast for the production yield is somewhat under 100 percent, the size of the production run probably should be somewhat larger than the order quantity to provide a good chance of fulfilling the order with acceptable items. (The difference between the run size and the order quantity is referred to as the *reject allowance*.) If an expensive setup is required for each production run, or if there is only time for one production run, the reject allowance may need to be quite large. However, an overly large value should be avoided to prevent excessive production costs.

Obtaining a reliable forecast of production yield is essential for choosing an appropriate value of the reject allowance.

This was the case for the *Albuquerque Microelectronics Operation*, a dedicated production source for radiation-hardened microchips. The first phase in the production of its microchips, the wafer fabrication process, was continuing to provide erratic production yields. For a given product, the yield typically would be quite small (0 to 40 percent) for the first several lots and then would gradually increase to a higher range (35 to 75 percent) for later lots. Therefore, a statistical forecasting method that considered this increasing trend was used to forecast the production yield.

Forecasting Economic Trends

With the possible exception of sales forecasting, the most extensive forecasting effort is devoted to forecasting economic trends on a regional, national, or even international level. How much will the nation's gross domestic product grow next quarter? Next year? What is the forecast for the rate of inflation? The unemployment rate? The balance of trade?

Statistical models to forecast economic trends (commonly called *econometric models*) have been developed in a number of governmental agencies, university research centers, large corporations, and consulting firms, both in the United States and elsewhere. Using historical data to project ahead, these econometric models typically consider a very large number of factors that help drive the economy. Some models include hundreds of variables and equations. However, except for their size and scope, these models resemble some of the statistical forecasting methods used by businesses for sales forecasting, etc.

These econometric models can be very influential in determining governmental policies. For example, the forecasts provided by the U.S. Congressional Budget Office strongly guide Congress in developing the federal budgets. These forecasts also help businesses in assessing the general economic outlook.

As an example on a smaller scale, the *U.S. Department of Labor* contracted with a consulting firm to develop the *unemployment insurance econometric forecasting model*

(UIEFM). The model is now in use by state employment security agencies around the nation. By projecting such fundamental economic factors as unemployment rates, wage levels, the size of the labor force covered by unemployment insurance, etc., UIEFM forecasts how much the state will need to pay in unemployment insurance. By projecting tax inflows into the state's unemployment insurance trust fund, UIEFM also forecasts trust fund balances over a 10-year period. Therefore, UIEFM has proved to be invaluable in managing state unemployment insurance systems and in guiding related legislative policies.

Forecasting Staffing Needs

One of the major trends in the American economy is a shifting emphasis from manufacturing to services. More and more of our manufactured goods are being produced outside the country (where labor is cheaper) and then imported. At the same time, an increasing number of American business firms are specializing in providing a service of some kind (e.g., travel, tourism, entertainment, legal aid, health services, financial, educational, design, maintenance, etc.). For such a company, forecasting "sales" becomes forecasting the demand for services, which then translates into forecasting staffing needs to provide those services.

For example, one of the fastest-growing service industries in the United States today is call centers. A call center receives telephone calls from the general public requesting a particular type of service. Depending on the center, the service might be providing technical assistance over the phone, or making a travel reservation, or filling a telephone order for goods, or booking services to be performed later, etc. There now are several hundred thousand call centers in the United States.

As with any service organization, an erroneous forecast of staffing requirements for a call center has serious consequences. Providing too few agents to answer the telephone leads to unhappy customers, lost calls, and perhaps lost business. Too many agents cause excessive personnel costs.

For example, consider a major OR study that involved personnel scheduling at *United Airlines*. With over 4,000 reservations sales representatives and support personnel at its 11 reservations offices, and about 1,000 customer service agents at its 10 largest airports, a computerized planning system was developed to design the work schedules for these employees. Although several other OR techniques (including linear programming) were incorporated into this system, statistical forecasting of staffing requirements also was a key ingredient. This system provided annual savings of over \$6 million as well as improved customer service and reduced support staff requirements.

Taco Bell is an international fast-food corporation that operates mostly through approximately 4,000 company-owned and franchised restaurants. The corporation has applied a variety of operations research techniques to develop a labor-management system that each of the restaurants uses to develop work schedules and assignments for its employees. A key component of the labor-management system is *statistical forecasting* of the number of customers arriving in each 15-minute interval throughout the hours of a restaurant's operation. This system is credited with saving \$13 million in labor costs while simultaneously improving the level of customer service.

L.L. Bean is a major retailer of high-quality outdoor goods and apparel. Over 70 percent of its total sales volume is generated through orders taken at the company's call center. Two 800 numbers are provided, one for placing orders and the second for making inquiries or reporting problems. Each of the company's agents is trained to answer just one of the 800 numbers. Therefore, separate statistical forecasting models were developed to forecast staffing requirements for the two 800 numbers on a weekly basis. The improved precision of these models is estimated to have saved L.L. Bean \$300,000 annually through enhanced scheduling efficiency.

■ **TABLE 27.1** Some applications of statistical forecasting methods

Organization	Quantity Being Forecasted	Issue of <i>Interfaces</i>
Merit Brass Co.	Sales of finished goods	Jan.–Feb. 1993
Hidroeléctrica Español	Energy demand	Jan.–Feb. 1990
American Airlines	Demand for different fare classes	Jan.–Feb. 1992
American Airlines	Need for spare parts to repair airplanes	July–Aug. 1989
Albuquerque Microelectronics	Production yield in wafer fabrication	March–April 1994
U.S. Department of Labor	Unemployment insurance payments	March–April 1988
United Airlines	Demand at reservations offices and airports	Jan.–Feb. 1986
Taco Bell	Number of customer arrivals	Jan.–Feb. 1998
L.L. Bean	Staffing needs at call center	Nov.–Dec. 1995

Other

Table 27.1 summarizes the actual applications of statistical forecasting methods presented in this section. The last column cites the issue of *Interfaces* which includes the article that describes each application in detail.

All five categories of forecasting applications discussed in this section use the types of forecasting methods presented in the subsequent sections. There also are other important categories (including forecasting weather, the stock market, and prospects for new products before market testing) that use specialized techniques that are not discussed here.

■ 27.2 JUDGMENTAL FORECASTING METHODS

Judgmental forecasting methods are, by their very nature, subjective, and they may involve such qualities as intuition, expert opinion, and experience. They generally lead to forecasts that are based upon qualitative criteria.

These methods may be used when no data are available for employing a statistical forecasting method. However, even when good data are available, some decision makers prefer a judgmental method instead of a formal statistical method. In many other cases, a combination of the two may be used.

Here is a brief overview of the main judgmental forecasting methods.

- 1. Manager's opinion:** This is the most informal of the methods, because it simply involves a single manager using his or her best judgment to make the forecast. In some cases, some data may be available to help make this judgment. In others, the manager may be drawing solely on experience and an intimate knowledge of the current conditions that drive the forecasted quantity.
- 2. Jury of executive opinion:** This method is similar to the first one, except now it involves a small group of high-level managers who pool their best judgment to collectively make the forecast. This method may be used for more critical forecasts for which several executives share responsibility and can provide different types of expertise.
- 3. Sales force composite:** This method is often used for sales forecasting when a company employs a sales force to help generate sales. It is a *bottom-up approach* whereby each salesperson provides an estimate of what sales will be in his or her region. These estimates then are sent up through the corporate chain of command, with managerial review at each level, to be aggregated into a corporate sales forecast.
- 4. Consumer market survey:** This method goes even further than the preceding one in adopting a *grass-roots approach* to sales forecasting. It involves surveying customers and potential customers regarding their future purchasing plans and how they would

respond to various new features in products. This input is particularly helpful for designing new products and then in developing the initial forecasts of their sales. It also is helpful for planning a marketing campaign.

5. **Delphi method:** This method employs a panel of experts in different locations who independently fill out a series of questionnaires. However, the results from each questionnaire are provided with the next one, so each expert then can evaluate this group information in adjusting his or her responses next time. The goal is to reach a relatively narrow spread of conclusions from most of the experts. The decision makers then assess this input from the panel of experts to develop the forecast. This involved process normally is used only at the highest levels of a corporation or government to develop long-range forecasts of broad trends.

The decision on whether to use one of these judgmental forecasting methods should be based on an assessment of whether the individuals who would execute the method have the background needed to make an informed judgment. Another factor is whether the expertise of these individuals or the availability of relevant historical data (or a combination of both) appears to provide a better basis for obtaining a reliable forecast.

The next eight sections discuss statistical forecasting methods based on relevant historical data.

27.3 TIME SERIES


Most statistical forecasting methods are based on using historical data from a *time series*.

A **time series** is a series of observations over time of some quantity of interest (a random variable). Thus, if X_i is the random variable of interest at time i , and if observations are taken at times¹ $i = 1, 2, \dots, t$, then the observed values $\{X_1 = x_1, X_2 = x_2, \dots, X_t = x_t\}$ are a time series.


For example, the recent monthly sales figures for a product comprises a time series, as illustrated in Fig. 27.1.

¹These times of observation sometimes are actually time periods (months, years, etc.), so we often will refer to the times as periods.

FIGURE 27.1
The evolution of the monthly sales of a product illustrates a time series.



■ **FIGURE 27.2**
Typical time series patterns, with random fluctuations around (a) a constant level, (b) a linear trend, and (c) a constant level plus seasonal effects.



Because a time series is a description of the past, a logical procedure for forecasting the future is to make use of these historical data. If the past data are indicative of what we can expect in the future, we can postulate an underlying mathematical model that is representative of the process. The model can then be used to generate forecasts.

In most realistic situations, we do not have complete knowledge of the exact form of the model that generates the time series, so an approximate model must be chosen. Frequently, the choice is made by observing the pattern of the time series. Several typical time series patterns are shown in Fig. 27.2. Figure 27.2a displays a typical time series if the generating process were represented by a **constant level** superimposed with random fluctuations. Figure 27.2b displays a typical time series if the generating process were represented by a **linear trend** superimposed with random fluctuations. Finally, Fig. 27.2c shows a time series that might be observed if the generating process were represented by a constant level superimposed with a **seasonal effect** together with random fluctuations. There are many other plausible representations, but these three are particularly useful in practice and so are considered in this chapter.

Once the form of the model is chosen, a mathematical representation of the generating process of the time series can be given. For example, suppose that the generating process is identified as a **constant-level model** superimposed with random fluctuations, as illustrated in Fig. 27.2a. Such a representation can be given by

$$X_i = A + e_i, \quad \text{for } i = 1, 2, \dots,$$

where X_i is the random variable observed at time i , A is the constant level of the model, and e_i is the random error occurring at time i (assumed to have expected value equal to zero and constant variance). Let

$$F_{t+1} = \text{forecast of the values of the time series at time } t + 1, \text{ given the observed values, } X_1 = x_1, X_2 = x_2, \dots, X_t = x_t.$$

Because of the random error e_{t+1} , it is impossible for F_{t+1} to predict the value $X_{t+1} = x_{t+1}$ precisely, but the goal is to have F_{t+1} estimate the constant level $A = E(X_{t+1})$ as closely as possible. It is reasonable to expect that F_{t+1} will be a function of at least some of the observed values of the time series.

■ 27.4 FORECASTING METHODS FOR A CONSTANT-LEVEL MODEL

We now present four alternative forecasting methods for the constant-level model introduced in the preceding paragraph. This model, like any other, is only intended to be an idealized representation of the actual situation. For the real time series, at least small shifts

in the value of A may be occurring occasionally. Each of the following methods reflects a different assessment of how recently (if at all) a significant shift may have occurred.

Last-Value Forecasting Method

By interpreting t as the *current time*, the last-value forecasting procedure uses the value of the time series observed at time t (x_t) as the forecast at time $t + 1$. Therefore,

$$F_{t+1} = x_t.$$

For example, if x_t represents the sales of a particular product in the quarter just ended, this procedure uses these sales as the forecast of the sales for the next quarter.

This forecasting procedure has the disadvantage of being imprecise; i.e., its variance is large because it is based upon a sample of size 1. It is worth considering only if (1) the underlying assumption about the constant-level model is “shaky” and the process is changing so rapidly that anything before time t is almost irrelevant or misleading or (2) the assumption that the random error e_t has constant variance is unreasonable and the variance at time t actually is much smaller than at previous times.

The last-value forecasting method sometimes is called the **naive method**, because statisticians consider it naive to use just a *sample size of one* when additional relevant data are available. However, when conditions are changing rapidly, it may be that the last value is the only relevant data point for forecasting the next value under current conditions. Therefore, decision makers who are anything but naive do occasionally use this method under such circumstances.

Averaging Forecasting Method

This method goes to the other extreme. Rather than using just a sample size of one, this method uses *all* the data points in the time series and simply *averages* these points. Thus, the forecast of what the next data point will turn out to be is

$$F_{t+1} = \sum_{i=1}^t \frac{x_i}{t}.$$

This estimate is an excellent one if the process is entirely stable, i.e., if the assumptions about the underlying model are correct. However, frequently there exists skepticism about the persistence of the underlying model over an extended time. Conditions inevitably change eventually. Because of a natural reluctance to use very old data, this procedure generally is limited to young processes.

Moving-Average Forecasting Method

Rather than using very old data that may no longer be relevant, this method averages the data for only the last n periods as the forecast for the next period, i.e.,

$$F_{t+1} = \sum_{i=t-n+1}^t \frac{x_i}{n}.$$

Note that this forecast is easily updated from period to period. All that is needed each time is to lop off the first observation and add the last one.

The *moving-average* estimator combines the advantages of the *last value* and *averaging* estimators in that it uses only recent history *and* it uses multiple observations. A disadvantage of this method is that it places as much weight on x_{t-n+1} as on x_t . Intuitively, one would expect a good method to place more weight on the most recent observation

than on older observations that may be less representative of current conditions. Our next method does just this.

Exponential Smoothing Forecasting Method

This method uses the formula

$$F_{t+1} = \alpha x_t + (1 - \alpha)F_t,$$

where α ($0 < \alpha < 1$) is called the **smoothing constant**. (The choice of α is discussed later.) Thus, the forecast is just a weighted sum of the last observation x_t and the preceding forecast F_t for the period just ended. Because of this recursive relationship between F_{t+1} and F_t , alternatively F_{t+1} can be expressed as

$$F_{t+1} = \alpha x_t + \alpha(1 - \alpha)x_{t-1} + \alpha(1 - \alpha)^2 x_{t-2} + \cdots.$$

In this form, it becomes evident that exponential smoothing gives the most weight to x_t and decreasing weights to earlier observations. Furthermore, the first form reveals that the forecast is simple to calculate because the data prior to period t need not be retained; all that is required is x_t and the previous forecast F_t .

Another alternative form for the exponential smoothing technique is given by

$$F_{t+1} = F_t + \alpha(x_t - F_t),$$

which gives a heuristic justification for this method. In particular, the forecast of the time series at time $t + 1$ is just the preceding forecast at time t plus the *product* of the forecasting error at time t and a discount factor α . This alternative form is often simpler to use.

A measure of effectiveness of exponential smoothing can be obtained under the assumption that the process is completely stable, so that X_1, X_2, \dots are independent, identically distributed random variables with variance σ^2 . It then follows that (for large t)

$$\text{var}[F_{t+1}] \approx \frac{\alpha\sigma^2}{2 - \alpha} = \frac{\sigma^2}{(2 - \alpha)\alpha},$$

so that the variance is statistically equivalent to a moving average with $(2 - \alpha)/\alpha$ observations. For example, if α is chosen equal to 0.1, then $(2 - \alpha)/\alpha = 19$. Thus, in terms of its variance, the exponential smoothing method with this value of α is *equivalent* to the moving-average method that uses 19 observations. However, if a change in the process does occur (e.g., if the mean starts increasing), exponential smoothing will react more quickly with better tracking of the change than the moving-average method.

An important drawback of exponential smoothing is that it lags behind a continuing trend; i.e., if the constant-level model is incorrect and the mean is increasing steadily, then the forecast will be several periods behind. However, the procedure can be easily adjusted for trend (and even seasonally adjusted).

Another disadvantage of exponential smoothing is that it is difficult to choose an appropriate smoothing constant α . Exponential smoothing can be viewed as a statistical filter that inputs raw data from a stochastic process and outputs smoothed estimates of a mean that varies with time. If α is chosen to be small, response to change is slow, with resultant smooth estimators. On the other hand, if α is chosen to be large, response to change is fast, with resultant large variability in the output. Hence, there is a need to compromise, depending upon the degree of stability of the process. It has been suggested that α should not exceed 0.3 and that a reasonable choice for α is approximately 0.1. This value can be increased temporarily if a change in the process is expected or when one is just starting the forecasting. At the start, a reasonable approach is to choose the forecast for period 2 according to

$$F_2 = \alpha x_1 + (1 - \alpha)(\text{initial estimate}),$$

where some initial estimate of the constant level A must be obtained. If past data are available, such an estimate may be the average of these data.

