Chapter **Twelve**

Forecast/Analytics Implementation

VLADGRIN/Getty Images

In this chapter, we discuss the forecasting/analytics process and provide a framework that will help you get the most out of any prediction effort. While every prediction problem has unique features, there is enough commonality in forecasting/analytics that guidelines can be helpful in several ways. First, the guidelines we provide will help you come to grips with some of the nuts-and-bolts issues related to data problems. Second, these guidelines will help you in making certain that the effort that goes into forecasting and analytics has the desired result in terms of the decision process. Finally, the guidelines discussed in this chapter will help you make logical choices regarding the technique(s) you should use for any particular situation.

LEARNING OBJECTIVES

After studying this chapter, you should be able to:

1. Discuss how forecasting has evolved from being purely judgmental to using highly complex methods and how predictive analytics can be considered an extension of forecast methods.
2. Explain the flow of the prediction process from raw data to action.
3. Discuss the two groups that must communicate well concerning forecasts and explain why this communication is important.
4. Explain the nine-step forecast/analytics process.
5. Explain the two major areas to consider when selecting a prediction method.

FORECASTING INVOLVES A DEFINITE FLOW

Both quantitative and qualitative information should be valued and, when possible, combined in preparing a forecast.

You have now learned numerous quantitative forecast methods. You have spent considerable time and effort developing a working knowledge of many quantitative techniques and how they can be implemented using a software package. Our own personal experiences, as well as the experiences of others, provide convincing evidence that quantitative forecasting and analytics methods outperform qualitative predictions. However, the best software cannot automatically take into account the specific industry, marketing, and economic knowledge that a business professional may have. To obtain the best forecast outcomes, both quantitative and qualitative information should be valued and, when possible, combined in preparing a forecast.

It is important for everyone involved with forecasting and analytics to be clear about the distinction between forecasts, plans, and goals. In a recent discussion, a veteran forecaster in the automobile industry commented: "I prepared what I thought was a logical and well-thought-out forecast, but when it was presented to management, the response was that the forecast was wrong and that I should go back and redo it." In this individual's case, what management wanted was a plan (what the company intends to do) or a goal (the company target) rather than an objective projection of what is likely, given the current business environment. This scenario is not uncommon. What it points out is a serious confusion on the part of many between a prediction (or projection), a plan, and a goal. The prediction should be one piece of objective information that plays a part in the development of plans and/or goals, but it should not be confused with the planning or goal-setting functions.

The forecast should be one piece of objective information that plays a part in the development of plans and/or goals, but it should not be confused with the planning or goal-setting functions.

The emergence of widely available and sophisticated forecasting software (such as the ForecastX Wizard™ you have been using) and predictive analytics software (such as XLMiner©) has made it possible for people to implement complex forecasting and analytics methods quickly and easily. However, there is danger in implementing a technique about which one does not have a reasonable level of understanding. For example, suppose that you are a brand manager who has some forecasting responsibility for certain brands but that this function is only about 10 percent of your overall workload. In this situation, you might be inclined to make relatively simple judgmental forecasts, or if you have come to realize that quantitative methods can improve forecast accuracy, you might be tempted to use an automated forecast "black box" to develop your forecasts. In either case, you are likely to have difficulty explaining and/or justifying the forecast to those to whom you report. However, if you have a basic understanding of forecast methods (which you have now developed), you can articulate the reasoning behind your forecast and how the quantitative methods employed are well suited to the type of data that represent sales of your products. You will be able to make qualitative judgments and adjustments to the forecasts and be able to explain why such adjustments may be necessary. You may not be able to derive the formulas for the Winters' exponential smoothing model or for developing an ARIMA forecast that does not overfit, but you know enough about how these methods work to know when they are appropriate.

Communication, cooperation, and collaboration are important if the forecasting effort is to be successful. Many times, the people who develop a prediction do so in a vacuum of sorts. They look at the data and prepare a forecast, which is then sent to users who have had little or no input into the forecast process. The forecast may not be in a form that is useful to the end user, or the units forecast may be inappropriate for their use, or the wrong series may have been forecast, or they may simply not have enough understanding of the forecast to use it properly.

Two particular groups that need to communicate well are the analysts (or data scientists) and the end users of a forecast (people in sales, marketing, finance, production and others). Each of these groups may have quite different perspectives on the forecasting process and the desired results. Collaboration among interested parties is essential for the forecasting process to truly meet an organizations' needs.

There may be substantial savings if inventory levels can be reduced.

For collaborative forecasting to be successful, all parties need to work together by treating the perspectives and biases of others as valuable inputs rather than as obstacles to overcome.[1] These days, the need for communication, cooperation, and collaboration goes beyond company boundaries. To maximize the benefits to be derived from the forecast process or analytics effort, communication, cooperation, and collaboration should involve the entire supply chain.

Everyone is well aware that inventory is expensive and there may be substantial savings if inventory levels can be reduced. Such reduction was the premise upon which "just-in-time" processes were developed.

THE FORECAST PROCESS

Data are the foundation, or base, for all predictions. There was a time when only numeric data were considered useful. Now useful data come in many forms. Much of the data used is still numeric, and the volume of available numeric data seems to grow almost without bounds. The "Internet of things" (IoT) drives a good part of this increased flow of numeric data. Cars talk to computers, refrigerators talk to tablets and phones, phones talk to thermostats, and on it goes. Sensors in one business can talk to sensors in another business, facilitating real time knowledge about inventories, production schedules, delays, and any other events that may affect forecasts. Adjustments can be made on the fly, saving time and money all along a supply chain.

An important way to look at the forecasting process is shown in Figure 12.1. Data are the starting point for all forecasts. Having reliable, clean, accurate data

[1] Sean Reese, "The Human Aspects of Collaborative Forecasting," *Journal of Business Forecasting*, Winter 2000–2001, pp. 3–9.

FIGURE 12.1
Data to Action in Forecasting
Forecasting begins with data that eventually drives actions.

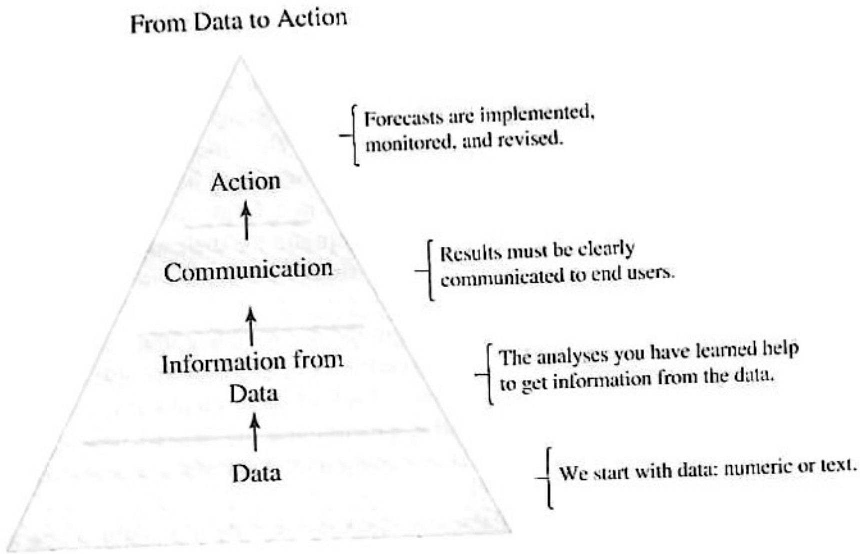

forms the bedrock of a successful forecasting process. Often people confuse data with information. Data are not information. Data are simply the raw materials that allow us to build forecasts. The raw data contain information, but the information is hidden in a forest of detail. The various tools you have learned allow you to glean information from the clutter of the details in the data.

Think about the thousands of SKUs that companies such as Walmart and Amazon must order, track, and sell. The sheer volume of the related data is so large that no human can make sense of what meaning is hidden in the vastness of the detail. This is where data analysis comes into play. Using various tools, one can dig into the data to find the information that is necessary to make sound business decisions.

A key role for analysts is to be able to communicate succinctly and with clarity the information derived from the data. This usually means no complex equations, limited quantitative jargon, and relatively few numbers (certainly not all the statistics your software can compute). Visualization using various graphics is the key to helping others understand the results of complex analyses. All predictions should be communicated in a manner that is easily understood by the end user. Long, complex tables will numb the mind. Try to boil the important results down to the essentials and present that information with clarity in terms the end user understands. Almost always, graphics are a great help in this communication process.

End users of forecasts will take action based on the forecast only if they trust and understand the forecast. Too often, analysts do a wonderful job of teasing information from the data, only to fail when it comes to communicating actionable results.

FIGURE 12.2
A Nine-Step Forecasting Process

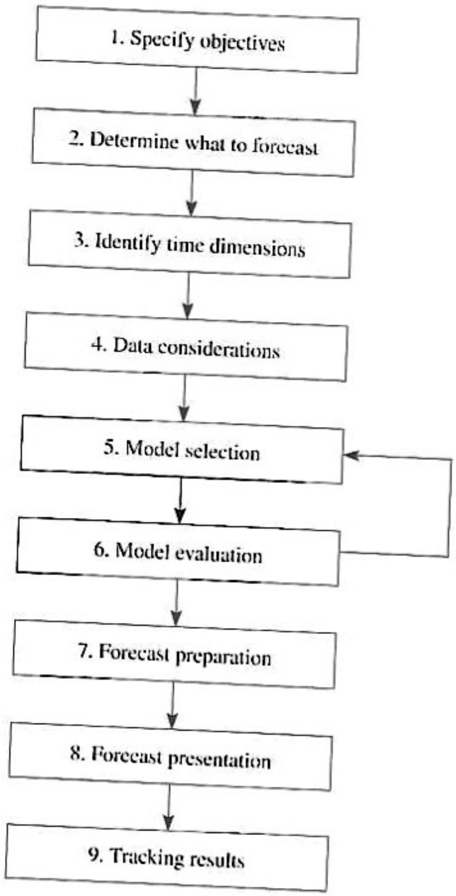

A Nine-Step Forecasting Process

Communication, cooperation, and collaboration are critical if forecasting is to have the desired positive effect on decisions.

In Chapter 2, we suggested a nine-step forecasting process. This is shown again in Figure 12.2. These steps begin and end with communication, cooperation, and collaboration between the managers who use the forecasts and the technicians who prepare them. This communication and cooperation are critical if forecasting is to have the desired positive effect on decisions.

Step 1. Specify Objectives

The objectives related to the decisions for which a prediction is important should be stated clearly. Management should articulate the role that the forecast will have in the decision process. If the decision will be the same regardless of the forecast, then any effort devoted to preparing the forecast is wasted. This may sound too obvious to deserve mention. However, it is not uncommon for a manager to request a forecast only to ignore it in the end. One reason that this happens is that the manager does not understand or have faith in the forecast. This issue will be addressed more fully in steps 7, 8, and 9, but a grounding of faith and understanding should begin here in step 1. If the manager who needs the information from a forecast and the data scientist who prepares the forecast take the opportunity to discuss the objectives and how the forecast will be used, there is increased likelihood that the ultimate forecast will be one that the manager understands and has faith in using.