The Excel files for this chapter in your OR Courseware includes a pair of Excel templates for each of the four forecasting methods presented in this section. In each use, one template (*without seasonality*) applies the method just as described here. The second template (*with seasonality*) also incorporates into the method the seasonal factors discussed in the next section.

The forecasting area of your IOR Tutorial also includes procedures for applying these four forecasting methods (and others). You enter the data (after making any needed seasonal adjustment yourself), and each procedure then shows a graph that includes both the data points (in blue) and the resulting forecasts (in red) for each period. You then have the opportunity to drag any of the data points to new values and immediately see how the subsequent forecasts would change. The purpose is to allow you to play with the data and gain a better feeling for how the forecasts perform with various configurations of data for each of the forecasting methods.

27.5 INCORPORATING SEASONAL EFFECTS INTO FORECASTING METHODS

It is fairly common for a time series to have a *seasonal pattern* with higher values at certain times of the year than others. For example, this occurs for the sales of a product that is a popular choice for Christmas gifts. Such a time series violates the basic assumption of a *constant-level model*, so the forecasting methods presented in the preceding section should not be applied directly.

Fortunately, it is relatively straightforward to make *seasonal adjustments* in such a time series so that these forecasting methods based on a constant-level model can still be applied. We will illustrate the procedure with the following example.

Example. The COMPUTER CLUB WAREHOUSE (commonly referred to as CCW) sells various computer products at bargain prices by taking telephone orders directly from customers at its call center. Figure 27.3 shows the average number of calls received per day in each of the four quarters of the past three years. Note how the call volume jumps up sharply in each Quarter 4 because of Christmas sales. There also is a tendency for the call volume to be a little higher in Quarter 3 than in Quarter 1 or 2 because of back-to-school sales.

To quantify these seasonal effects, the second column of Table 27.2 shows the average daily call volume for each quarter over the past three years. Underneath this column, the *overall average* over all four quarters is calculated to be 7,529. Dividing the average for each quarter by this overall average gives the *seasonal factor* shown in the third column.

In general, the **seasonal factor** for any period of a year (a quarter, a month, etc.) measures how that period compares to the overall average for an entire year. Specifically, using historical data, the seasonal factor is calculated to be

$$\text{Seasonal factor} = \frac{\text{average for the period}}{\text{overall average}}.$$

Your OR Courseware includes an Excel template for calculating these seasonal factors.

The Seasonally Adjusted Time Series

It is much easier to analyze a time series and detect new trends if the data are first adjusted to remove the effect of seasonal patterns. To remove the seasonal effects from the time series shown in Fig. 27.3, each of these average daily call volumes

27.5 INCORPORATING SEASONAL EFFECTS INTO FORECASTING METHODS 27-11




FIGURE 27.3 The average number of calls received per day at the CCW call center in each of the four quarters of the past three years.

TABLE 27.2 Calculation of the seasonal factors for the CCW problem

Quarter	Three-Year Average	Seasonal Factor
1	7,019	$\frac{7,019}{7,529} = 0.93$
2	6,784	$\frac{6,784}{7,529} = 0.90$
3	7,434	$\frac{7,434}{7,529} = 0.99$
4	8,880	$\frac{8,880}{7,529} = 1.18$
Total = 30,117		
Average = $\frac{30,117}{4} = 7,529$.		


needs to be divided by the corresponding seasonal factor given in Table 27.2. Thus, the formula is

$$\text{Seasonally adjusted call volume} = \frac{\text{actual call volume}}{\text{seasonal factor}}$$

Applying this formula to all 12 call volumes in Fig. 27.3 gives the seasonally adjusted call volumes shown in column F of the spreadsheet in Fig. 27.4.

In effect, these seasonally adjusted call volumes show what the call volumes would have been if the calls that occur because of the time of the year (Christmas shopping, back-to-school shopping, etc.) had been spread evenly throughout the year instead. Compare the plots in Figs. 27.4 and 27.3. After considering the smaller vertical scale in Fig. 27.4, note how much less fluctuation this figure has than Fig. 27.3 because of removing seasonal effects. However, this figure still is far from completely flat because fluctuations in call

	A	B	C	D	E	F	G	H	I	J	
1	Seasonally Adjusted Time Series for CCW										
2											
3				Seasonal	Actual	Seasonally Adjusted					
4		Year	Quarter	Factor	Call Volume	Call Volume					
5		1	1	0.93	6809	7322					
6		1	2	0.90	6465	7183					
7		1	3	0.99	6569	6635					
8		1	4	1.18	8266	7005					
9		2	1	0.93	7257	7803					
10		2	2	0.90	7064	7849					
11		2	3	0.99	7784	7863					
12		2	4	1.18	8724	7393					
13		3	1	0.93	6992	7518					
14		3	2	0.90	6822	7580					
15		3	3	0.99	7949	8029					
16		3	4	1.18	9650	8178					



	F
5	=E5/D5
6	=E6/D6
7	=E7/D7
8	=E8/D8
9	:
10	:

■ **FIGURE 27.4**

The seasonally adjusted time series for the CCW problem obtained by dividing each actual average daily call volume in Fig. 27.3 by the corresponding seasonal factor obtained in Table 27.2.

volume occur for other reasons beside just seasonal effects. For example, hot new products attract a flurry of calls. A jump also occurs just after the mailing of a catalog. Some random fluctuations occur without any apparent explanation. Figure 27.4 enables seeing and analyzing these fluctuations in sales volumes that are not caused by seasonal effects.

The General Procedure

After seasonally adjusting a time series, any of the forecasting methods presented in the preceding section (or the next section) can then be applied. Here is an outline of the general procedure.

1. Use the following formula to seasonally adjust each value in the time series:

$$\text{Seasonally adjusted value} = \frac{\text{actual value}}{\text{seasonal factor}}$$

2. Select a time series forecasting method.
3. Apply this method to the seasonally adjusted time series to obtain a forecast of the next *seasonally adjusted* value (or values).
4. Multiply this forecast by the corresponding seasonal factor to obtain a forecast of the next *actual* value (without seasonal adjustment).

As mentioned at the end of the preceding section, an Excel template that incorporates seasonal effects is available in your OR Courseware for each of the forecasting methods to assist you with combining the method with this procedure.

27.6 AN EXPONENTIAL SMOOTHING METHOD FOR A LINEAR TREND MODEL

Recall that the constant-level model introduced in Sec. 27.3 assumes that the sequence of random variables $\{X_1, X_2, \dots, X_t\}$ generating the time series has a constant expected value denoted by A , where the goal of the forecast F_{t+1} is to estimate A as closely as possible. However, as was illustrated in Fig. 27.2b, some time series violate this assumption by having a continuing trend where the expected values of successive random variables keep changing in the same direction. Therefore, a forecasting method based on the constant-level model (perhaps after adjusting for seasonal effects) would do a poor job of forecasting for such a time series because it would be continually lagging behind the trend. We now turn to another model that is designed for this kind of time series.

Suppose that the generating process of the observed time series can be represented by a *linear trend* superimposed with *random fluctuations*, as illustrated in Fig. 27.2b. Denote the slope of the linear trend by B , where the slope is called the **trend factor**. The model is represented by

$$X_i = A + Bi + e_i, \quad \text{for } i = 1, 2, \dots,$$

where X_i is the random variable that is observed at time i , A is a constant, B is the trend factor, and e_i is the random error occurring at time i (assumed to have expected value equal to zero and constant variance).

For a real time series represented by this model, the assumptions may not be completely satisfied. It is common to have at least small shifts in the values of A and B occasionally. It is important to detect these shifts relatively quickly and reflect them in the forecasts. Therefore, practitioners generally prefer a forecasting method that places substantial weight on recent observations and little if any weight on old observations. The exponential smoothing method presented next is designed to provide this kind of approach.

Adapting Exponential Smoothing to This Model

The exponential smoothing method introduced in Sec. 27.4 can be adapted to include the trend factor incorporated into this model. This is done by also using exponential smoothing to estimate this trend factor.

Let

$$T_{t+1} = \text{exponential smoothing estimate of the trend factor } B \text{ at time } t + 1, \text{ given the observed values, } X_1 = x_1, X_2 = x_2, \dots, X_t = x_t.$$

Given T_{t+1} , the forecast of the value of the time series at time $t + 1$ (F_{t+1}) is obtained simply by adding T_{t+1} to the formula for F_{t+1} given in Sec. 27.4, so

$$F_{t+1} = \alpha x_t + (1 - \alpha)F_t + T_{t+1}.$$

To motivate the procedure for obtaining T_{t+1} , note that the model assumes that

$$B = E(X_{i+1}) - E(X_i), \quad \text{for } i = 1, 2, \dots$$

Thus, the standard statistical estimator of B would be the *average* of the observed differences, $x_2 - x_1, x_3 - x_2, \dots, x_t - x_{t-1}$. However, the exponential smoothing approach recognizes that the parameters of the stochastic process generating the time series (including A and B) may actually be gradually shifting over time so that the most recent observations are the most reliable ones for estimating the current parameters. Let

$$L_{t+1} = \text{latest trend at time } t + 1 \text{ based on the last two values } (x_t \text{ and } x_{t-1}) \text{ and the last two forecasts } (F_t \text{ and } F_{t-1}).$$

The exponential smoothing formula used for L_{t+1} is

$$L_{t+1} = \alpha(x_t - x_{t-1}) + (1 - \alpha)(F_t - F_{t-1}).$$

Then T_{t+1} is calculated as

$$T_{t+1} = \beta L_{t+1} + (1 - \beta)T_t,$$

where β is the **trend smoothing constant** which, like α , must be between 0 and 1. Calculating L_{t+1} and T_{t+1} in order then permits calculating F_{t+1} with the formula given in the preceding paragraph.

Getting started with this forecasting method requires making two initial estimates about the status of the time series just prior to beginning forecasting. These initial estimates are

x_0 = initial estimate of the *expected value* of the time series (A) if the conditions just prior to beginning forecasting were to remain unchanged without any trend;

T_1 = initial estimate of the *trend* of the time series (B) just prior to beginning forecasting.

The resulting forecasts for the first two periods are

$$F_1 = x_0 + T_1,$$

$$L_2 = \alpha(x_1 - x_0) + (1 - \alpha)(F_1 - x_0),$$

$$T_2 = \beta L_2 + (1 - \beta)T_1,$$

$$F_2 = \alpha x_1 + (1 - \alpha)F_1 + T_2.$$

The above formulas for L_{t+1} , T_{t+1} , and F_{t+1} then are used directly to obtain subsequent forecasts.

Since the calculations involved with this method are relatively involved, a computer commonly is used to implement the method. The Excel files for this chapter in your OR Courseware include two Excel templates (one without seasonal adjustments and one with) for this method. In addition, the forecasting area in your IOR Tutorial includes a procedure of this method that also enables you to investigate graphically the effect of making changes in the data.

Application of the Method to the CCW Example

Reconsider the example involving the Computer Club Warehouse (CCW) that was introduced in the preceding section. Figure 27.3 shows the time series for this example (representing the average daily call volume quarterly for 3 years) and then Fig. 27.4 gives the seasonally adjusted time series based on the seasonal factors calculated in Table 27.2. We now will assume that these seasonal factors were determined *prior* to these three years of data and that the company then was using *exponential smoothing with trend* to forecast the average daily call volume quarter by quarter over the 3 years based on these data. CCW management has chosen the following initial estimates and smoothing constants:

$$x_0 = 7,500, \quad T_1 = 0, \quad \alpha = 0.3, \quad \beta = 0.3.$$

Working with the seasonally adjusted call volumes given in Fig. 27.4, these initial estimates lead to the following seasonally adjusted forecasts.

$$\text{Y1, Q1: } F_1 = 7,500 + 0 = 7,500.$$

$$\text{Y1, Q2: } L_2 = 0.3(7,322 - 7,500) + 0.7(7,500 - 7,500) = -53.4.$$

$$T_2 = 0.3(-53.4) + 0.7(0) = -16.$$

$$F_2 = 0.3(7,322) + 0.7(7,500) - 16 = 7,431.$$

$$\text{Y1, Q3: } L_3 = 0.3(7,183 - 7,322) + 0.7(7,431 - 7,500) = -90.$$

$$T_3 = 0.3(-90) + 0.7(-16) = -38.2.$$

$$F_3 = 0.3(7,183) + 0.7(7,431) - 38.2 = 7,318.$$

⋮

27.6 AN EXPONENTIAL SMOOTHING METHOD

Exponential Smoothing with Trend Forecasting Method with Seasonality for CCW										
Year	Quarter	True Value	Seasonally Adjusted Value	Latest Trend	Estimated Trend	Seasonally Adjusted Forecast	Actual Forecast	Forecasting Error	Smoothing Constant	
1	1	6,809	7,322		0	7,500	6,975	166	$\alpha =$	0.3
7	1	6,465	7,183	-54	-16	7,430	6,687	222	$\beta =$	0.3
8	1	6,569	6,635	-90	-38	7,318	7,245	676	Initial Estimate	
9	1	8,266	7,005	-243	-100	7,013	8,276	10	Average =	7,500
10	2	7,257	7,803	-102	-100	6,910	6,427	830	Trend =	0
11	2	7,064	7,849	167	-20	7,158	6,442	622	Type of Seasonality	
12	2	7,784	7,863	187	42	7,407	7,333	451	Quarterly	
13	2	8,724	7,393	179	83	7,627	9,000	276	Seasonal Factor	
14	3	6,992	7,518	13	62	7,619	7,085	93	Quarter	1
15	3	6,822	7,580	32	53	7,642	6,877	55	2	0.93
16	3	7,949	8,029	34	47	7,670	7,594	355	3	0.90
17	3	9,650	8,178	155	80	7,858	9,272	378	4	0.99
18	4		#N/A	176	108	8,062	7,498			1.18
19	4		#N/A			#N/A				

Mean Absolute Deviation	
MAD =	345
Mean Square Error	
MSE =	180,796

	E	F	G	H	I	J
3	Seasonally			Seasonally		
4	Adjusted	Latest	Estimated	Adjusted	Actual	Forecasting
5	Value	Trend	Trend	Forecast	Forecast	Error
6	=D6/M16		=InitialEstimateTrend	=InitialEstimateAverage+InitialEstimateTrend	=M16*H6	=ABS(D6-I6)
7	=D7/M17	=Alpha*(E6-InitialEstimateAverage)+(1-Alpha)*(H6-InitialEstimateAverage)	=Beta*F7+(1-Beta)*G6	=Alpha*E6+(1-Alpha)*H6+G7	=M17*H7	=ABS(D7-I7)
8	=D8/M18	=Alpha*(E7-E6)+(1-Alpha)*(H7-H6)	=Beta*F8+(1-Beta)*G7	=Alpha*E7+(1-Alpha)*H7+G8	=M18*H8	=ABS(D8-I8)
9	=D9/M19	=Alpha*(E8-E7)+(1-Alpha)*(H8-H7)	=Beta*F9+(1-Beta)*G8	=Alpha*E8+(1-Alpha)*H8+G9	=M19*H9	=ABS(D9-I9)
10	=D10/M16	=Alpha*(E9-E8)+(1-Alpha)*(H9-H8)	=Beta*F10+(1-Beta)*G9	=Alpha*E9+(1-Alpha)*H9+G10	=M16*H10	=ABS(D10-I10)
11	=D11/M17	=Alpha*(E10-E9)+(1-Alpha)*(H10-H9)	=Beta*F11+(1-Beta)*G10	=Alpha*E10+(1-Alpha)*H10+G11	=M17*H11	=ABS(D11-I11)

Range Name	Cells
ActualForecast	I6:I30
Alpha	M5
Beta	M6
ForecastingError	J6:J30
InitialEstimateAverage	M9
InitialEstimateTrend	M10
MAD	M30
MSE	M33
SeasonalFactor	M16:M27
SeasonallyAdjustedForecast	H6:H30
SeasonallyAdjustedValue	E6:E30
TrueValue	D6:D30
TypeOfSeasonality	M13

	L	M
30	MAD =	=AVERAGE(ForecastingError)

	L	M
33	MSE =	=SUMSQ(ForecastingError)/COUNT(ForecastingError)

■ FIGURE 27.5

The Excel template in your OR Courseware for the exponential smoothing with trend method with seasonal adjustments is applied here to the CCW problem.

The Excel template in Fig. 27.5 shows the results from these calculations for all 12 quarters over the 3 years, as well as for the upcoming quarter. The middle of the figure shows the plots of all the seasonally adjusted call volumes and seasonally adjusted forecasts. Note how each trend up or down in the call volumes causes the forecasts to gradually trend in the same direction, but then the trend in the forecasts takes a couple of quarters to turn around when the trend in call volumes suddenly reverses direction. Each number in column I is calculated by multiplying the seasonally adjusted forecast in column H by the corresponding seasonal factor in column M to obtain the forecast of the actual value (not seasonally adjusted) for the average daily call volume. Column J then shows the resulting *forecasting errors* (the absolute value of the difference between columns D and I).