Step 2. Determine What to Forecast

Once your overall objectives are clear, you must decide exactly what to forecast. For example, it is not sufficient to say that you want a sales forecast. Do you want a forecast of sales revenue or unit sales? Do you want an annual forecast or a quarterly, monthly, or weekly forecast? Is the prediction about which specific customers will "churn"? Is the forecast needed on a global basis, by global regions (such as Asia and North America), on a national level, on a sales territory level? What level of product aggregation is desired? All of these issues must be made clear at the start. It is generally better to base sales forecasts on units rather than dollars so that price changes do not cloud actual variations in unit sales. The unit sales forecast can then be converted to a dollar figure easily enough. If the effect of price on sales is important, you may want to use a regression-based technique that incorporates causality. Good communication between the forecast user and the analysts who prepare the forecast is important in making certain that the appropriate variables are being forecast.

Step 3. Identify Time Dimensions

There are two types of time dimensions to consider. First, one must establish the length of the forecast horizon. For annual forecasts, this might be from one to five years or more, although forecasts beyond a few years are likely to be influenced by unforeseen events that are not incorporated into the model used. Quarterly forecasts are probably best used for one or two years (four to eight quarters), as are monthly forecasts (perhaps as long as 12 to 18 months). The objectives dictate the time interval (year, quarter, and so forth) that is appropriate in preparing the forecast. For inventory control, short time periods are often necessary, whereas an annual forecast may be sufficient for the preparation of an estimated profit-and-loss statement for the coming year.

Second, the manager and the forecaster must agree on the urgency of the forecast. Is it needed tomorrow? Is there ample time to explore alternative methods? Proper planning is appropriate here. If a forecasting process is integrated into ongoing operations, then the forecasting personnel can plan an appropriate schedule, which will contribute to better forecasts.

Step 4. Data Considerations

The data necessary in preparing a forecast may come from within or may be external. Let us first consider internal data. Some people may believe that internal data are readily available and easy to incorporate into the forecasting process. It is surprising how often this turns out to be far from correct. Data may be available in a technical sense yet not readily available to the person who needs them to prepare the forecast. Or the data may be available but not expressed in the right unit of measurement (e.g., in sales dollars rather than units sold).

Data are often aggregated across both variables and time, but it is best to have disaggregated data. For example, data may be kept for refrigerator sales in total but not by type of refrigerator, type of customer, or region. In addition, what data are maintained may be kept in quarterly or monthly form for only a few years and annually thereafter. Such aggregation of data limits what can be forecast and may limit the appropriate pool of forecasting techniques. Data storage has become relatively inexpensive, so all data should be kept for as long as possible. Communication and cooperation among the personnel involved in database maintenance, forecast preparation, and forecast use can help alleviate many unnecessary problems in this regard.

External data are available from a wide variety of sources, many of which have been discussed in Chapter 1 and some of which are discussed in Chapter 5 (Prevedere, for example). Data from national, state, and local government agencies are generally available at low cost. The more local the level of government unit, the more likely it is that the data will not be available as quickly as you might like or in the desired detail. Other sources of secondary data include industry or trade associations and private companies, such as some of the major banks. Most secondary data are available on the Internet.

Step 5. Model Selection

There are many methods to select from when you set out to make any forecast. There are subjective or judgmental methods, some of which were reviewed in Chapter 1, and a growing set of quantitative methods is becoming available. The most widely used of these quantitative methods have been discussed in the previous chapters. The emergence of data and text mining in a forecasting environment is an exciting new horizon. Some of the things that should be included in making the selection are:

1. The type and quantity of data available
2. The pattern that the data have exhibited in the past
3. The urgency of the forecast
4. The length of the forecast horizon
5. The technical background of the people preparing and using the forecast

This issue of selecting the appropriate methods to use is of sufficient importance that we will come back to it in the next section. There, we provide specific guidelines for each of the methods discussed in the text.

Step 6. Model Evaluation

Once the methods that we want to use have been selected, we need to do some initial evaluation of how well they work. For quantitative methods, we should apply the techniques to historical series and evaluate how well they work in a retrospective sense. We have referred to this as an evaluation of the "fit" of the model. If they do not work well in the historical context, there is little reason to believe that they will perform any better in the unknown domain of the future.

If we have sufficient historical data, a good approach to model testing is to use a "holdout" period for evaluation. For example, suppose we have quarterly data on sales for 10 years. We might use only the earliest nine years (36 data points) and make a forecast for the tenth year. If the model performs well when the forecast values are compared with the known values for the four quarters of year 10, we have reason to believe that the technique may also work well when the forecast period is indeed unknown. Out-of-sample evaluations such as this provide a preliminary measure of potential forecast "accuracy." Recall that in predictive analytics, the concept of a "holdout" is built into the basic process as you always partition the data into "training" and "validation" partitions. It is the lift and misclassification results for the validation partition that are most useful for model evaluation.

Once you are satisfied with a model based on historical and holdout period (or validation partition) evaluations, you should respecify the model using all the available data (historical and holdout) and then use it for your actual forecast.

Once you are satisfied with a model based on historical and holdout period evaluations, you should respecify the model using all the available data (historical and holdout) and then use it for your actual forecast.

Suppose a technique turns out not to perform well when tested. The purpose of testing is, at least in part, to help us avoid applying a method that does not work well in our unique situation. Therefore, we should go back to step 5 and select another method that is appropriate to the problem at hand. It is not always possible to tell ahead of time how well a particular method will actually perform in a specific forecasting environment. We can apply reasoned judgment to our initial selection, but ultimately, the proof is in the pudding. We must apply the method to see whether it performs adequately for the purpose at hand.

Step 7. Forecast Preparation

At this point, some method or set of methods has been selected for use in developing the forecast, and from testing, you have reasonable expectations that the methods will perform well. We recommend using more than one forecasting method when possible, and it is desirable for these to be of different types (e.g., a regression model and Holt's exponential smoothing, rather than two different regression models). The methods chosen should be used to prepare a range of forecasts. You might, for example, prepare a worst-case forecast, a best-case forecast, and a most-likely forecast. The latter may be based on a combination of forecasts (ensembles) developed by following the procedures suggested in the Appendix to Chapter 5. In predictive analytics, the concept of ensemble models is again built into the process and most of the commercial software; consider boosting and bagging as well as alternative algorithms. And always remember that more data (and better data) will almost always make better predictions.