Forecasting More Than One Time Period Ahead

We have focused thus far on forecasting what will happen in the *next* time period (the next quarter in the case of CCW). However, decision makers sometimes need to forecast further into the future. How can the various forecasting methods be adapted to do this?

In the case of the methods for a constant-level model presented in Sec. 27.4, the forecast for the next period F_{t+1} also is the best available forecast for subsequent periods as well. However, when there is a *trend* in the data, as we are assuming in this section, it is important to take this trend into account for long-range forecasts. *Exponential smoothing with trend* provides a straightforward way of doing this. In particular, after determining the *estimated trend* T_{t+1} , this method's forecast for n time periods into the future is

$$F_{t+n} = \alpha x_t + (1 - \alpha)F_t + nT_{t+1}.$$

27.7 TIME SERIES FORECASTING WITH CB PREDICTOR

As described in Selected Reference 11, there now are literally dozens of commercial software packages available for performing the kinds of forecasting techniques presented in the preceding sections. We now will use one of these packages, an Excel add-in called CB Predictor, to illustrate what they can do. CB Predictor is one module of the popular Crystal Ball software suite featured in Sec. 20.6 and Chap. 28. (See Sec. 20.6 for information about how to access Crystal Ball.)

CB Predictor will perform most of the forecasting techniques discussed in this chapter, as well as a few others. This requires entering the historical data in a spreadsheet, such as shown in Fig. 27.6 for the CCW problem.

The next step is to select a cell within the body of the data (e.g., somewhere in column D of Fig. 27.6), and choose CB Predictor from the Crystal Ball tab (for Excel 2007) or the CB Tools menu (for earlier versions of Excel if CB Predictor is installed as part of the Crystal Ball package) or the Tools menu (for earlier versions of Excel if CB Predictor is installed as a stand-alone package). This brings up the CB Predictor dialogue box shown in Fig. 27.7. This dialogue box contains four panes (Input Data, Data Attributes, Method Gallery, and Results) used to perform the 10-step process of forecasting with CB Predictor.

Figure 27.7 shows the entries that are needed in this dialogue box for the CCW problem. In the first pane (Input Data), CB Predictor has made a (correct) guess as to the appropriate cell range for the data (C3:D15) in Step 1. If this guess had not been correct, the needed corrections should be made. Next, the arrangement of the data is specified in Step 2. For CCW, the entries of the dot and checkmarks shown in the figure indicate that the data are in columns, the first column has dates, and the first row has headers. If desired, you can view a graph of the data in Step 3.

	A	B	C	D
1	Crystal Ball Predictor for CCW			
2				
3		Year	Quarter	Call Volume
4		1	Q1	6,809
5		1	Q2	6,465
6		1	Q3	6,569
7		1	Q4	8,266
8		2	Q1	7,257
9		2	Q2	7,064
10		2	Q3	7,784
11		2	Q4	8,724
12		3	Q1	6,992
13		3	Q2	6,822
14		3	Q3	7,949
15		3	Q4	9,650

FIGURE 27.6

The first step in doing time-series forecasting with CB Predictor is to enter the historical data in a spreadsheet, as shown here for the CCW problem.

Click on Next or the Data Attributes tab to move to the second pane. Step 4 specifies the type of data and whether to apply seasonality. For CCW, the data are in quarters and there is seasonality over each cycle of four quarters, so these entries have been made in Fig. 27.7. Step 5 is used only if you wish to perform linear regression (the subject of Sec. 27.10).

In the third pane (Method Gallery), Step 6 allows you to choose the time-series forecasting methods that CB Predictor should try. The entries of the checkmarks for Step 6 in Fig. 27.7 indicate that the decision has been made for the CCW problem to try all the available methods. Most of the techniques displayed in the Methods Gallery are similar to techniques that were covered in the preceding sections, although are given somewhat different names. (Three of the techniques were not covered.) Table 27.3 identifies the relationships between the methods in CB Predictor and the methods discussed in the preceding sections. However, each of the CB Predictor techniques that involves exponential smoothing differs from its counterpart in the preceding sections by eliminating the need for the user to provide any inputs (seasonal factors, initial estimates, and smoothing constants) other than the raw data. Instead, CB Predictor uses the raw data to provide the best fit for all these inputs as well as the forecasts (so considerable raw data is needed to accomplish all this well). This is very helpful for an inexperienced user, but also deprives an experienced user of the opportunity to provide any desired inputs. Therefore, CB Predictor includes an option to override this feature and manually provide smoothing constants. (This is done by double clicking on any technique in the dialogue box, which brings up a Parameters dialogue box, and then entering user-defined parameters.) CB Predictor will try all the techniques that are checked in the Methods Gallery and use the one that fits the historical data the best.

In the final pane (Results), Steps 7, 8 and 9 are used to specify how many periods to forecast, the percentiles for a confidence interval, and whether to paste the forecast into the spreadsheet. Figure 27.7 shows the specification made for the CCW problem. Once these steps are completed, clicking on Preview can generate a preview graph of the forecast. The resulting forecast is shown in Fig. 27.8. The historical data for years 1 through 3 are shown as a light line, while the forecasts generated are shown with a dark line. In year 4, the forecasts are displayed along with a 5 percent and 95 percent confidence interval (light lines). These light lines specify the range of values that fall within the confidence interval. Given the percentiles of 5 percent and 95 percent for the confidence

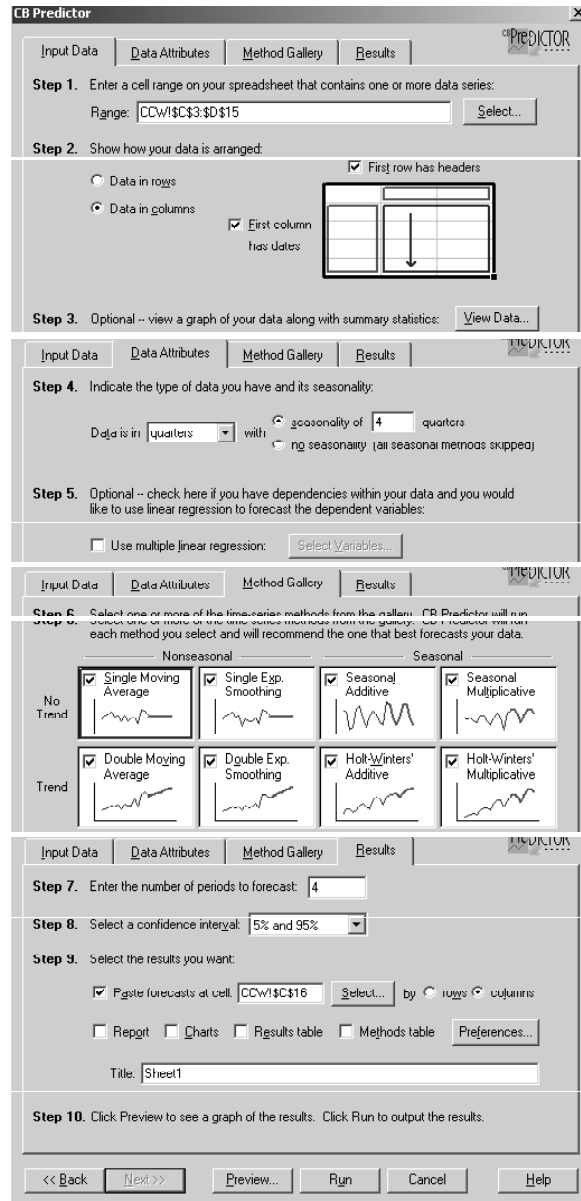


FIGURE 27.7
The CB Predictor dialogue box has the four panes (Input Data, Data Attributes, Method Gallery, and Results) displayed here, where the entries shown are for the CCW problem.

interval, each piece of actual year 4 data should have only a 5 percent chance of falling below the lower line and a 5 percent chance of rising above the upper line if the time series is not disrupted by some event.

Clicking on Run in the final pane generates the forecast for the number of periods specified in Step 7 and (if Step 9 is checked) pastes the forecast directly into the spreadsheet. The result is shown in Fig. 27.9, where cells D16:D19 give the forecasts for all four quarters of year 4. Note that the seasonality that has been well established in the first three years is clearly reflected in these quarterly forecasts for year 4. For each of the

FIGURE 27.8
 The CB Predictor preview graph for the CCW problem uses dark lines to show the forecasts generated retrospectively (Years 1 through 3) and for each quarter of the next year (Year 4). The light lines for Years 1 through 3 display the historical data, whereas the two light lines in Year 4 show a confidence interval (with percentiles of 5 percent and 95 percent) for next year's forecasts.

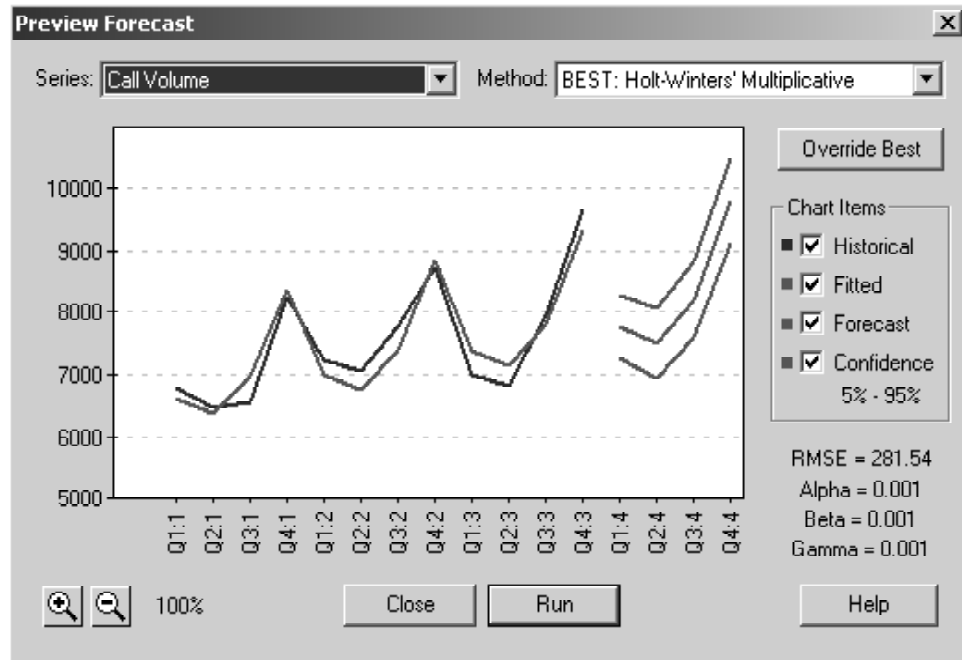


TABLE 27.3 The relationship between CB Predictor techniques and the forecasting techniques in this chapter

CB Predictor Technique	Related Technique in This Chapter
Single moving average	Moving average
Double moving average	Not covered
Single exponential smoothing	Exponential smoothing
Double exponential smoothing	Exponential smoothing with trend
Seasonal additive	Not covered
Holt-Winters additive	Not covered
Seasonal multiplicative	Exponential smoothing with seasonality
Holt-Winters multiplicative	Exponential smoothing with seasonality and trend

four quarters of a year, there is a generally upward trend in the call volume data over the first three years (especially in the third and fourth quarters). This too is clearly reflected in the forecasts for each quarter of year 4. The confidence intervals for these forecasts shown in Fig. 27.8 provide additional valuable information by indicating the degree of uncertainty in these forecasts.

The top right-hand side of the Preview Forecast in Fig. 27.8 indicates that the technique that fits the historical data best is the Holt-Winters multiplicative method (called exponential smoothing with seasonality and trend in this chapter), so this is the one that provides the forecasts shown in Fig. 27.9. The bottom right-hand side of Fig. 27.8 also reveals the curious fact that this method concluded that the best values for all the smoothing constants—alpha, beta, and gamma (the smoothing constant for seasonal factors)—are essentially 0. This means that the forecasting method essentially uses only an initial estimate, initial trend, and initial seasonality factors (all calculated based on the first three

	A	B	C	D
1	Crystal Ball Predictor for CCW			
2				
3		Year	Quarter	Call Volume
4		1	Q1	6,809
5		1	Q2	6,465
6		1	Q3	6,569
7		1	Q4	8,266
8		2	Q1	7,257
9		2	Q2	7,064
10		2	Q3	7,784
11		2	Q4	8,724
12		3	Q1	6,992
13		3	Q2	6,822
14		3	Q3	7,949
15		3	Q4	9,650
16		4	Q1	7,791
17		4	Q2	7,515
18		4	Q3	8,203
19		4	Q4	9,799

FIGURE 27.9
Using the data from Figure 27.6, cells D16:D19 give the forecast obtained by CB Predictor for each quarter of the coming year (Year 4) for the CCW problem. These are the forecasts plotted with the dark line in the Year 4 portion of the preview graph in Fig. 27.8.

years of data) to forecast the future, without updating these estimates using exponential smoothing. We suspect that with more historical data, the method would pick nonzero smoothing constants to take advantage of exponential smoothing.

In general, 12 data points are a rather small number for obtaining forecasts when using relatively sophisticated methods such as exponential smoothing with seasonality and trend. A lot more good data would be needed to take full advantage of the power of such methods for generating relatively precise forecasts. For the CCW problem, this suggests that using monthly instead of quarterly call volume data over the past three years (36 data points instead of only 12) probably would have provided a better basis for generating forecasts for year 4.

27.8 FORECASTING ERRORS

Several forecasting methods now have been presented. How does one choose the appropriate method for any particular application? Identifying the underlying model that best fits the time series (constant-level, linear trend, etc., perhaps in combination with seasonal effects) is an important first step. Assessing how *stable* the parameters of the model are, and so how much reliance can be placed on older data for forecasting, also helps to narrow down the selection of the method. However, the final choice between two or three methods may still not be clear. Some measure of performance is needed.

The goal is to generate forecasts that are as accurate as possible, so it is natural to base a measure of performance on the *forecasting errors*.

The **forecasting error** (also called the *residual*) for any period t is the absolute value of the deviation of the forecast for period t (F_t) from what then turns out to be the observed value of the time series for period t (x_t). Thus, letting E_t denote this error,

$$E_t = |x_t - F_t|.$$

For example, column J of the spreadsheet in Fig. 27.5 gives the forecasting errors when applying *exponential smoothing with trend* to the CCW example.

Given the forecasting errors for n time periods ($t = 1, 2, \dots, n$), two popular measures of performance are available. One, called the **mean absolute deviation (MAD)** is simply the average of the errors, so

$$\text{MAD} = \frac{\sum_{t=1}^n E_t}{n}.$$

This is the measure shown by MAD(M30) in Fig. 27.5. The other measure, called the **mean square error (MSE)**, instead averages the *square* of the forecasting errors, so

$$\text{MSE} = \frac{\sum_{t=1}^n E_t^2}{n}.$$

This measure is provided by MSE (M33) in Fig. 27.5.

The advantages of MAD are its ease of calculation and its straightforward interpretation. However, the advantage of MSE is that it imposes a relatively large penalty for a large forecasting error that can have serious consequences for the organization while almost ignoring inconsequentially small forecasting errors. In practice, managers often prefer to use MAD, whereas statisticians generally prefer MSE.

Either measure of performance might be used in two different ways. One is to compare alternative forecasting methods in order to choose one with which to begin forecasting. This is done by applying the methods *retrospectively* to the time series in the past (assuming such data exist). This is a very useful approach as long as the future behavior of the time series is expected to resemble its past behavior. Similarly, this retrospective testing can be used to help select the parameters for a particular forecasting method, e.g., the smoothing constant(s) for exponential smoothing. Second, after the real forecasting begins with some method, one of the measures of performance (or possibly both) normally would be calculated periodically to monitor how well the method is performing. If the performance is disappointing, the same measure of performance can be calculated for alternative forecasting methods to see if any of them would have performed better.

27.9 BOX-JENKINS METHOD

In practice, a forecasting method often is chosen without adequately checking whether the underlying model is an appropriate one for the application. The beauty of the **Box-Jenkins method** is that it carefully coordinates the model and the procedure. (Practitioners often use this name for the method because it was developed by G.E.P. Box and G.M. Jenkins. An alternative name is the *ARIMA method*, which is an acronym for autoregressive integrated moving average.) This method employs a systematic approach to identifying an appropriate model, chosen from a rich class of models. The historical data are used to test the validity of the model. The model also generates an appropriate forecasting procedure.

To accomplish all this, the Box-Jenkins method requires a great amount of past data (a minimum of 50 time periods), so it is used only for major applications. It also is a sophisticated and complex technique, so we will provide only a conceptual overview of the method. (See Selected References 2 and 6 for further details.)

The Box-Jenkins method is iterative in nature. First, a model is chosen. To choose this model, we must compute autocorrelations and partial autocorrelations and examine their patterns. An *autocorrelation* measures the correlation between time series values separated by

a fixed number of periods. This fixed number of periods is called the *lag*. Therefore, the autocorrelation for a lag of two periods measures the correlation between every other observation; i.e., it is the correlation between the original time series and the same series moved forward two periods. The *partial autocorrelation* is a conditional autocorrelation between the original time series and the same series moved forward a fixed number of periods, holding the effect of the other lagged times fixed. Good estimates of both the autocorrelations and the partial autocorrelations for all lags can be obtained by using a computer to calculate the *sample* autocorrelations and the *sample* partial autocorrelations. (These are “good” estimates because we are assuming large amounts of data.)

From the autocorrelations and the partial autocorrelations, we can identify the functional form of one or more possible models because a rich class of models is characterized by these quantities. Next we must estimate the parameters associated with the model by using the historical data. Then we can compute the residuals (the forecasting errors when the forecasting is done retrospectively with the historical data) and examine their behavior. Similarly, we can examine the behavior of the estimated parameters. If both the residuals and the estimated parameters behave as expected under the presumed model, the model appears to be validated. If they do not, then the model should be modified and the procedure repeated until a model is validated. At this point, we can obtain an actual forecast for the next period.

For example, suppose that the sample autocorrelations and the sample partial autocorrelations have the patterns shown in Fig. 27.10. The sample autocorrelations appear to decrease exponentially as a function of the time lags, while the same partial autocorrelations have spikes at the first and second time lags followed by values that seem to be of negligible magnitude. This behavior is characteristic of the functional form


$$X_t = B_0 + B_1X_{t-1} + B_2X_{t-2} + e_t.$$

Assuming this functional form, we use the time series data to estimate B_0 , B_1 , and B_2 . Denote these estimates by b_0 , b_1 , and b_2 , respectively. Together with the time series data, we then obtain the residuals

$$x_t - (b_0 + b_1x_{t-1} + b_2x_{t-2}).$$

If the assumed functional form is adequate, the residuals and the estimated parameters should behave in a predictable manner. In particular, the sample residuals should behave approximately as independent, normally distributed random variables, each having mean 0 and variance σ^2 (assuming that e_t , the random error at time period t , has mean 0 and variance σ^2). The estimated parameters should be uncorrelated and significantly different from zero. Statistical tests are available for this diagnostic checking.

■ **FIGURE 27.10**
Plot of sample autocorrelation and partial autocorrelation versus time lags.



The Box-Jenkins procedure appears to be a complex one, and it is. Fortunately, computer software is available. The programs calculate the sample autocorrelations and the sample partial autocorrelations necessary for identifying the form of the model. They also estimate the parameters of the model and do the diagnostic checking. These programs, however, cannot accurately identify one or more models that are compatible with the autocorrelations and the partial autocorrelations. Expert human judgment is required. This expertise can be acquired, but it is beyond the scope of this text. Although the Box-Jenkins method is complicated, the resulting forecasts are extremely accurate and, when the time horizon is short, better than most other forecasting methods. Furthermore, the procedure produces a measure of the forecasting error.

27.10 CAUSAL FORECASTING WITH LINEAR REGRESSION

In the preceding seven sections, we have focused on *time series forecasting methods*, i.e., methods that forecast the next value in a time series based on its previous values. We now turn to another type of approach to forecasting.

Causal Forecasting

In some cases, the variable to be forecasted has a rather direct relationship with one or more other variables whose values will be known at the time of the forecast. If so, it would make sense to base the forecast on this relationship. This kind of approach is called *causal forecasting*.

Causal forecasting obtains a forecast of the quantity of interest (the *dependent variable*) by relating it directly to one or more other quantities (the *independent variables*) that drive the quantity of interest.

Table 27.4 shows some examples of the kinds of situations where causal forecasting sometimes is used. In each of the first three cases, the indicated dependent variable can be expected to go up or down rather directly with the independent variable(s) listed in the rightmost column. The last case also applies when some quantity of interest (e.g., sales of a product) tends to follow a steady trend upward (or downward) with the passage of time (the independent variable that drives the quantity of interest).

As one specific example, Sec. 27.1 includes a description of American Airline's elaborate system for forecasting its need for expensive spare parts (its "rotatable" parts) to continue operating its fleet of several hundred airplanes. This system uses causal forecasting, where the demand for spare parts is the dependent variable and the number of flying hours is the independent variable. This makes sense because the demand for spare parts should be roughly proportional to the number of flying hours for the fleet.

■ **TABLE 27.4** Possible examples of causal forecasting

Type of Forecasting	Possible Dependent Variable	Possible Independent Variables
Sales	Sales of a product	Amount of advertising
Spare parts	Demand for spare parts	Usage of equipment
Economic trends	Gross domestic product	Various economic factors
Any quantity	This same quantity	Time

Linear Regression

We will focus on the type of causal forecasting where the mathematical relationship between the dependent variable and the independent variable(s) is assumed to be a linear one (plus some random fluctuations). The analysis in this case is referred to as *linear regression*.

To illustrate the linear regression approach, suppose that a publisher of textbooks is concerned about the initial press run for her books. She sells books both through bookstores and through mail orders. This latter method uses an extensive advertising campaign on line, as well as through publishing media and direct mail. The advertising campaign is conducted prior to the publication of the book. The sales manager has noted that there is a rather interesting linear relationship between the number of mail orders and the number sold through bookstores during the first year. He suggests that this relationship be exploited to determine the initial press run for subsequent books.

Thus, if the number of mail order sales for a book is denoted by X and the number of bookstore sales by Y , then the random variables X and Y exhibit a *degree of association*. However there is *no functional relationship* between these two random variables; i.e., given the number of mail order sales, one does not expect to determine *exactly* the number of bookstore sales. For any given number of mail order sales, there is a range of possible bookstore sales, and vice versa.

What, then, is meant by the statement, “The sales manager has noted that there is a rather interesting linear relationship between the number of mail orders and the number sold through bookstores during the first year”? Such a statement implies that the *expected value* of the number of bookstore sales is linear with respect to the number of mail order sales, i.e.,

$$E[Y|X = x] = A + Bx.$$

Thus, if the number of mail order sales is x for many different books, the average number of corresponding bookstore sales would tend to be approximately $A + Bx$. This relationship between X and Y is referred to as a **degree of association model**.

As already suggested in Table 27.4, other examples of this degree of association model can easily be found. A college admissions officer may be interested in the relationship between a student’s performance on the college entrance examination and subsequent performance in college. An engineer may be interested in the relationship between tensile strength and hardness of a material. An economist may wish to predict a measure of inflation as a function of the cost of living index, and so on.

The degree of association model is not the only model of interest. In some cases, there exists a **functional relationship** between two variables that may be linked linearly. In a forecasting context, one of the two variables is time, while the other is the variable of interest. In Sec. 27.6, such an example was mentioned in the context of the generating process of the time series being represented by a linear trend superimposed with random fluctuations, i.e.,

$$X_t = A + Bt + e_t,$$

where A is a constant, B is the slope, and e_t is the random error, assumed to have expected value equal to zero and constant variance. (The symbol X_t can also be read as X given t or as $X|t$.) It follows that

$$E(X_t) = A + Bt.$$

Note that both the degree of association model and the *exact functional relationship* model lead to the same linear relationship, and their subsequent treatment is almost identical. Hence, the publishing example will be explored further to illustrate how to treat both kinds of models, although the special structure of the model

$$E(X_t) = A + Bt,$$

with t taking on integer values starting with 1, leads to certain simplified expressions. In the standard notation of regression analysis, X represents the **independent variable** and Y represents the **dependent variable** of interest. Consequently, the notational expression for this special time series model now becomes

$$Y_t = A + Bt + e_t.$$

Method of Least Squares

Suppose that bookstore sales and mail order sales are given for 15 books. These data appear in Table 27.5, and the resulting plot is given in Fig. 27.11.

It is evident that the points in Fig. 27.11 do not lie on a straight line. Hence, it is not clear where the line should be drawn to show the linear relationship. Suppose that an arbitrary line, given by the expression $\tilde{y} = a + bx$, is drawn through the data. A measure of how well this line fits the data can be obtained by computing the *sum of squares* of the vertical deviations of the actual points from the fitted line. Thus, let y_i represent the bookstore sales of the i th book and x_i the corresponding mail order sales. Denote by \tilde{y}_i the point on the fitted line corresponding to the mail order sales of x_i . The proposed measure of fit is then given by


$$Q = (y_1 - \tilde{y}_1)^2 + (y_2 - \tilde{y}_2)^2 + \dots + (y_{15} - \tilde{y}_{15})^2 = \sum_{i=1}^{15} (y_i - \tilde{y}_i)^2.$$

The usual method for identifying the “best” fitted line is the **method of least squares**. This method chooses that line $a + bx$ that makes Q a minimum. Thus, a and b are obtained simply by setting the partial derivatives of Q with respect to a and b equal to zero and solving the resulting equations. This method yields the solution

$$b = \frac{\sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})}{\sum_{i=1}^n (x_i - \bar{x})^2} = \frac{\sum_{i=1}^n x_i y_i - \left(\sum_{i=1}^n x_i\right) \left(\sum_{i=1}^n y_i\right) / n}{\sum_{i=1}^n x_i^2 - \left(\sum_{i=1}^n x_i\right)^2 / n}$$

■ **TABLE 27.5** Data for the mail-order and bookstore sales example

Mail-Order Sales	Bookstore Sales
1,310	4,360
1,313	4,590
1,320	4,520
1,322	4,770
1,338	4,760
1,340	5,070
1,347	5,230
1,355	5,080
1,360	5,550
1,364	5,390
1,373	5,670
1,376	5,490
1,384	5,810
1,395	6,060
1,400	5,940



■ **FIGURE 27.11**
Plot of mail order sales
versus bookstore sales from
Table 27.4.

and

$$a = \bar{y} - b\bar{x},$$

where

$$\bar{x} = \sum_{i=1}^n \frac{x_i}{n}$$

and

$$\bar{y} = \sum_{i=1}^n \frac{y_i}{n}.$$

(Note that \bar{y} is not the same as $\tilde{y} = a + bx$ discussed in the preceding paragraph.)

For the publishing example, the data in Table 27.5 and Fig. 27.11 yield

$$\bar{x} = 1,353.1,$$

$$\bar{y} = 5,219.3,$$

$$\sum_{i=1}^{15} (x_i - \bar{x})(y_i - \bar{y}) = 214,543.9,$$

$$\sum_{i=1}^{15} (x_i - \bar{x})^2 = 11,966,$$

$$a = -19,041.9,$$

$$b = 17.930.$$

Hence, the least-squares estimate of bookstore sales \tilde{y} with mail order sales x is given by

$$\tilde{y} = -19,041.9 + 17.930x,$$

and this is the line drawn in Fig. 27.11. Such a line is referred to as a **regression line**.

An Excel template called Linear Regression is available in your OR Courseware for calculating a regression line in this way. A procedure in the forecasting area of your IOR Tutorial also will perform this calculation for you, as well as enable you to graphically investigate the effect of making changes in the data.

This regression line is useful for forecasting purposes. For a given value of x , the corresponding value of y represents the forecast.

The decision maker may be interested in some measure of uncertainty that is associated with this forecast. This measure is easily obtained provided that certain assumptions can be made. Therefore, for the remainder of this section, it is assumed that

1. A random sample of n pairs $(x_1, Y_1), (x_2, Y_2), \dots, (x_n, Y_n)$ is to be taken.
2. The Y_i are normally distributed with mean $A + Bx_i$ and variance σ^2 (independent of i).

The assumption that Y_i is normally distributed is not a critical assumption in determining the uncertainty in the forecast, but the assumption of constant variance is crucial. Furthermore, an estimate of this variance is required.

An unbiased estimate of σ^2 is given by $s_{y|x}^2$, where

$$s_{y|x}^2 = \sum_{i=1}^n \frac{(y_i - \tilde{y}_i)^2}{n - 2}.$$

Confidence Interval Estimation of $E(Y|x = x^*)$

A very important reason for obtaining the linear relationship between two variables is to use the line for future decision making. From the regression line, it is possible to estimate $E(Y|x)$ by a *point estimate* (the forecast) and a *confidence interval estimate* (a measure of forecast uncertainty).

For example, the publisher might want to use this approach to estimate the expected number of bookstore sales corresponding to mail order sales of, say, 1,400, by both a point estimate and a confidence interval estimate for forecasting purposes.

A point estimate of $E(Y|x = x^*)$ is given by

$$\tilde{y}^* = a + bx^*,$$

where x^* denotes the given value of the independent variable and \tilde{y}^* is the corresponding point estimate.

The endpoints of a $(100)(1 - \alpha)$ percent confidence interval for $E(Y|x = x^*)$ are given by

$$a + bx^* - t_{\alpha/2;n-2}s_{y|x} \sqrt{\frac{1}{n} + \frac{(x^* - \bar{x})^2}{\sum_{i=1}^n (x_i - \bar{x})^2}}$$

and

$$a + bx^* + t_{\alpha/2;n-2}s_{y|x} \sqrt{\frac{1}{n} + \frac{(x^* - \bar{x})^2}{\sum_{i=1}^n (x_i - \bar{x})^2}},$$

where $s_{y|x}^2$ is the estimate of σ^2 , and $t_{\alpha/2;n-2}$ is the $100\alpha/2$ percentage point of the t distribution with $n - 2$ degrees of freedom as given in Table 27.6. Note that the interval is narrowest where $x^* = \bar{x}$, and it becomes wider as x^* departs from the mean.

In the publishing example with $x^* = 1,400$, $s_{y|x}^2$ is computed from the data in Table 27.5 to be 17,030, so $s_{y|x} = 130.5$. If a 95 percent confidence interval is required, Table 27.6 gives $t_{0.025;13} = 2.160$. The earlier calculation of a and b yields

$$a + bx^* = -19,041.9 + 17.930(1,400) = 6,060$$

as the point estimate of $E(Y|1,400)$, that is, the forecast. Consequentially, the confidence limits corresponding to mail order sales of 1,400 are

$$\begin{aligned} \text{Lower confidence limit} &= 6,060 - 2.160(130.5) \sqrt{\frac{1}{15} + \frac{46.9^2}{11,966}} \\ &= 5,919, \end{aligned}$$

$$\begin{aligned} \text{Upper confidence limit} &= 6,060 + 2.160(130.5) \sqrt{\frac{1}{15} + \frac{46.9^2}{11,966}} \\ &= 6,201. \end{aligned}$$

The fact that the confidence interval was obtained at a data point ($x = 1,400$) is purely coincidental.

The Excel template for linear regression in your OR Courseware does most of the computational work involved in calculating these confidence limits. In addition to computing a and b (the regression line), it calculates $s_{y|x}^2$, \bar{x} , and $\sum_{i=1}^n (x_i - \bar{x})^2$.

Predictions

The confidence interval statement for the expected number of bookstore sales corresponding to mail order sales of 1,400 may be useful for budgeting purposes, but it is not too useful for making decisions about the *actual* press run. Instead of obtaining bounds on the *expected* number of bookstore sales, this kind of decision requires bounds on what the *actual* bookstore sales will be, i.e., a **prediction interval** on the value that the random variable (bookstore sales) takes on. This measure is a *different* measure of forecast uncertainty.