Step 8. Forecast Presentation

For a forecast to be used as intended, it must be presented to management clearly, in a way that provides an understanding of how the forecast was obtained and that elicits confidence in the forecast. It does not matter how much work is put into

"The forecaster must be capable of communicating the findings in language which the functional managers can understand. . . ."

developing the forecast. It does not matter how confident the preparer is in the results. It does not matter how sophisticated the methodology may be. What matters is whether or not the manager understands and has confidence in the forecast. All too often, quantitative analyses are put on a shelf and do not play the role in decision making that they should, because the results are not effectively presented to management. Decades ago Mark J. Lawless, who has been involved with forecasting within a number of corporations, including Chrysler, NCR, Ponderosa, and Hanson Industries Housewares Group, commented that:

> In communicating the forecast results to management, the forecaster must be capable of communicating the findings in language which the functional managers can understand and which is compatible with the corporate culture.[2]

The forecast should be communicated to management both in written form and in an oral presentation. Visuals should dominate. The written document should be at a level that is appropriate to the reader. In most cases, the managers who read the forecast document will have little interest in technical matters. They need the information necessary so that they use the forecast appropriately. They do not need the amount of background and detail to be able to prepare the forecast themselves.

Tables should be kept relatively short. Rarely would it be desirable to include an entire history of the data used and historical forecasts. The most recent observations and forecasts are usually sufficient. The long series should, however, be shown graphically and should include both actual and forecast values. In such graphic displays, colors and/or patterns can be used effectively to distinguish actual and forecast values.

The oral presentation should follow the same form and be made at about the same level as the written document. Generous use should be made of flip charts, slides, overheads, or projections of computer displays to heighten interest and involvement in the presentation. This oral presentation provides an excellent opportunity for discussion and clarification, which helps the manager gain a more complete understanding of the forecast and confidence in its usefulness.

Step 9. Tracking Results

Neither the preparer nor the user is done with the forecast after the presentation and incorporation of results into the relevant decisions. The *process* continues. Deviations from the forecast and the actual events should be discussed in an open, objective, and positive manner. The objectives of such discussions should be to understand why errors occurred, to determine whether the magnitude of the errors was sufficient to have made a difference in the decisions that were based on the forecast, and to reevaluate the entire process with the intent of improving performance in the next round of forecasts. Input from both managers and technicians is important for the continual refinement of the forecasting process.

[2] Lawless, Mark J. "Effective Sales Forecasting: A Management Tool." *Journal of Business Forecasting* 9, no. 1 (Spring 1990), 10.

It is important to stress once more the critical role that communication and cooperation between managers and technicians play in building and maintaining a successful forecasting process. This is true whether forecasts are prepared in house or by outside suppliers. Without a commitment to communication and cooperation, it is not likely that any organization can get a maximum return from the forecasting effort.

CHOOSING A FORECASTING TECHNIQUE

There a great many techniques that can be used to make forecasts. In this text, you have learned about many of the most commonly used methods. These methods are not only commonly used, but they are also the basis upon which many other methods have been developed. You started learning about some subjective (qualitative or judgmental) methods that were once the core of forecasting. More recently, quantitative methods have come to dominate the forecasting landscape.

Quantitative methods have evolved from simplistic methods, such as moving averages, to highly complex tools, such as data and text mining. Figure 12.3 illustrates the increasing complexity of forecasting methods. With the increasing complexity, we also gain richer information.

Now that you have an understanding of a variety of forecasting techniques, you need a general framework that will help you determine when to use each method.