The two endpoints of a prediction interval are given by the expressions

$$a + bx_+ - t_{\alpha/2;n-2}s_{y|x} \sqrt{1 + \frac{1}{n} + \frac{(x_+ - \bar{x})^2}{\sum_{i=1}^n (x_i - \bar{x})^2}}$$

■ **TABLE 27.6** 100 α percentage points of Student's t distribution

P{Student's t with ν Degrees of Freedom \geq Tabled Value} = α										
$\nu \backslash \alpha$	0.40	0.25	0.10	0.05	0.025	0.01	0.005	0.0025	0.001	0.0005
1	0.325	1.000	3.078	6.314	12.706	31.821	63.657	127.32	318.31	636.62
2	0.289	0.816	1.886	2.920	4.303	6.965	9.925	14.089	22.327	31.598
3	0.277	0.765	1.638	2.353	3.182	4.541	5.841	7.453	10.214	12.924
4	0.271	0.741	1.533	2.132	2.776	3.747	4.604	5.598	7.173	8.610
5	0.267	0.727	1.476	2.015	2.571	3.365	4.032	4.773	5.893	6.869
6	0.265	0.718	1.440	1.943	2.447	3.143	3.707	4.317	5.208	5.959
7	0.263	0.711	1.415	1.895	2.365	2.998	3.499	4.029	4.785	5.408
8	0.262	0.706	1.397	1.860	2.306	2.896	3.355	3.833	4.501	5.041
9	0.261	0.703	1.383	1.833	2.262	2.821	3.250	3.690	4.297	4.781
10	0.260	0.700	1.372	1.812	2.228	2.764	3.169	3.581	4.144	4.587
11	0.260	0.697	1.363	1.796	2.201	2.718	3.106	3.497	4.025	4.437
12	0.259	0.695	1.356	1.782	2.179	2.681	3.055	3.428	3.930	4.318
13	0.259	0.694	1.350	1.771	2.160	2.650	3.012	3.372	3.852	4.221
14	0.258	0.692	1.345	1.761	2.145	2.624	2.977	3.326	3.787	4.140
15	0.258	0.691	1.341	1.753	2.131	2.602	2.947	3.286	3.733	4.073
16	0.258	0.690	1.337	1.746	2.120	2.583	2.921	3.252	3.686	4.015
17	0.257	0.689	1.333	1.740	2.110	2.567	2.898	3.222	3.646	3.965
18	0.257	0.688	1.330	1.734	2.101	2.552	2.878	3.197	3.610	3.922
19	0.257	0.688	1.328	1.729	2.093	2.539	2.861	3.174	3.579	3.883
20	0.257	0.687	1.325	1.725	2.086	2.528	2.845	3.153	3.552	3.850
21	0.257	0.686	1.323	1.721	2.080	2.518	2.831	3.135	3.527	3.819
22	0.256	0.686	1.321	1.717	2.074	2.508	2.819	3.119	3.505	3.792
23	0.256	0.685	1.319	1.714	2.069	2.500	2.807	3.104	3.485	3.767
24	0.256	0.685	1.318	1.711	2.064	2.492	2.797	3.091	3.467	3.745
25	0.256	0.684	1.316	1.708	2.060	2.485	2.787	3.078	3.450	3.725
26	0.256	0.684	1.315	1.706	2.056	2.479	2.779	3.067	3.435	3.707
27	0.256	0.684	1.314	1.703	2.052	2.473	2.771	3.057	3.421	3.690
28	0.256	0.683	1.313	1.701	2.048	2.467	2.763	3.047	3.408	3.674
29	0.256	0.683	1.311	1.699	2.045	2.462	2.756	3.038	3.396	3.659
30	0.256	0.683	1.310	1.697	2.042	2.457	2.750	3.030	3.385	3.646
40	0.255	0.681	1.303	1.684	2.021	2.423	2.704	2.971	3.307	3.551
60	0.254	0.679	1.296	1.671	2.000	2.390	2.660	2.915	3.232	3.460
120	0.254	0.677	1.289	1.658	1.980	2.358	2.617	2.860	3.160	3.373
∞	0.253	0.674	1.282	1.645	1.960	2.326	2.576	2.807	3.090	3.291

Source: Table 12 of *Biometrika Tables for Statisticians*, vol. I, 3d ed., 1966, by permission of the Biometrika Trustees.

and

$$a + bx_+ + t_{\alpha/2;n-2} s_{y|x} \sqrt{1 + \frac{1}{n} + \frac{(x_+ - \bar{x})^2}{\sum_{i=1}^n (x_i - \bar{x})^2}}$$

For any given value of x (denoted here by x_+), the probability is $1 - \alpha$ that the value of the future Y_+ associated with x_+ will fall in this interval.

Thus, in the publishing example, if x_+ is 1,400, then the corresponding 95 percent prediction interval for the number of bookstore sales is given by $6,060 \pm 315$, which is naturally wider than the confidence interval for the expected number of bookstore sales, $6,060 \pm 141$.

This method of finding a prediction interval works fine if it is only being done once. However, it is not feasible to use the same data to find multiple prediction intervals with various values of x_+ in this way and then specify a probability that *all* these predictions will be correct. For example, suppose that the publisher wants prediction intervals for several different books. For each individual book, she still is able to use these expressions to find the prediction interval and then make the prediction that the bookstore sales will be within this interval, where the probability is $1 - \alpha$ that the prediction will be correct. However, what she cannot do is specify a probability that *all* these predictions will be correct. The reason is that these predictions are all based upon the same statistical data, so the predictions are not statistically independent. If the predictions were independent and if k future bookstore sales were being predicted, with each prediction being made with probability $1 - \alpha$, then the probability would be $(1 - \alpha)^k$ that *all* k predictions of future bookstore sales will be correct. Unfortunately, the predictions are *not* independent, so the actual probability cannot be calculated, and $(1 - \alpha)^k$ does not even provide a reasonable approximation.

This difficulty can be overcome by using **simultaneous tolerance intervals**. Using this technique, the publisher can take the mail order sales of any book, find an interval (based on the previously determined linear regression line) that will contain the actual bookstore sales with probability at least $1 - \alpha$, and repeat this for any number of books having the same or different mail order sales. Furthermore, the probability is P that *all* these predictions will be correct. An alternative interpretation is as follows. If every publisher followed this procedure, each using his or her own linear regression line, then $100P$ percent of the publishers (on average) would find that at least $100(1 - \alpha)$ percent of their bookstore sales fell into the predicted intervals. The expression for the endpoints of each such tolerance interval is given by

$$a + bx_+ - c^{**} s_{y|x} \sqrt{\frac{1}{n} + \frac{(x_+ - \bar{x})^2}{\sum_{i=1}^n (x_i - \bar{x})^2}}$$

and

$$a + bx_+ + c^{**} s_{y|x} \sqrt{\frac{1}{n} + \frac{(x_+ - \bar{x})^2}{\sum_{i=1}^n (x_i - \bar{x})^2}},$$

where c^{**} is given in Table 27.7.

Thus, the publisher can predict that the bookstore sales corresponding to known mail order sales will fall in these tolerance intervals. Such statements can be made for as many books as the publisher desires. Furthermore, the probability is P that at least $100(1 - \alpha)$ percent of bookstore sales corresponding to mail order sales will fall in these intervals. If P is chosen as 0.90 and $\alpha = 0.05$, the appropriate value of c^{**} is 11.625. Hence, the number of bookstore sales corresponding to mail order sales of 1,400 books is predicted to fall in the interval $6,060 \pm 759$. If another book had mail order sales of 1,353, the bookstore sales are predicted to fall in the interval $5,258 \pm 390$, and so on. At least 95 percent of the bookstore sales will fall into their predicted intervals, and these statements are made with confidence 0.90.

To summarize, we now have described three *measures of forecast uncertainty*. The first (in the preceding subsection) is a *confidence interval* on the *expected value* of the random variable Y (for example, bookstore sales) given the observed value x of the independent variable X (for example, mail order sales). The second is a *prediction interval* on the *actual value* that Y will take on, given x . The third is *simultaneous tolerance intervals* on a *succession of actual values* that Y will take on given a succession of observed values of X .

■ **TABLE 27.7** Values of c^{**}

n	$\alpha = 0.50$	$\alpha = 0.25$	$\alpha = 0.10$	$\alpha = 0.05$	$\alpha = 0.01$	$\alpha = 0.001$
$P = 0.90$						
4	7.471	10.160	13.069	14.953	18.663	23.003
6	5.380	7.453	9.698	11.150	14.014	17.363
8	5.037	7.082	9.292	10.722	13.543	16.837
10	4.983	7.093	9.366	10.836	13.733	17.118
12	5.023	7.221	9.586	11.112	14.121	17.634
14	5.101	7.394	9.857	11.447	14.577	18.232
16	5.197	7.586	10.150	11.803	15.057	18.856
18	5.300	7.786	10.449	12.165	15.542	19.484
20	5.408	7.987	10.747	12.526	16.023	20.140
$P = 0.95$						
4	10.756	14.597	18.751	21.445	26.760	32.982
6	6.652	9.166	11.899	13.669	17.167	21.266
8	5.933	8.281	10.831	12.484	15.750	19.568
10	5.728	8.080	10.632	12.286	15.553	19.369
12	5.684	8.093	10.701	12.391	15.724	19.619
14	5.711	8.194	10.880	12.617	16.045	20.050
16	5.771	8.337	11.107	12.898	16.431	20.559
18	5.848	8.499	11.357	13.204	16.845	21.097
20	5.937	8.672	11.619	13.521	17.272	21.652
$P = 0.99$						
4	24.466	33.019	42.398	48.620	60.500	74.642
6	10.444	14.285	18.483	21.215	26.606	32.920
8	8.290	11.453	14.918	17.166	21.652	26.860
10	7.567	10.539	13.796	15.911	20.097	24.997
12	7.258	10.182	13.383	15.479	19.579	24.403
14	7.127	10.063	13.267	15.355	19.485	24.316
16	7.079	10.055	13.306	15.410	19.582	24.467
18	7.074	10.111	13.404	15.552	19.794	24.746
20	7.108	10.198	13.566	15.745	20.065	25.122

Source: Reprinted by permission from G. J. Lieberman and R. G. Miller, "Simultaneous Tolerance Intervals in Regression," *Biometrika*, **50**(1 and 2): 164, 1963.

■ **27.11 FORECASTING IN PRACTICE**

You now have seen the major forecasting methods used in practice. We conclude with a brief look at how widely the various methods are used.

To begin, consider the actual forecasting applications discussed in Sec. 27.1 and summarized in Table 27.1. Most of this table is repeated here in Table 27.8, but the rightmost column now identifies which forecasting method was used in each application. Not surprisingly, these major forecasting projects chose one of the more sophisticated statistical forecasting methods. (The most sophisticated is the ARIMA method, commonly called the Box-Jenkins method, that is described in Sec. 27.9.)

Every company needs to do at least some forecasting, but their methods often are not as sophisticated as with these major projects. Some insight into their general approach

■ **TABLE 27.8** The forecasting methods used in the actual applications presented in Section 27.1

Organization	Quantity Being Forecasted	Forecasting Model
Merit Brass Co.	Sales of finished goods	Exponential smoothing
Hidroeléctrica Español	Energy demand	ARIMA (Box-Jenkins), etc.
American Airlines	Demand for different fare classes	Exponential smoothing
American Airlines	Need for spare parts to repair airplanes	Linear regression
Albuquerque Microelectronics	Production yield in wafer fabrication	Exponential smoothing with a linear trend
U.S. Department of Labor	Unemployment insurance payments	Linear regression
United Airlines	Demand at reservation offices and airports	ARIMA (Box-Jenkins)
Taco Bell	Number of customer arrivals	Moving average
L.L. Bean	Staffing needs at call center	ARIMA (Box-Jenkins)

was provided by a survey conducted some years ago¹ of sales forecasting practices at 500 U.S. corporations. Although this survey published in 1994 now is somewhat out of date, we believe that its results are still somewhat reflective of current forecasting practices.

This survey indicates that at that time, judgmental forecasting methods were somewhat more widely used than statistical methods. The main reasons given for using judgmental methods were accuracy and difficulty in obtaining the data required for statistical methods. Comments also were made that upper management is not familiar with quantitative techniques, that judgmental methods create a sense of ownership, and that these methods add a commonsense element to the forecast.

Among the judgmental methods, the most popular was a jury of executive opinion. This was especially true for companywide or industry sales forecasts but also holds true by a small margin over manager's opinion when forecasting sales of individual products or families of products.

Statistical forecasting methods also are fairly widely used, especially in companies with high sales. Compared to earlier surveys, familiarity with such methods is increasing. (Given that statistical forecasting methods now have been regularly taught in business schools and management seminars for many years, we anticipate that this trend of increasing familiarity with such methods has continued since the time of the survey.) However, many survey respondents cited better data availability as the improvement they most wanted to see in their organizations. The availability of good data is crucial for the use of these methods. (Fortunately, the rapid advances in information technology since the survey was conducted has led to much better data availability in many companies.)

The survey indicates that the moving-average method and linear regression were the most widely used statistical forecasting methods. The moving-average method was more popular for short- and medium-range forecasts (less than a year), as well as for forecasting sales of individual products and families of products. Linear regression was more popular for longer-range forecasts and for forecasting either companywide or industry sales. Both exponential smoothing and the last-value method also received considerable use. However, the highest dissatisfaction is with the last-value method, and its popularity was decreasing compared to earlier surveys.

When statistical forecasting methods were used, it was fairly common to also use judgmental methods to adjust the forecasts. (This continues to be fairly common practice.)

As managers have continued to become more familiar with statistical methods, and more used to using the computer to compile data and implement OR techniques, we anticipate that the usage of statistical forecasting methods is continuing to grow. However,

¹N. R. Sanders and K. B. Manrodt, "Forecasting Practices in U.S. Corporations: Survey Results," *Interfaces*, **24**: 92–100, March–April 1994.

there always will be an important role for judgmental methods, both alone and in combination with statistical methods.

Another important trend in recent years has been an increasing availability and usage of sophisticated software packages for applying statistical forecasting methods. (See Selected Reference 11 for a survey of these packages.) Selected Reference 9 provides a recent survey of the use, satisfaction, and performance of forecasting software in practice. Most of the U.S. corporations responding to the survey reported using software for their forecasts, although this sometimes involved using only spreadsheets or internally developed forecasting software. Those using commercially available software packages reported both the best forecasting performance and the greatest satisfaction with the features of the software.

27.12 CONCLUSIONS

The future success of any business depends heavily on the ability of its management to forecast well. Judgmental forecasting methods often play an important role in this process. However, the ability to forecast well is greatly enhanced if historical data are available to help guide the development of a statistical forecasting method. By studying these data, an appropriate model can be structured. A forecasting method that behaves well under the model should be selected. This method may require choosing one or more parameters—e.g., the smoothing constant α in exponential smoothing—and the historical data may prove useful in making this choice. After forecasting begins, the performance should be monitored carefully to assess whether modifications should be made in the method.

SELECTED REFERENCES

1. Armstrong, J. E. (ed.): *Handbook of Forecasting Principles*, Kluwer Academic Publishers (now Springer), Boston, 2001.
2. Box, G. E. P., and G. M. Jenkins: *Time Series Analysis, Forecasting and Control*, Holden-Day, San Francisco, 1976.
3. Bunn, D., and G. Wright: "Interaction of Judgmental and Statistical Methods: Issues and Analysis," *Management Science*, **37**: 501–518, 1991.
4. Franses, P. H.: "Do We Think We Make Better Forecasts Than in the Past? A Survey of Academics," *Interfaces*, **34**(6): 466–468, Nov.–Dec. 2004.
5. Hillier, F. S., and M. S. Hillier: *Introduction to Management Science: A Modeling and Case Studies Approach with Spreadsheets*, McGraw-Hill/Irwin, Burr Ridge, IL, 3rd ed. 2008, chap. 10.
6. Hoff, J. C.: *A Practical Guide to Box-Jenkins Forecasting*, Lifetime Learning Publications, Belmont, CA, 1983.
7. Marshall, K. T., and R. M. Oliver: *Decision Making and Forecasting*, McGraw-Hill, New York, 1995.
8. Sanders, N. R., and K. B. Manrodt: "Forecasting Practices in U.S. Corporations: Survey Results," *Interfaces*, **24**(2): 92–100, March–April 1994.
9. Sanders, N. R., and K. B. Manrodt: "Forecasting Software in Practice: Use, Satisfaction, and Performance," *Interfaces*, **33**(5): 90–93, Sept.–Oct. 2003.
10. Rust, R. T., D. Simester, R. J. Brodie, and V. Nilikant: "Model Selection Criteria: An Investigation of Relative Accuracy, Posterior Probabilities, and Combinations of Criteria," *Management Science*, **41**: 322–333, 1995.
11. Yurkiewicz, J.: "Forecasting Software Survey," *OR/MS Today*, **35**(3): 54–63, June 2008.

■ LEARNING AIDS FOR THIS CHAPTER ON THIS WEBSITE

“Ch. 27—Forecasting” Excel Files:

Template for *Seasonal Factors*
 Templates for *Last-Value Method* (with and without Seasonality)
 Templates for *Averaging Method* (with and without Seasonality)
 Templates for *Moving-Average Method* (with and without Seasonality)
 Templates for *Exponential Smoothing Method* (with and without Seasonality)
 Templates for *Exponential Smoothing with Trend* (with and without Seasonality)
 Template for *Linear Regression*

Procedures in IOR Tutorial:

Last Value Method
 Averaging Method
 Moving Average Method
 Exponential Smoothing
 Exponential Smoothing with Trend
 Linear Regression

An Excel Add-in:

CB Predictor (part of Crystal Ball 2000, Professional Edition)

“Ch. 27—Forecasting” LINGO File for Selected Examples

See Appendix 1 for documentation of the software.

■ PROBLEMS

To the left of each of the following problems (or their parts), we have inserted a T whenever the corresponding template listed above can be helpful. (Either CB Predictor or some of the above procedures in your IOR Tutorial should be used for certain problems, but this will be specified in the statement of the problem whenever needed.)

27.1-1. Select one of the applications of statistical forecasting methods listed in Table 27.1. Read the article describing the application in the indicated issue of *Interfaces*. Write a two-page summary of the application and the benefits it provided.

27.1-2. Select three of the applications of statistical forecasting methods listed in Table 27.1. Read the articles describing the applications in the indicated issues of *Interfaces*. For each one, write a one-page summary of the application and the benefits it provided.

27.4-1. The Hammaker Company’s newest product has had the following sales during its first five months: 5 17 29 41 39. The sales manager now wants a forecast of sales in the next month. (Use hand calculations rather than an Excel template.)

- (a) Use the last-value method.
- (b) Use the averaging method.
- (c) Use the moving-average method with the 3 most recent months.
- (d) Given the sales pattern so far, do any of these methods seem inappropriate for obtaining the forecast? Why?

27.4-2. Sales of stoves have been going well for the Good-Value Department Store. These sales for the past five months have been 15 18 12 17 13. Use the following methods to obtain a forecast of sales for the next month. (Use hand calculations rather than an Excel template.)

- (a) The last-value method.
- (b) The averaging method.
- (c) The moving-average method with 3 months.
- (d) If you feel that the conditions affecting sales next month will be the same as in the last five months, which of these methods do you prefer for obtaining the forecast? Why?

27.4-3. You are using the moving-average forecasting method based upon the last four observations. When making the forecast for the last period, the oldest of the four observations was 1,945 and the forecast was 2,083. The true value for the last period then turned out to be 1,977. What is your new forecast for the next period?

27.4-4. You are using the moving-average forecasting method based upon sales in the last three months to forecast sales for the next month. When making the forecast for last month, sales for the third month before were 805. The forecast for last month was 782 and then the actual sales turned out to be 793. What is your new forecast for next month?

27.4-5. After graduating from college with a degree in mathematical statistics, Ann Preston has been hired by the Monty Ward

Company to apply statistical methods for forecasting the company's sales. For one of the company's products, the moving-average method based upon sales in the 10 most recent months already is being used. Ann's first task is to update last month's forecast to obtain the forecast for next month. She learns that the forecast for last month was 1,551 and that the actual sales then turned out to be 1,532. She also learns that the sales for the tenth month before last month was 1,632. What is Ann's forecast for next month?

27.4-6. The J.J. Bone Company uses exponential smoothing to forecast the average daily call volume at its call center. The forecast for last month was 782, and then the actual value turned out to be 792. Obtain the forecast for next month for each of the following values of the smoothing constant: $\alpha = 0.1, 0.3, 0.5$.

27.4-7. You are using exponential smoothing to obtain monthly forecasts of the sales of a certain product. The forecast for last month was 2,083, and then the actual sales turned out to be 1,973. Obtain the forecast for next month for each of the following values of the smoothing constant: $\alpha = 0.1, 0.3, 0.5$.

27.4-8. If α is set equal to 0 or 1 in the exponential smoothing expression, what happens to the forecast?

27.4-9. A company uses exponential smoothing with $\alpha = \frac{1}{2}$ to forecast demand for a product. For each month, the company keeps a record of the forecast demand (made at the end of the preceding month) and the actual demand. Some of the records have been lost; the remaining data appear in the table below.

	January	February	March	April	May	June
Forecast			400	380	390	380
Actual	400		360	—	—	

- (a) Using only data in the table for March, April, May, and June, determine the actual demands in April and May.
- (b) Suppose now that a clerical error is discovered; the actual demand in January was 432, not 400, as shown in the table. Using only the actual demands going back to January (even though the February actual demand is unknown), give the corrected forecast for June.

27.5-1. Figure 27.3 shows CCW's average daily call volume for each quarter of the past three years, and column F of Fig. 27.4 gives the seasonally adjusted call volumes. Management now wonders what these seasonally adjusted call volumes would have been if the company had started using seasonal factors two years ago rather than applying them retrospectively now. (Use hand calculations rather than an Excel template.)

- (a) Use only the call volumes in Year 1 to determine the seasonal factors for Year 2 (so that the "average" call volume for each quarter is just the actual call volume for that quarter in Year 1).
- (b) Use these seasonal factors to determine the seasonally adjusted call volumes for Year 2.

- (c) Use the call volumes in Year 1 and 2 to determine the seasonal factors for Year 3.
- (d) Use the seasonal factors obtained in part (c) to determine the seasonally adjusted call volumes for Year 3.

27.5-2. Even when the economy is holding steady, the unemployment rate tends to fluctuate because of seasonal effects. For example, unemployment generally goes up in Quarter 3 (summer) as students (including new graduates) enter the labor market. The unemployment rate then tends to go down in Quarter 4 (fall) as students return to school and temporary help is hired for the Christmas season. Therefore, using seasonal factors to obtain a seasonally adjusted unemployment rate is helpful for painting a truer picture of economic trends.

Over the past 10 years, one state's average unemployment rates (not seasonally adjusted) in Quarters 1, 2, 3, and 4 have been 6.2 percent, 6.0 percent, 7.5 percent, and 5.5 percent, respectively. The overall average has been 6.3 percent. (Use hand calculations below rather than an Excel template.)

- (a) Determine the seasonal factors for the four quarters.
- (b) Over the next year, the unemployment rates (not seasonally adjusted) for the four quarters turn out to be 7.8 percent, 7.4 percent, 8.7 percent, and 6.1 percent. Determine the seasonally adjusted unemployment rates for the four quarters. What does this progression of rates suggest about whether the state's economy is improving?

27.5-3. Ralph Billett is the manager of a real estate agency. He now wishes to develop a forecast of the number of houses that will be sold by the agency over the next year.

The agency's quarter-by-quarter sales figures over the last three years are shown below.

Quarter	Year 1	Year 2	Year 3
1	23	19	21
2	22	21	26
3	31	27	32
4	26	24	28

- (Use hand calculations below rather than an Excel template.)
- (a) Determine the seasonal factors for the four quarters.
 - (b) After considering seasonal effects, use the last-value method to forecast sales in Quarter 1 of next year.
 - (c) Assuming that each of the quarterly forecasts is correct, what would the last-value method forecast as the sales in each of the four quarters next year?
 - (d) Based on his assessment of the current state of the housing market, Ralph's best judgment is that the agency will sell 100 houses next year. Given this forecast for the year, what is the quarter-by-quarter forecast according to the seasonal factors?

27.5-4. A manufacturer sells a certain product in batches of 100 to wholesalers. The following table shows the quarterly sales figure for this product over the last several years.

Quarter of 2000	Sales	Quarter of 2001	Sales	Quarter of 2002	Sales	Quarter of 2003	Sales	Quarter of 2004	Sales
1	6,900	1	8,200	1	9,400	1	11,400	1	8,800
2	6,700	2	7,000	2	9,200	2	10,000	2	7,600
3	7,900	3	7,300	3	9,800	3	9,400	3	7,500
4	7,100	4	7,500	4	9,900	4	8,400	4	—

The company incorporates seasonal effects into its forecasting of future sales. It then uses exponential smoothing (with seasonality) with a smoothing constant of $\alpha = 0.1$ to make these forecasts. When starting the forecasting, it uses the average sales over the past four quarters to make the initial estimate of the seasonally adjusted constant level A for the underlying constant-level model.

- T (a) Suppose that the forecasting started at the beginning of 2001. Use the data for 2000 to determine the seasonal factors and then determine the forecast of sales for each quarter of 2001.
- T (b) Suppose that the forecasting started at the beginning of 2002. Use the data for both 2000 and 2001 to determine the seasonal factors and then determine the forecast of sales for each quarter of 2002.
- T (c) Suppose that the forecasting started at the beginning of 2004. Use the data for 2000 through 2003 to determine the seasonal factors and then determine the forecast of sales for each quarter of 2004.
- (d) Under the assumptions of the constant-level model, the forecast obtained for any period of one year also provides the best available forecast at that time for the same period in any subsequent year. Use the results from parts (a), (b), and (c) to record the forecast of sales for Quarter 4 of 2004 when entering Quarter 4 of 2001, 2002, and 2004, respectively.
- (e) Evaluate whether it is important to incorporate seasonal effects into the forecasting procedure for this particular product.
- (f) Evaluate how well the constant-level assumption of the constant-level model (after incorporating seasonal effects) appears to hold for this particular product.

27.6-1. Look ahead at the scenario described in Prob. 27.8-3. Notice the steady trend upward in the number of applications over the past three years—from 4,600 to 5,300 to 6,000. Suppose now that the admissions office of Ivy College had been able to foresee this kind of trend and so had decided to use exponential smoothing with trend to do the forecasting. Suppose also that the initial estimates just over three years ago had been *expected value* = 3,900 and *trend* = 700. Then, with any values of the smoothing constants, the forecasts obtained by this forecasting method would have been exactly correct for all three years.

Illustrate this fact by doing the calculations to obtain these forecasts when the smoothing constant is $\alpha = 0.25$ and the trend smoothing constant is $\beta = 0.25$. (Use hand calculations rather than an Excel template.)

27.6-2. Exponential smoothing with trend, with a smoothing constant of $\alpha = 0.2$ and a trend smoothing constant of $\beta = 0.3$, is being used to forecast values in a time series. At this point,

the last two values have been 535 and then 550. The last two forecasts have been 530 and then 540. The last estimate of the trend factor has been 10. Use this information to forecast the next value in the time series. (Use hand calculations rather than an Excel template.)

27.6-3. The Healthwise Company produces a variety of exercise equipment. Healthwise management is very pleased with the increasing sales of its newest model of exercise bicycle. The sales during the last two months have been 4,655 and then 4,935.

Management has been using exponential smoothing with trend, with a smoothing constant of $\alpha = 0.1$ and a trend smoothing constant of $\beta = 0.2$, to forecast sales for the next month each time. The forecasts for the last two months were 4,720 and then 4,975. The last estimate of the trend factor was 240.

Calculate the forecast of sales for next month. (Use hand calculations rather than an Excel template.)

T **27.6-4.** The Pentel Microchip Company has started production of its new microchip. The first phase in this production is the wafer fabrication process. Because of the great difficulty in fabricating acceptable wafers, many of these tiny wafers must be rejected because they are defective. Therefore, management places great emphasis on continually improving the wafer fabrication process to increase its *production yield* (the percentage of wafers fabricated in the current lot that are of acceptable quality for producing microchips).

So far, the production yields of the respective lots have been 15, 21, 24, 32, 37, 41, 40, 47, 51, 53 percent. Use exponential smoothing with trend to forecast the production yield of the next lot. Begin with initial estimates of 10 percent for the expected value and 5 percent for the trend. Use smoothing constants of $\alpha = 0.2$ and $\beta = 0.2$.

27.7-1. Look ahead at Prob. 27.8-9. For each of the three cases, use CB Predictor to determine the best value of α by selecting only *single exponential smoothing* in step 6 and clicking Preview in step 10. Also obtain a forecast of the next observation in each of these time series.

27.7-2. Look ahead at Prob. 27.8-10. For each of the three cases, use CB Predictor to determine the best values of α and β by selecting only *double exponential smoothing* in step 6 and clicking Preview in step 10. Also obtain a forecast of the next observation in each of these time series.

27.7-3. Look ahead at the scenario described in Prob. 27.8-11. Use CB Predictor to compare the moving-average method, exponential

smoothing method, and exponential smoothing with trend (with the best choice of smoothing constants), so that CB Predictor then can select the best of these methods for this particular series of demands. Also obtain the forecast for next year with the selected method.

27.8-1. You have been forecasting sales the last four quarters. These forecasts and the true values that subsequently were obtained are shown below.

Quarter	Forecast	True Value
1	327	345
2	332	317
3	328	336
4	330	311

- (a) Calculate MAD.
- (b) Calculate MSE.

27.8-2. Sharon Johnson, sales manager for the Alvarez-Baines Company, is trying to choose between two methods for forecasting sales that she has been using during the past five months. During these months, the two methods obtained the forecasts shown below for the company's most important product, where the subsequent actual sales are shown on the right.

Month	Forecast		Actual Sales
	Method 1	Method 2	
1	5,324	5,208	5,582
2	5,405	5,377	4,906
3	5,195	5,462	5,755
4	5,511	5,414	6,320
5	5,762	5,549	5,153

- (a) Calculate and compare MAD for these two forecasting methods.
- (b) Calculate and compare MSE for these two forecasting methods.
- (c) Sharon is uncomfortable with choosing between these two methods based on such limited data, but she also does not want to delay further before making her choice. She does have similar sales data for the three years prior to using these forecasting methods the past five months. How can these older data be used to further help her evaluate the two methods and choose one?

27.8-3. Three years ago, the admissions office for Ivy College began using exponential smoothing with a smoothing constant of 0.25 to forecast the number of applications for admission each year. Based on previous experience, this process was begun with an initial estimate of 5,000 applications. The actual number of applications then turned out to be 4,600 in the first year. Thanks to new favorable ratings in national surveys, this number grew to 5,300 in the second year and 6,000 last year. (Use hand calculations below rather than an Excel template.)

- (a) Determine the forecasts that were made for each of the past three years.
- (b) Calculate MAD for these three years.

- (c) Calculate MSE for these three years.
- (d) Determine the forecast for next year.

27.8-4. Ben Swanson, owner and manager of Swanson's Department Store, has decided to use statistical forecasting to get a better handle on the demand for his major products. However, Ben now needs to decide which forecasting method is most appropriate for each category of product. One category is major household appliances, such as washing machines, which have a relatively stable sales level. Monthly sales of washing machines last year are shown below.

Month	Sales	Month	Sales	Month	Sales
January	23	May	22	September	21
February	24	June	27	October	29
March	22	July	20	November	23
April	28	August	26	December	28

- (a) Considering that the sales level is relatively stable, which of the most basic forecasting methods—the last-value method or the averaging method or the moving-average method—do you feel would be most appropriate for forecasting future sales? Why?
- T (b) Use the last-value method retrospectively to determine what the forecasts would have been for the last 11 months of last year. What is MAD?
- T (c) Use the averaging method retrospectively to determine what the forecasts would have been for the last 11 months of last year. What is MAD?
- T (d) Use the moving-average method with $n = 3$ retrospectively to determine what the forecasts would have been for the last 9 months of last year. What is MAD?
- (e) Use their MAD values to compare the three methods.
- (f) Use their MSE values to compare the three methods.
- (g) Do you feel comfortable in drawing a definitive conclusion about which of the three forecasting methods should be the most accurate in the future based on these 12 months of data?

T **27.8-5.** Reconsider Prob. 27.8-4. Ben Swanson now has decided to use the exponential smoothing method to forecast future sales of washing machines, but he needs to decide on which smoothing constant to use. Using an initial estimate of 24, apply this method retrospectively to the 12 months of last year with $\alpha = 0.1, 0.2, 0.3, 0.4,$ and 0.5 .

- (a) Compare MAD for these five values of the smoothing constant α .
- (b) Calculate and compare MSE for these five values of α .

27.8-6. Reconsider Prob. 27.8-4. For each of the forecasting methods specified in parts (b), (c), and (d), use the corresponding procedure in the forecasting area of your IOR Tutorial to obtain the requested forecasts. Then use the accompanying graph that plots both the sales data and forecasts to answer the following questions for these forecasting methods.

- (a) Based on your examination of the graphs for the three forecasting methods, which method do you feel is doing the best job of forecasting with the given data? Why?

- (b) Ben Swanson now has found that an error was made in determining the sales for April, but he has not yet obtained the corrected sales figure. For each of the three forecasting methods, Ben wants to know which of the original monthly forecasts would change now because of changing the sales figure for April. Answer this question by dragging vertically the blue dot that corresponds to April sales and observing which of the red dots (corresponding to monthly forecasts) move.
- (c) Repeat part (b) if the sales for April change from 28 to 16.
- (d) Repeat part (b) if the sales for April change from 28 to 40.

27.8-7. Management of the Jackson Manufacturing Corporation wishes to choose a statistical forecasting method for forecasting total sales for the corporation. Total sales (in millions of dollars) for each month of last year are shown below.

Month	Sales	Month	Sales	Month	Sales
January	126	May	153	September	147
February	137	June	154	October	151
March	142	July	148	November	159
April	150	August	145	December	166

- (a) Note how the sales level is shifting significantly from month to month—first trending upward and then dipping down before resuming an upward trend. Assuming that similar patterns would continue in the future, evaluate how well you feel each of the five forecasting methods introduced in Secs. 27.4 and 27.6 would perform in forecasting future sales.
- T (b) Apply the last-value method, the averaging method, and the moving-average method (with $n = 3$) retrospectively to last year's sales and compare their MAD values. Then compare their MSE values.
- T (c) Using an initial estimate of 120, apply the exponential smoothing method retrospectively to last year's sales with $\alpha = 0.1, 0.2, 0.3, 0.4,$ and 0.5 . Compare both MAD and MSE for these five values of the smoothing constant α .
- T (d) Using initial estimates of 120 for the expected value and 10 for the trend, apply exponential smoothing with trend retrospectively to last year's sales. Use all combinations of the smoothing constants where $\alpha = 0.1$ or 0.3 or 0.5 and $\beta = 0.1$ or 0.3 or 0.5 . Compare both MAD and MSE for these nine combinations.
- (e) Which one of the above forecasting methods would you recommend that management use? Using this method, what is the forecast of total sales for January of the new year?

27.8-8. Reconsider Prob. 27.8-7. For each of the forecasting method specified in parts (b), (c), and (d) (with smoothing constants $\alpha = 0.5$ and $\beta = 0.5$ as needed), use the corresponding procedure in the forecasting area of your IOR Tutorial to obtain the requested forecasts. Then use the accompanying graph that plots both the sales data and forecasts to answer the following questions for these forecasting methods.