FIGURE 12.3 Complexity of and Information from Forecast Methods

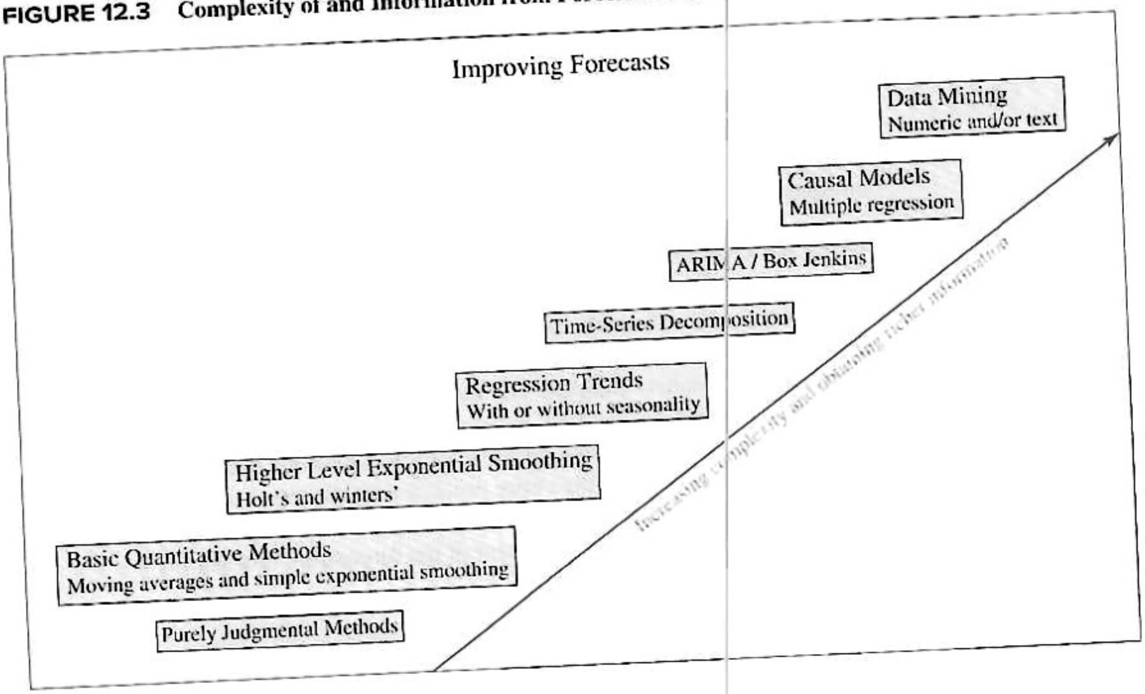

There are few hard-and-fast rules in this regard, but there are guidelines to assist in making the determination. If you understand how to use the methods discussed in this text, you have a good start toward determining when each method is likely to be useful. For example, if you are preparing a quarterly forecast of sales for a product that exhibits considerable seasonality, you would want to use one of the methods that is designed to handle such seasonal fluctuations.

In this section, we evaluate the forecasting methods presented earlier in the text relative to the underlying conditions for which they are most likely to be useful. There are many characteristics of a forecasting situation that might be considered in selecting an appropriate method. We will focus attention on two major areas: data and time. For data, we consider the type and quantity of data that are available as well as any pattern that may exist in the data (e.g., trend, cycle, and/or seasonality). Time includes the amount of historical data and the forecast horizon. We begin with the methods discussed in Chapter 1 and progress sequentially through the text, ending with data and text mining techniques. Table 12.1 provides a quick reference summary of the data and time issues.

Sales Force Composite (SFC)

In using the sales force composite method, little or no historical data are necessary. The data required are the current estimates of salespeople regarding expected sales for the forecast horizon. Historical data may be considered by the sales force, but not necessarily. Thus, this method may not reflect patterns in the data unless they are obvious to the sales force (e.g., holiday season sales of jewelry). The method may, however, provide early warning signals of pending change (positive or negative) because of the closeness of the sales force to the customer. The SFC method is probably best used for short- to medium-term forecasts.[3] The preparation time is relatively short once a system for gathering data from the sales force is in place.

Customer Surveys (CS)

Forecasts that are based on surveys of buyers' intentions require no historical data, and thus, the past plays no explicit role in forecasting the future. Customer surveys are most appropriate for medium- to long-term forecasting. For example, a natural gas utility has used this method to help in long-term planning by gathering survey data on customers' plans for future energy use, including long-term capital expansion plans. The time necessary to develop, conduct, and analyze a survey research project can be relatively extensive. Rarely can such a project be completed in less than two to three months. If the same survey is used year after year, however, this time can be shortened considerably. CS is not a method to consider if there is a sense of urgency in getting the forecast.

[3] Short-term, medium-term, and long-term forecasts will be mentioned throughout this section. Short-term forecasts include up to three months, medium-term forecasts cover four months to about two years, and long-term forecasts are for periods longer than two years

TABLE 12.1 Guide to Selecting an Appropriate Forecasting Method

Forecasting Method	Data Pattern	Quantity of Historical Data No. of Observations	Forecast Horizon
Subjective Methods			
Sales force composite	Any	Little	Short to medium
Customer surveys	Not applicable	None	Medium to long
Jury of executive opinion	Any	Little	Any
Delphi	Any	Little	Long
Quantitative Methods			
Naive	Stationary[a]	1 (or number of seasons)	Very short
Moving averages	Stationary[a]	Number equal to the number of periods in the moving average	Very short
Exponential Smoothing			
Simple	Stationary[a,b]	5 to 10	Short
Adaptive response	Stationary[a,b]	10 to 15	Short
Holt's	Linear trend[b]	10 to 15	Short to medium
Winters'	Trend and seasonality	At least 4 or 5 per season	Short to medium
Bass model	S-curve	Small, 3 to 10	Short, medium, and long
Regression-Based			
Trend	Linear and nonlinear Trend with or without seasonality	Minimum of 10 with 4 or 5 per season if data are seasonal	Short to medium
Causal	Can handle nearly all data patterns	We desire 10 per independent variable.	Short, medium, and long
Time-Series Decomposition			
	Can handle trend, seasonal, and cyclical patterns	Enough to see 2 peaks and troughs	Short, medium, and long
ARIMA			
	Stationary[a]	Minimum of 50	Short, medium, and long
Data Mining			
	Any	Used with large databases	Prediction usually for near-term use
Text Mining			
	Any	Used with large databases	Prediction usually for near-term use

[a] Or data that have been transformed to a stationary series.
[b] May be used for seasonal data if the data are first deseasonalized.

Jury of Executive Opinion (JEO)

The executives included do not need a formal data set. They need only the body of experience that they have developed to make judgments concerning the most likely value of the forecast variable during the period of interest. Historical data patterns may or may not be reflected in the opinions expressed, although regular patterns such as seasonality are very likely to receive attention, albeit implicit attention. A JEO may be used for any forecast horizon and is generally a relatively quick procedure. This method requires a substantial base of expertise on the part of the participants.

Delphi Method

The Delphi method does not require a historical data series, other than what is in the knowledge base of the panel members, and therefore does not necessarily reflect patterns that may have existed in the past. It is most often applicable for long-range forecasting but can be applied to medium-term projects as well. In these respects, it is much like JEO. However, the time to develop the Delphi forecast can be considerable unless the responses of panel members stabilize quickly. Use of the Internet speeds the flow of information and thus shortens the time considerably. The Delphi approach, as well as a jury of executive opinion and customer surveys, are sometimes useful in forecasting the sales of new products.

Naive

The basic naive model requires only one historical value as a basis for the forecast. An extended naive model that takes the most recent trend into account requires just two past values. This method is best suited to situations in which the data are stationary. Seasonality can sometimes be accounted for in a reasonably stationary series using a seasonal time lag. The naive approach is suited only for very short-term forecasts.