- (a) Based on your examination of the graphs for the five forecasting methods, which method do you feel is doing the best job of forecasting with the given data? Why?

- (b) Management now has been informed that an error was made in calculating the sales for April, but a corrected sales figure has not yet been obtained. Therefore, for each of the five forecasting methods, management wants to know which of the original monthly forecasts would change now because of changing the sales figure for April. Answer this question by dragging vertically the blue dot that corresponds to April sales and observing which of the red dots (corresponding to monthly forecasts) move.
- (c) Repeat part (b) if the sales for April change from 150 to 125.
- (d) Repeat part (b) if the sales for April change from 150 to 175.

T **27.8-9.** Choosing an appropriate value of the smoothing constant α is a key decision when applying the exponential smoothing method. When relevant historical data exist, one approach to making this decision is to apply the method retrospectively to these data with different values of α and then choose the value of α that gives the smallest MAD. Use this approach for choosing α with each of the following time series representing monthly sales. In each case, use an initial estimate of 50 and compare $\alpha = 0.1, 0.2, 0.3, 0.4,$ and 0.5 .

- (a) 51 48 52 49 53 49 48 51 50 49
- (b) 52 50 53 51 52 48 52 53 49 52
- (c) 50 52 51 55 53 56 52 55 54 53

T **27.8-10.** The choice of the smoothing constants α and β has a considerable effect on the accuracy of the forecasts obtained by using exponential smoothing with trend. For each of the following time series, set $\alpha = 0.2$ and then compare MAD obtained with $\beta = 0.1, 0.2, 0.3, 0.4,$ and 0.5 . Begin with initial estimates of 50 for the expected value and 2 for the trend.

- (a) 52 55 55 58 59 63 64 66 67 72 73 74
- (b) 52 55 59 61 66 69 71 72 73 74 73 74
- (c) 52 53 51 50 48 47 49 52 57 62 69 74

27.8-11. The Andes Mining Company mines and ships copper ore. The company's sales manager, Juanita Valdes, has been using the moving-average method based on the last three years of sales to forecast the demand for the next year. However, she has become dissatisfied with the inaccurate forecasts being provided by this method.

Here are the annual demands (in tons of copper ore) over the past 10 years: 382 405 398 421 426 415 443 451 446 464

- (a) Explain why this pattern of demands inevitably led to significant inaccuracies in the moving-average forecasts.
- T (b) Determine the moving-average forecasts for the past 7 years. What is MAD? What is the forecast for next year?
- T (c) Determine what the forecasts would have been for the past 10 years if the exponential smoothing method had been used instead with an initial estimate of 380 and a smoothing constant of $\alpha = 0.5$. What is MAD? What is the forecast for next year?
- T (d) Determine what the forecasts would have been for the past 10 years if exponential smoothing with trend had been used instead. Use initial estimates of 370 for the expected value and 10 for the trend, with smoothing constants $\alpha = 0.25$ and $\beta = 0.25$.
- (e) Based on the MAD values, which of these three methods do you recommend using hereafter?

PROBLEMS

27.8-12. Reconsider Prob. 27.8-11. For each of the forecasting methods specified in parts (b), (c), and (d), use the corresponding procedure in the forecasting area of your IOR Tutorial to obtain the requested forecasts. After examining the accompanying graph that plots both the demand data and forecasts, write a one-sentence description for each method of whether its plot of forecasts tends to lie below or above or at about the same level as the demands being forecasted. Then use these conclusions to select one of the methods to recommend using hereafter.

27.8-13. The Centerville Water Department provides water for the entire town and outlying areas. The number of acre-feet of water consumed in each of the four seasons of the three preceding years is shown below.

Season	Year 1	Year 2	Year 3
Winter	25	27	24
Spring	47	46	49
Summer	68	72	70
Fall	42	39	44

- T (a) Determine the seasonal factors for the four seasons.
- T (b) After considering seasonal effects, use the last-value method to forecast water consumption next winter.
- (c) Assuming that each of the forecasts for the next three seasons is correct, what would the last-value method forecast as the water consumption in each of the four seasons next year?
- T (d) After considering seasonal effects, use the averaging method to forecast water consumption next winter.
- T (e) After considering seasonal effects, use the moving-average method based on four seasons to forecast water consumption next winter.
- T (f) After considering seasonal effects, use the exponential smoothing method with an initial estimate of 46 and a smoothing constant of $\alpha = 0.1$ to forecast water consumption next winter.
- T (g) Compare the MAD values of these four forecasting methods when they are applied retrospectively to the last three years.
- T (h) Compare the MSE values of these four forecasting methods when they are applied retrospectively to the last three years.

27.8-14. Reconsider Prob. 27.5-3. Ralph Billett realizes that the last-value method is considered to be the naive forecasting method, so he wonders whether he should be using another method. Therefore, he has decided to use the available Excel templates that consider seasonal effects to apply various statistical forecasting methods retrospectively to the past three years of data and compare their MAD values.

- T (a) Determine the seasonal factors for the four quarters.
- T (b) Apply the last-value method.
- T (c) Apply the averaging method.
- T (d) Apply the moving-average method based on the four most recent quarters of data.
- T (e) Apply the exponential smoothing method with an initial estimate of 25 and a smoothing constant of $\alpha = 0.25$.
- T (f) Apply exponential smoothing with trend with smoothing constants of $\alpha = 0.25$ and $\beta = 0.25$. Use initial estimates of 25 for the expected value and 0 for the trend.

- T (g) Compare the MAD values for these methods. Use the one with the smallest MAD to forecast sales in Quarter 1 of next year.
- (h) Use the forecast in part (g) and the seasonal factors to make long-range forecasts now of the sales in the remaining quarters of next year.

T **27.8-15.** Transcontinental Airlines maintains a computerized forecasting system to forecast the number of customers in each fare class who will fly on each flight in order to allocate the available reservations to fare classes properly. For example, consider *economy-class customers* flying in midweek on the noon flight from New York to Los Angeles. The following table shows the average number of such passengers during each month of the year just completed. The table also shows the seasonal factor that has been assigned to each month based on historical data.

Month	Average Number	Seasonal Factor	Month	Average Number	Seasonal Factor
January	68	0.90	July	94	1.17
February	71	0.88	August	96	1.15
March	66	0.91	September	80	0.97
April	72	0.93	October	73	0.91
May	77	0.96	November	84	1.05
June	85	1.09	December	89	1.08

- (a) After considering seasonal effects, compare both the MAD and MSE values for the last-value method, the averaging method, the moving-average method (based on the most recent three months), and the exponential smoothing method (with an initial estimate of 80 and a smoothing constant of $\alpha = 0.2$) when they are applied retrospectively to the past year.
- (b) Use the forecasting method with the smallest MAD value to forecast the average number of these passengers flying in January of the new year.

27.8-16. Reconsider Prob. 27.8-15. The economy is beginning to boom so the management of Transcontinental Airlines is predicting that the number of people flying will steadily increase this year over the relatively flat (seasonally adjusted) level of last year. Since the forecasting methods considered in Prob. 27.8-15 are relatively slow in adjusting to such a trend, consideration is being given to switching to exponential smoothing with trend.

Subsequently, as the year goes on, management's prediction proves to be true. The following table shows the average number of the passengers under consideration in each month of the new year.

Month	Average Number	Month	Average Number	Month	Average Number
January	75	May	85	September	94
February	76	June	99	October	90
March	81	July	107	November	106
April	84	August	108	December	110

- T (a) Repeat part (a) of Prob. 27.8-15 for the two years of data.

- T (b) After considering seasonal effects, apply exponential smoothing with trend to just the new year. Use initial estimates of 80 for the expected value and 2 for the trend, along with smoothing constants of $\alpha = 0.2$ and $\beta = 0.2$. Compare MAD for this method to the MAD values obtained in part (a). Then do the same with MSE.
- T (c) Repeat part (b) when exponential smoothing with trend is begun at the beginning of the first year and then applied to both years, just like the other forecasting methods in part (a). Use the same initial estimates and smoothing constants except change the initial estimate of trend to 0.
- (d) Based on these results, which forecasting method would you recommend that Transcontinental Airlines use hereafter?
- (e) One month later (January), the average number of passengers is 97. Use CB Predictor to apply exponential smoothing with seasonality and trend to all the data. Identify the values of the smoothing constants (including the smoothing constant for seasonal factors) that CB Predictor concludes are the best values. Also obtain forecasts (as of the end of January) of the average number of passengers in each subsequent month of the third year.

27.8-17. Quality Bikes is a wholesale firm that specializes in the distribution of bicycles. In the past, the company has maintained ample inventories of bicycles to enable filling orders immediately, so informal rough forecasts of demand were sufficient to make the decisions on when to replenish inventory. However, the company's new president, Marcia Salgo, intends to run a tighter ship. Scientific inventory management is to be used to reduce inventory levels and minimize total variable inventory costs. At the same time, Marcia has ordered the development of a computerized forecasting system based on statistical forecasting that considers seasonal effects. The system is to generate three sets of forecasts—one based on the moving-average method, a second based on the exponential smoothing method, and a third based on exponential smoothing with trend. The average of these three forecasts for each month is to be used for inventory management purposes.

The following table gives the available data on monthly sales of 10-speed bicycles over the past three years. The last column also shows monthly sales this year, which is the first year of operation of the new forecasting system.

Month	Past Sales			Current Sales This Year
	Year 1	Year 2	Year 3	
January	352	317	338	364
February	329	331	346	343
March	365	344	383	391
April	358	386	404	437
May	412	423	431	458
June	446	472	459	494
July	420	415	433	468
August	471	492	518	555
September	355	340	309	387
October	312	301	335	364
November	567	629	594	662
December	533	505	527	581

- T (a) Determine the seasonal factors for the 12 months based on past sales.
- T (b) After considering seasonal effects, apply the moving-average method based on the most recent three months to forecast monthly sales this year.
- T (c) After considering seasonal effects, apply the exponential smoothing method to forecast monthly sales this year. Use an initial estimate of 420 and a smoothing constant of $\alpha = 0.2$.
- T (d) After considering seasonal effects, apply exponential smoothing with trend to forecast monthly sales this year. Use initial estimates of 420 for the expected value and 0 for the trend, along with smoothing constants of $\alpha = 0.2$ and $\beta = 0.2$.
- (e) Compare both the MAD and MSE values obtained in parts (b), (c), and (d).
- (f) Calculate the combined forecast for each month by averaging the forecasts for that month obtained in parts (b), (c), and (d). Then calculate the MAD for these combined forecasts.
- (g) Based on these results, what is your recommendation for how to do the forecasts next year?
- (h) Combining the data on past sales for years 1, 2, 3 and the current sales for this year, use CB Predictor to compare exponential smoothing with seasonality and exponential smoothing with seasonality and trend (with the best choice of smoothing constants), so that CB Predictor then can select the best of these methods for this particular set of sales data. Also obtain the forecasts (as of the end of this year) for each month of next year with the selected method.

27.8-18. Reconsider the sales data for a certain product given in Prob. 27.5-4. The company's management now has decided to recontinue incorporating seasonal effects into its forecasting procedure for this product because there does not appear to be a substantial seasonal pattern. Management also is concerned that exponential smoothing may not be the best forecasting method for this product and so has decided to test and compare several forecasting methods. Each method is to be applied retrospectively to the given data and then its MSE is to be calculated. The method with the smallest value of MSE will be chosen to begin forecasting.

Apply this retrospective test and calculate MSE for each of the following methods. (Also obtain the forecast for the upcoming quarter with each method.)

- T (a) The *moving-average* method based on the last four quarters, so start with a forecast for the fifth quarter.
- T (b) The *exponential smoothing* method with $\alpha = 0.1$. Start with a forecast for the third quarter by using the sales for the second quarter as the latest observation and the sales for the first quarter as the initial estimate.
- T (c) The *exponential smoothing method* with $\alpha = 0.3$. Start as described in part (b).
- T (d) The *exponential smoothing with trend* method with $\alpha = 0.3$ and $\beta = 0.3$. Start with a forecast for the third quarter by using the sales for the second quarter as the initial estimate of the *expected value* of the time series (A) and the difference (sales for second quarter *minus* sales for first quarter) as the initial estimate of the *trend* of the time series (B).

PROBLEMS

- (e) Compare MSE for these methods. Which one has the smallest value of MSE?

27.8-19. Follow the instructions of Prob. 27.8-18 for a product with the following sales history.

Quarter	Sales	Quarter	Sales	Quarter	Sales
1	546	5	647	9	736
2	528	6	594	10	724
3	530	7	665	11	813
4	508	8	630	12	—

27.10-1. Long a market leader in the production of heavy machinery, the Spellman Corporation recently has been enjoying a steady increase in the sales of its new lathe. The sales over the past 10 months are shown below.

Month	Sales	Month	Sales
1	430	6	514
2	446	7	532
3	464	8	548
4	480	9	570
5	498	10	591

Because of this steady increase, management has decided to use *causal forecasting*, with the month as the independent variable and sales as the dependent variable, to forecast sales in the coming months.

- (a) Plot these data on a two-dimensional graph with the month on the horizontal axis and sales on the vertical axis.
- T (b) Find the formula for the linear regression line that fits these data.
- (c) Plot this line on the graph constructed in part (a).
- (d) Use this line to forecast sales in month 11.
- (e) Use this line to forecast sales in month 20.
- (f) What does the formula for the linear regression line indicate is roughly the average growth in sales per month?

27.10-2. Reconsider Probs. 27.8-3 and 27.6-1. Since the number of applications for admission submitted to Ivy College has been increasing at a steady rate, causal forecasting can be used to forecast the number of applications in future years by letting the year be the independent variable and the number of applications be the dependent variable.

- (a) Plot the data for Years 1, 2, and 3 on a two-dimensional graph with the year on the horizontal axis and the number of applications on the vertical axis.
- (b) Since the three points in this graph line up in a straight line, this straight line is the linear regression line. Draw this line.
- T (c) Find the formula for this linear regression line.
- (d) Use this line to forecast the number of applications for each of the next five years (Years 4 through 8).
- (e) As these next years go on, conditions change for the worse at Ivy College. The favorable ratings in the national surveys that had propelled the growth in applications turn unfavorable.

Consequently, the number of applications turn out to be 6,300 in Year 4 and 6,200 in Year 5, followed by sizable drops to 5,600 in Year 6 and 5,200 in Year 7. Does it still make sense to use the forecast for Year 8 obtained in part (d)? Explain.

- T (f) Plot the data for all seven years. Find the formula for the linear regression line based on all these data and plot this line. Use this formula to forecast the number of applications for Year 8. Does the linear regression line provide a close fit to the data? Given this answer, do you have much confidence in the forecast it provides for Year 8? Does it make sense to continue to use a linear regression line when changing conditions cause a large shift in the underlying trend in the data?
- T (g) Apply exponential smoothing with trend to all seven years of data to forecast the number of applications in Year 8. Use initial estimates of 3,900 for the expected value and 700 for the trend, along with smoothing constants of $\alpha = 0.5$ and $\beta = 0.5$. When the underlying trend in the data stays the same, causal forecasting provides the best possible linear regression line (according to the method of least squares) for making forecasts. However, when changing conditions cause a shift in the underlying trend, what advantage does exponential smoothing with trend have over causal forecasting?

27.10-3. Reconsider Prob. 27.8-11. Despite some fluctuations from year to year, note that there has been a basic trend upward in the annual demand for copper ore over the past 10 years. Therefore, by projecting this trend forward, causal forecasting can be used to forecast demands in future years by letting the year be the independent variable and the demand be the dependent variable.

- (a) Plot the data for the past 10 years (Years 1 through 10) on a two-dimensional graph with the year on the horizontal axis and the demand on the vertical axis.
- T (b) Find the formula for the linear regression line that fits these data.
- (c) Plot this line on the graph constructed in part (a).
- (d) Use this line to forecast demand next year (Year 11).
- (e) Use this line to forecast demand in Year 15.
- (f) What does the formula for the linear regression line indicate is roughly the average growth in demand per year?
- (g) Use the linear regression procedure in the forecasting area of your IOR Tutorial to generate a graph of the data and the linear regression line. Then experiment with the data to see how the linear regression line shifts as you drag any of the data points up or down.

27.10-4. Luxury Cruise Lines has a fleet of ships that travel to Alaska repeatedly every summer (and elsewhere during other times of the year). A considerable amount of advertising is done each winter to help generate enough passenger business for that summer. With the coming of a new winter, a decision needs to be made about how much advertising to do this year.

The following table shows the amount of advertising (in thousands of dollars) and the resulting sales (in thousands of passengers booked for a cruise) for each of the past five years.