Moving Averages

Moving averages are most appropriate when the data are stationary and do not exhibit seasonality. Relatively few historical data are necessary. The number of past observations must be at least equal to the number of periods in the moving average. For example, if a four-period moving average is used, you need at least four historical data points. Moving averages are normally used to forecast just one period ahead.

Simple Exponential Smoothing (SES)

Historical data are necessary to establish the best weighting factor in simple exponential smoothing, but thereafter, only the most recent observed and forecasted values are required. Five to ten past values are sufficient to determine the weighting factor. The data series should be stationary (i.e., have no trend and no seasonality) when SES is used. This method is appropriate for short-term forecasting. While the arithmetic work can be done by hand, a computer can be helpful in determining the best weighting factor. Once the weighting factor is known, forecasts can be developed very quickly.

Adaptive–Response-Rate Single Exponential Smoothing (ADRES)

The adaptive–response-rate single exponential smoothing model may be used when the data are stationary and exhibit no seasonality but have a shift in level. Ten to 15 historical observations should be available when ADRES is used, and forecasts should be for only a short forecast horizon, typically one or two periods ahead.

Holt's Exponential Smoothing (HES)

As with SES, Holt's exponential smoothing model requires historical data to determine weighting values, but only the very recent past is required to apply the model. It is desirable to have at least 10 to 15 historical observations in determining the two weights. HES can be used effectively with data series that exhibit a positive or negative trend; thus, this method has a much wider scope of application than SES. However, it should not be used when the data contain a seasonal pattern unless the data have been deseasonalized. HES is appropriate for short- and medium-term forecasts and, like SES, can be implemented rapidly once the weights have been selected.

Winters' Exponential Smoothing (WES)

Sufficient historical data to determine the weights are necessary in using Winters' exponential smoothing model. A minimum of four or five observations per season should be used (i.e., for quarterly data, 16 or 20 observations should be used). Because this method incorporates both trend and seasonal components, it is applicable to a wide spectrum of data patterns. Like HES, this method is most appropriate for short- to medium-term forecasts. Once the weights have been determined, the process of making a forecast moves quickly. The use of professional forecasting software (such as ForecastX™) is recommended for the process of selecting the best values for the weights in the WES model.

Regression-Based Trend Models

The data requirement for using a regression-based trend depends to a considerable extent on the consistency in the trend and whether or not the trend is linear. We look for enough data that the t-statistic for the slope term (i.e., the trend) is significant (a t-value of 2 or more in absolute value is a handy rule of thumb). For a simple linear trend, 10 observations may be quite sufficient. A simple trend model can be effective when the series being forecast has no pattern other than the trend. Such a model is appropriate for short- to medium-term forecasts and can be developed and implemented relatively quickly.

Regression-Based Trend Models with Seasonality

To include seasonality in a regression-based trend model, it is desirable to have at least four or five observations per season. Thus, for quarterly data, a minimum of 16 observations would be appropriate. For monthly data, 48 or more observations should be used. Regular seasonal patterns in the series are often modeled quite well by using dummy variables. As with simple trend models, linear or nonlinear forms can be used; the models are best for short- to medium-term forecasts, and the time necessary for preparation is short. A computer regression program is a virtual necessity, however.

Regression Models with Causality

The quantity of data required for the development of a causal regression model depends on the number of independent variables in the model and on how much contribution each of those variables makes in explaining variation in the dependent variable. One rule of thumb is that you should expect to have a minimum of 10 observations per independent variable. Thus, for a model with three independent variables, you should have at least 30 observations. However, in practical applications, the length of the data set may be less. No matter how many observations are used, a statistical evaluation should be the guide for model acceptability. Developing and maintaining a database for multiple-regression models can be a significant undertaking. The effort may be worthwhile, however, since multiple-regression models are often effective in dealing with complex data patterns and may even help identify turning points. Seasonality can be handled by using dummy variables. Causal regression models can be useful for short-, medium-, or long-term forecasts. Because the causal variables must usually be forecast as well, regression models may take more effort to develop.

Time-Series Decomposition (TSD)

The quantity of data needed for time-series decomposition should be enough for you to see at least two peaks and two troughs in the cycle factor, if the cycle factor is important. If the cycle factor does not appear important (i.e., has not been far above or below 1.0 during the historical period), then the quantity of data needed should be determined by what is necessary to adequately identify the seasonal pattern. A rule of thumb would be at least four or five observations per season (e.g., for quarterly data, you should have at least 16 to 20 observations). TSD is quite good at picking up patterns in the data. The challenge is for the analyst to successfully project the patterns through the forecast horizon. This is generally fairly easy for the trend and seasonal pattern but more difficult for the cyclical pattern. TSD is especially appropriate for short-term and medium-term forecasting. If the cycle pattern is not important or if it can be projected with confidence, the method can also be used effectively for long-term forecasts. This method may be one of the best in terms of being able to identify and incorporate turning points. Doing so is dependent on the analyst's ability to correctly interpret when the cycle factor may turn up or down.

ARIMA

A long data series (at least 50 data points—more if data are seasonal) is necessary to make use of the ARIMA models. These models can handle variability in the data as long as the series is stationary or can be transformed to a stationary series. This method can be applied to short-, medium-, or long-term forecast horizons. Because of the complexity of model identification, forecast preparation can take an extended period of time. This complexity also means that the preparer needs a highly sophisticated technical background. Users of ARIMA forecasts must also be quite sophisticated, because even achieving a basic understanding of the method is not easy. It is rare to find a manager who has a good feel for how an ARIMA forecast is developed and rarer still to find a manager capable of explaining the forecast derivation to others who must use the results. This may be part of the reason that ARIMA models have often had relatively low ratings in terms of importance, accuracy, and use by business managers.

Data Mining

The data used in predictive analytics is usually not time series data like that used in most of the forecasting models. The size of the data set may also differ in that analytics depends more upon very large data sets; in most situations, the more data, the better the prediction. Two very different types of data mining algorithms were examined; classification algorithms and clustering algorithms. Classification is the most used type of analytics algorithm; it is used extensively in business and many different forms of classification were examined. Ensembles of different classification algorithms or modifications that created ensembles such as boosting and bagging are often used to improve accuracy. The clustering type models are less used but also have their place in business usefulness.

Text Mining

The types of data used in text mining were a stark departure from the data previously examined. Text was unstructured, not arranged in neat columns and rows with only numbers populating the various locations. Text, since it is so available for analysis, offered a new and larger frontier for prediction. The bag of words analysis that we examined is useful and will often yield useful results but that is only the frontier of examining unstructured data. Video, photo, and audio data will also offer possibilities for building predictive models. Text is just the tip of the data that data scientists will find useful and predictive in the future.

SPECIAL FORECASTING CONSIDERATIONS

In the text, a number of situations have been discussed for which special forecasting techniques are appropriate. Four of these are: (1) situations in which we must make forecasts if "events" of some type influence the forecast; (2) situations in which we have multiple forecasts, each of which may contain valuable information that we do not want to ignore; (3) situations in which we need to forecast a new product for which we have little historical information; and (4) situations in which we need to predict some outcome and we have very large, often somewhat unrelated, databases that hold hidden keys to the likely outcome. Here, we review some important aspects of each of these four.

Event Modeling

When forecasting sales or demand in a highly promoted market, using event modeling can often significantly improve forecast accuracy. Event modeling is a feature within many forecasting programs, such as ForecastX™. This feature allows the user to specify the time of one or more special events, such as irregular promotions and natural disasters, in the calibration data. For each type of special event, the effect is estimated and the data adjusted so that the events do not distort the trend and seasonal patterns of the time series.

The method of event modeling follows in the same pattern as the other smoothing models except that the event model adds a smoothing equation for each of the events being considered. Event models are analogous to seasonal models: just as each month is assigned its own index for seasonality, so, too, each event type is assigned its own index. Event adjustments are created through the use of an indicator variable that assigns an integer for each event type to the period during which it recurs. An example of integer value assignment would be that 0 indicates a period in which no event has occurred, 1 indicates a period in which a free-standing advertising insert was used, 2 indicates a period in which instantly redeemable coupons were used, and so on. The event indicator variable must be defined for each historical period and future period in the forecast horizon.

Combining Forecasts (Ensembles)

Instead of choosing the best model from among two or more alternatives, a more reasoned approach, according to the empirical evidence, is to combine the forecasts in order to obtain a forecast that is more accurate than any of the separate predictions. Any time a particular forecast is ignored because it is not the "best" forecast produced, it is likely that valuable independent information contained in the discarded forecast has been lost. The information lost may be of two types:

1. Some variables included in the discarded forecast may not be included in the "best" forecast.
2. The discarded forecast may make use of a type of relationship ignored by the "best" forecast.

In the first of these cases, it is quite possible for individual forecasts to be based on different information; thus, ignoring any one of these forecasts would necessarily exclude the explanatory power unique to the information included in the discarded model. In the second situation, it is often the case that different assumptions are made in different models about the form of the relationship between the variables. Each of the different forms of relationship tested, however, may have some explanatory value. Choosing only the "best" of the relationships could exclude functional information. To prevent this loss of useful information requires some method for combining the two forecasts into a single *better* forecast. We should expect that combinations of forecasts that use very different models are likely to be effective in reducing forecast error.

Combining forecasts is not guaranteed to reduce error, but one does not know until forecasts are combined and the combination results compared with results for individual forecasts. When combining forecasts, it is important that one check to be sure that doing so will not create a forecast bias.

New-Product Forecasting (NPF)

Most products for which we are likely to have to prepare a sales forecast are products with a substantial amount of sales history for which the methods you have learned in earlier chapters will work quite well. However, often we are faced with new, or substantially altered, products with little sales history. These new products pose particularly difficult issues for a forecaster. Understanding the concept of a product life cycle (PLC) can be helpful in developing a forecast for a new product. During the introductory stage of the product life cycle, only consumers who are *innovators* are likely to buy the product. Sales start low and increase slowly. Near the end of this stage, sales start to increase at an increasing rate. As the product enters the growth stage of the PLC, sales are still increasing at an increasing rate as *early adopters* enter the market. In this stage, the rate of growth in sales starts to decline. Near the end of the growth stage, sales growth starts to level off substantially as the product enters the maturity stage. Businesses may employ marketing strategies to extend this stage; however, all products eventually reach the stage of decline in sales and are, at some point, removed from the market.

Product life cycles are not uniform in exact shape or duration and vary from industry to industry. Think, for example, about products that are fashion items or fads in comparison with products that have real staying power in the marketplace. Fashion items and products that would be considered fads typically have a steep introductory stage followed by short growth and maturity stages and a decline that is also very steep. High-tech products, such as cell phones, often have life cycles that are relatively short in comparison with low-technology products. For high-tech electronic products, life cycles may be as short as six to nine months.

Methods such as analog forecasts, test marketing, and product clinics are often useful for new-product forecasting. The Bass model for sales of new products is probably the most notable model for new-product forecasting. The Bass model was originally developed for application only to durable goods. However, it has been adapted for use in forecasting a wide variety of products with short product life cycles and new products with limited historical data.

Data Mining

Sometimes people think of forecasting only in the context of time-series data. In some manner, past data are used to help predict the likely outcomes in the future. These include univariate time-series methods, such as exponential smoothing, as well as causal models, such as multiple regression. We have seen that at times regression models may be useful with cross-sectional data to predict some outcome, such as sales volume. Data mining is another technique that has been developed to help one predict outcomes when there is a great deal of data available that might contain hidden information.