Amount of advertising (\$1,000s)	225	400	350	275	450
Sales (thousands of passengers)	16	21	20	17	23

- (a) To use causal forecasting to forecast sales for a given amount of advertising, what needs to be the dependent variable and the independent variable?
- (b) Plot the data on a graph.
- T (c) Find the formula for the linear regression line that fits these data. Then plot this line on the graph constructed in part (b).
- (d) Forecast the sales that would be attained by expending \$300,000 on advertising.
- (e) Estimate the amount of advertising that would need to be done to attain a booking of 22,000 passengers.
- (f) According to the linear regression line, about how much increase in sales can be attained on the average per \$1,000 increase in the amount of advertising?

27.10-5. Reconsider Prob. 27.10-4. Use the linear regression procedure in the forecasting area of your IOR Tutorial to generate the linear regression line. On the resulting graph that shows this line and the five data points (as blue dots), note that the leftmost data point, the middle data point, and the rightmost data point all lie very close to the line. You can see how the linear regression line shifts as any one of these data points moves up or down by moving your mouse onto the blue dot at this point and dragging it vertically.

For each of these three data points, determine whether the linear regression line shifts above this point or shifts below it or still passes essentially through it when the following change is made in one of these data points (but none of the others).

- (a) Change the sales from 16 to 19 when the amount of advertising is 225.
- (b) Change the sales from 23 to 26 when the amount of advertising is 450.
- (c) Change the sales from 20 to 23 when the amount of advertising is 350.

27.10-6. To support its large fleet, North American Airlines maintains an extensive inventory of spare parts, including wing flaps. The number of wing flaps needed in inventory to replace damaged wing flaps each month depends partially on the number of flying hours for the fleet that month, since increased usage increases the chances of damage.

The following table shows both the number of replacement wing flaps needed and the number of thousands of flying hours for the entire fleet for each of several recent months.

Thousands of flying hours	162	149	185	171	138	154
Number of wing flaps needed	12	9	13	14	10	11

- (a) Identify the dependent variable and the independent variable for doing causal forecasting of the number of wing flaps needed for a given number of flying hours.
- (b) Plot the data on a graph.

- T (c) Find the formula for the linear regression line.
- (d) Plot this line on the graph constructed in part (b).
- (e) Forecast the average number of wing flaps needed in a month in which 150,000 flying hours are planned.
- (f) Repeat part (e) for 200,000 flying hours.
- (g) Use the linear regression procedure in the forecasting area of your IOR Tutorial to generate a graph of the data and the linear regression line. Then experiment with the data to see how the linear regression line shifts as you drag any of the data points up or down.

T 27.10-7. Joe Barnes is the owner of Standing Tall, one of the major roofing companies in town. Much of the company's business comes from building roofs on new houses. Joe has learned that general contractors constructing new houses typically will subcontract the roofing work about 2 months after construction begins. Therefore, to help him develop long-range schedules for his work crews, Joe has decided to use county records on the number of housing construction permits issued each month to forecast the number of roofing jobs on new houses he will have 2 months later.

Joe has now gathered the following data for each month over the past year, where the second column gives the number of housing construction permits issued in that month and the third column shows the number of roofing jobs on new houses that were subcontracted out to Standing Tall in that month.

Month	Permits	Jobs	Month	Permits	Jobs
January	323	19	July	446	34
February	359	17	August	407	37
March	396	24	September	374	33
April	421	23	October	343	30
May	457	28	November	311	27
June	472	32	December	277	22

Use a causal forecasting approach to develop a forecasting procedure for Joe to use hereafter.

27.10-8. The following data relate road width x and accident frequency y . Road width (in feet) was treated as the independent variable, and values y of the random variable Y , in accidents per 10^8 vehicle miles, were observed.

Number of Observations = 7		x	y
$\sum_{i=1}^7 x_i = 354$	$\sum_{i=1}^7 y_i = 481$	26	92
		30	85
		44	78
$\sum_{i=1}^7 x_i^2 = 19,956$	$\sum_{i=1}^7 y_i^2 = 35,451$	50	81
		62	54
		68	51
$\sum_{i=1}^7 x_i y_i = 22,200$		74	40

Assume that Y is normally distributed with mean $A + Bx$ and constant variance for all x and that the sample is random. Interpolate if necessary.

- (a) Fit a least-squares line to the data, and forecast the accident frequency when the road width is 55 feet.
- (b) Construct a 95 percent prediction interval for Y_+ , a future observation of Y , corresponding to $x_+ = 55$ feet.
- (c) Suppose that two future observations on Y , both corresponding to $x_+ = 55$ feet, are to be made. Construct prediction intervals for both of these observations so that the probability is at least 95 percent that both future values of Y will fall into them simultaneously. [Hint: If k predictions are to be made, such as given in part (d), each with probability $1 - \alpha$, then the probability is at least $1 - k\alpha$ that all k future observations will fall into their respective intervals.]
- (d) Construct a simultaneous tolerance interval for the future value of Y corresponding to $x_+ = 55$ feet with $P = 0.90$ and $1 - \alpha = 0.95$.

T 27.10-9. The following data are observations y_i on a dependent random variable Y taken at various levels of an independent variable x . [It is assumed that $E(Y_i|x_i) = A + Bx_i$, and the Y_i are independent normal random variables with mean 0 and variance σ^2 .]

x_i	0	2	4	6	8
y_i	0	4	7	13	16

- (a) Estimate the linear relationship by the method of least squares, and forecast the value of Y when $x = 10$.
- (b) Find a 95 percent confidence interval for the expected value of Y at $x^* = 10$.
- (c) Find a 95 percent prediction interval for a future observation to be taken at $x_+ = 10$.
- (d) For $x_+ = 10$, $P = 0.90$, and $1 - \alpha = 0.95$, find a simultaneous tolerance interval for the future value of Y_+ . Interpolate if necessary.

T 27.10-10. If a particle is dropped at time $t = 0$, physical theory indicates that the relationship between the distance traveled r and the time elapsed t is $r = gt^k$ for some positive constants g and k . A transformation to linearity can be obtained by taking logarithms:

$$\log r = \log g + k \log t.$$

By letting $y = \log r$, $A = \log g$, and $x = \log t$, this relation becomes $y = A + kx$. Due to random error in measurement, however, it can be stated only that $E(Y|x) = A + kx$. Assume that Y is normally distributed with mean $A + kx$ and variance σ^2 .

A physicist who wishes to estimate k and g performs the following experiment: At time 0 the particle is dropped. At time t the distance r is measured. He performs this experiment five times, obtaining the following data (where all logarithms are to base 10).

$y = \log r$	$x = \log t$
-3.95	-2.0
-2.12	-1.0
0.08	0.0
2.20	+1.0
3.87	+2.0

- (a) Obtain least-squares estimates for k and $\log g$, and forecast the distance traveled when $\log t = +3.0$.
- (b) Starting with a forecast for $\log r$ when $\log t = 0$, use the exponential smoothing method with an initial estimate of $\log r = -3.95$ and $\alpha = 0.1$, that is,

$$\text{Forecast of } \log r \text{ (when } \log t = 0) = 0.1(-2.12) + 0.9(-3.95),$$

- to forecast each $\log r$ for all integer $\log t$ through $\log t = +3.0$.
- (c) Repeat part (b), except adjust the exponential smoothing method to incorporate a trend factor into the underlying model as described in Sec. 27.6. Use an initial estimate of trend equal to the slope found in part (a). Let $\beta = 0.1$.

27.10-11. Suppose that the relation between Y and x is given by

$$E(Y|x) = Bx,$$

where Y is assumed to be normally distributed with mean Bx and known variance σ^2 . Also n independent pairs of observations are taken and are denoted by $x_1, y_1; x_2, y_2; \dots; x_n, y_n$. Find the least-squares estimate of B .

CASE

CASE 27.1 Finagling the Forecasts

Mark Lawrence—the man with two first names—has been pursuing a vision for more than two years. This pursuit began when he became frustrated in his role as director of human resources at Cutting Edge, a large company manufacturing computers and computer peripherals. At that time, the human resources department under his direction provided records and benefits administration to the 60,000 Cutting Edge employees throughout the United States, and 35 separate records and benefits administration centers existed across the country. Employees contacted these records and benefits centers to obtain

information about dental plans and stock options, to change tax forms and personal information, and to process leaves of absence and retirements. The decentralization of these administration centers caused numerous headaches for Mark. He had to deal with employee complaints often since each center interpreted company policies differently—communicating inconsistent and sometimes inaccurate answers to employees. His department also suffered high operating costs, since operating 35 separate centers created inefficiency.

His vision? To centralize records and benefits administration by establishing one administration center. This centralized records and benefits administration center would perform

two distinct functions: data management and customer service. The data management function would include updating employee records after performance reviews and maintaining the human resource management system. The customer service function would include establishing a call center to answer employee questions concerning records and benefits and to process records and benefits changes over the phone.

One year after proposing his vision to management, Mark received the go-ahead from Cutting Edge corporate headquarters. He prepared his “to do” list—specifying computer and phone systems requirements, installing hardware and software, integrating data from the 35 separate administration centers, standardizing record-keeping and response procedures, and staffing the administration center. Mark delegated the systems requirements, installation, and integration jobs to a competent group of technology specialists. He took on the responsibility of standardizing procedures and staffing the administration center.

Mark had spent many years in human resources and therefore had little problem with standardizing record-keeping and response procedures. He encountered trouble in determining the number of representatives needed to staff the center, however. He was particularly worried about staffing the call center since the representatives answering phones interact directly with customers—the 60,000 Cutting Edge employees. The customer service representatives would receive extensive training so that they would know the records and benefits policies backward and forward—enabling them to answer questions accurately and process changes efficiently. Overstaffing would cause Mark to suffer the high costs of training unneeded representatives and paying the surplus representatives the high salaries that go along with such an intense job. Understaffing would cause Mark to continue to suffer the headaches from customer complaints—something he definitely wanted to avoid.

The number of customer service representatives Mark needed to hire depends on the number of calls that the records

and benefits call center would receive. Mark therefore needed to forecast the number of calls that the new centralized center would receive. He approached the forecasting problem by using judgmental forecasting. He studied data from one of the 35 decentralized administration centers and learned that the decentralized center had serviced 15,000 customers and had received 2,000 calls per month. He concluded that since the new centralized center would service four times the number of customers—60,000 customers—it would receive four times the number of calls—8,000 calls per month.

Mark slowly checked off the items on his “to do” list, and the centralized records and benefits administration center opened one year after Mark had received the go-ahead from corporate headquarters.

Now, after operating the new center for 13 weeks, Mark’s call center forecasts are proving to be terribly inaccurate. The number of calls the center receives is roughly three times as large as the 8,000 calls per month that Mark had forecasted. Because of demand overload, the call center is slowly going to hell in a handbasket. Customers calling the center must wait an average of 5 minutes before speaking to a representative, and Mark is receiving numerous complaints. At the same time, the customer service representatives are unhappy and on the verge of quitting because of the stress created by the demand overload. Even corporate headquarters has become aware of the staff and service inadequacies, and executives have been breathing down Mark’s neck demanding improvements.

Mark needs help, and he approaches you to forecast demand for the call center more accurately.

Luckily, when Mark first established the call center, he realized the importance of keeping operational data, and he provides you with the number of calls received on each day of the week over the last 13 weeks. The data (shown below) begins in week 44 of the last year and continues to week 5 of the current year. Mark indicates that the days where no calls were received were holidays.

	Monday	Tuesday	Wednesday	Thursday	Friday
Week 44	1,130	851	859	828	726
Week 45	1,085	1,042	892	840	799
Week 46	1,303	1,121	1,003	1,113	1,005
Week 47	2,652	2,825	1,841	0	0
Week 48	1,949	1,507	989	990	1,084
Week 49	1,260	1,134	941	847	714
Week 50	1,002	847	922	842	784
Week 51	823	0	0	401	429
Week 52/1	1,209	830	0	1,082	841
Week 2	1,362	1,174	967	930	853
Week 3	924	954	1,346	904	758
Week 4	886	878	802	945	610
Week 5	910	754	705	729	772

- (a) Mark first asks you to forecast daily demand for the next week using the data from the past 13 weeks. You should make the forecasts for all the days of the next week now (at the end of week 5), but you should provide a different forecast for each day of the week by treating the forecast for a single day as being the actual call volume on that day.
- (1) From working at the records and benefits administration center, you know that demand follows “seasonal” patterns within the week. For example, more employees call at the beginning of the week when they are fresh and productive than at the end of the week when they are planning for the weekend. You therefore realize that you must account for the seasonal patterns and adjust the data that Mark gave you accordingly. What is the seasonally adjusted call volume for the past 13 weeks?
 - (2) Using the seasonally adjusted call volume, forecast the daily demand for the next week using the last-value forecasting method.
 - (3) Using the seasonally adjusted call volume, forecast the daily demand for the next week using the averaging forecasting method.
 - (4) Using the seasonally adjusted call volume, forecast the daily demand for the next week using the moving-average forecasting method. You decide to use the five most recent days in this analysis.
 - (5) Using the seasonally adjusted call volume, forecast the daily demand for the next week using the exponential smoothing forecasting method. You decide to use a smoothing constant of 0.1 because you believe that demand without seasonal effects remains relatively stable. Use the daily call volume average over the past 13 weeks for the initial estimate.
- (b) After 1 week, the period you have forecasted passes. You realize that you are able to determine the accuracy of your forecasts because you now have the actual call volumes from the week you had forecasted. The actual call volumes are shown next.

	Monday	Tuesday	Wednesday	Thursday	Friday
Week 6	723	677	521	571	498

For each of the forecasting methods, calculate the mean absolute deviation for the method and evaluate the performance of the method. When calculating the mean absolute deviation, you should use the actual forecasts you found in part (a) above. You should not recalculate the forecasts based on the actual values. In your evaluation, provide an explanation for the effectiveness or ineffectiveness of the method.

- (c) You realize that the forecasting methods that you have investigated do not provide a great degree of accuracy, and you decide to use a creative approach to forecasting that combines the statistical and judgmental approaches. You know that Mark had used data from one of the 35 decentralized records and benefits administration centers to perform his original forecasting. You therefore suspect that call volume data exist for this decentralized center. Because the decentralized centers performed the same functions as the new centralized center currently performs, you decide that the call volumes from the decentralized center will help you forecast the call volumes for the new centralized center. You simply need to understand how the decentralized volumes relate to the new centralized volumes. Once you understand this relationship, you can use the call volumes from the decentralized center to forecast the call volumes for the centralized center.

You approach Mark and ask him whether call center data exist for the decentralized center. He tells you that data exist, but they do not exist in the format that you need. Case volume data—not call volume data—exist. You do not understand the

distinction, so Mark continues his explanation. There are two types of demand data—case volume data and call volume data. Case volume data count the actions taken by the representatives at the call center. Call volume data count the number of calls answered by the representatives at the call center. A case may require one call or multiple calls to resolve it. Thus, the number of cases is always less than or equal to the number of calls.

You know you only have case volume data for the decentralized center, and you certainly do not want to compare apples and oranges. You therefore ask if case volume data exist for the new centralized center. Mark gives you a wicked grin and nods his head. He sees where you are going with your forecasts, and he tells you that he will have the data for you within the hour.

At the end of the hour, Mark arrives at your desk with two data sets: weekly case volumes for the decentralized center and weekly case volumes for the centralized center. You ask Mark if he has data for daily case volumes, and he tells you that he does not. You therefore first have to forecast the weekly demand for the next week and then break this weekly demand into daily demand.

The decentralized center was shut down last year when the new centralized center opened, so you have the decentralized case data spanning from week 44 of two years ago to week 5 of last year. You compare this decentralized data to the centralized data spanning from week 44 of last year to week 5 of this year. The weekly case volumes are shown in the table below.

	Decentralized Case Volume	Centralized Case Volume
Week 44	612	2,052
Week 45	721	2,170
Week 46	693	2,779
Week 47	540	2,334
Week 48	1,386	2,514

	Decentralized Case Volume	Centralized Case Volume
Week 49	577	1,713
Week 50	405	1,927
Week 51	441	1,167
Week 52/1	655	1,549
Week 2	572	2,126
Week 3	475	2,337
Week 4	530	1,916
Week 5	595	2,098

- (1) Find a mathematical relationship between the decentralized case volume data and the centralized case volume data.
- (2) Now that you have a relationship between the weekly decentralized case volume and the weekly centralized case volume, you are able to forecast the weekly case volume for the new center. Unfortunately, you do not need the weekly case volume; you need the daily call volume. To calculate call volume from case volume, you perform further analysis and determine that each case generates an average of 1.5 calls. To calculate daily call volume from weekly call volume, you decide to use the seasonal factors as conversion factors. Given the following case volume data from the decentralized center for week 6 of last year, forecast the daily call volume for the new center for week 6 of this year.

	Week 6
Decentralized case volume	613

- (3) Using the actual call volumes given in part (b), calculate the mean absolute deviation and evaluate the effectiveness of this forecasting method.
- (d) Which forecasting method would you recommend Mark use and why? As the call center continues its operation, how would you recommend improving the forecasting procedure?

(Note: Data files for this case are provided on the book's website for your convenience.)