Data mining techniques work often with very large and somewhat unrelated databases. There was a time when decision makers had too little data upon which to base decisions. Now that has changed dramatically, and decision makers have so much data that it is difficult to find the information content from the data. This is where data mining becomes a useful tool.

Data mining has become a new application for some types of forecasting in which we have huge amounts of data but we know little about the structural relationships contained therein. Data mining is a tool that helps us uncover relationships that are often quite unexpected yet useful in making predictions. For example, a California retailer found through data mining that shoppers who buy diapers are also likely to buy beer.[4] Such knowledge would not be likely to be uncovered using more simplistic data analysis but can be useful in predicting sales of both items and in developing new ways to structure marketing communications involving both products.

Suppose you wanted to forecast the number of sports cars an insurance company would insure. It is obvious to us that one factor would be the price (premium) charged for coverage, which in turn would be influenced by the number of claims filed by sports car owners. Conventional wisdom might suggest that sports car owners would have more claims for accidents and/or thefts. However, through data mining, Farmers Group found that sports cars owned by people who also owned another vehicle have fewer insurance claims. As a result, they restructured their premiums in these situations with a resulting increase in premium revenue of over $4 million in two years without having a substantial increase in claims.[5] It was only possible to make the prediction about the potential new market by using data mining.

Text Mining

Text mining is truly the frontier of prediction. The suggestions you receive from Amazon and the ability of firms to react almost immediately to customer complaints (or compliments) are likely due to some text mining algorithm. The uses of digital assistants such as Siri and Alexa are examples of text mining (in which the audio has been converted to text and then analyzed) that have become commonplace. As we obtain access to more data from the Internet of Things, sensors, and the mining of social media, the uses of text mining will grow. And with more data to work with, the usefulness of the predictions will likely become better.

[4] Donald R. Cooper and Pamela S. Schindler, *Marketing Research*, McGraw-Hill/Irwin, 2006, p. 261.

[5] Carl McDaniel and Roger Gates, *Marketing Research Essentials*, 6th ed. New York: John Wiley & Sons, 2008, pp. 79–80.

USING PROCAST™ IN FORECASTX™ TO MAKE FORECASTS

We generally recommend that you think carefully about your forecast objectives and about the nature of your data to select a forecast method. However, there are situations in which you might allow ForecastX™ to select the method for you. Suppose you have hundreds or thousands of SKUs to forecast. Rather than analyzing each individually, you might want to let the software take on that task. There is a way to do this within ForecastX™ by using a feature called **ProCast™**.

ProCast™ will search through a subset of methods to determine the method that will work best for your data. You can select the criteria to use in selecting the best method. In the example shown below, we selected to minimize the absolute error. There is a downside to using **ProCast™** in that you may not be familiar with the method selected and so would have trouble explaining it to an end user. Also, causal models would not be included in the decision set because ForecastX™ would not have access to the desired independent variables.

To use Procast™, begin by opening your data file in Excel. Place your cursor in any cell with the data to be forecast. In the example below, we selected cell B3.

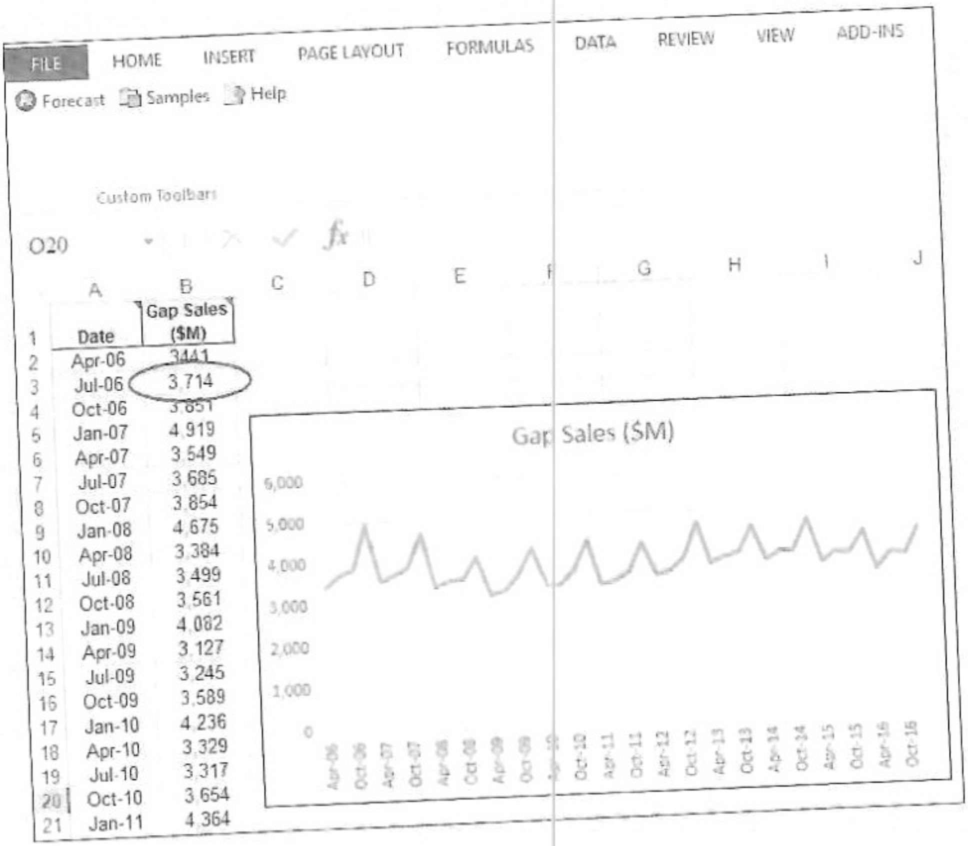

Then start ForecastX™. In the **Data Capture** dialog box, identify the data you want to use, as shown below.

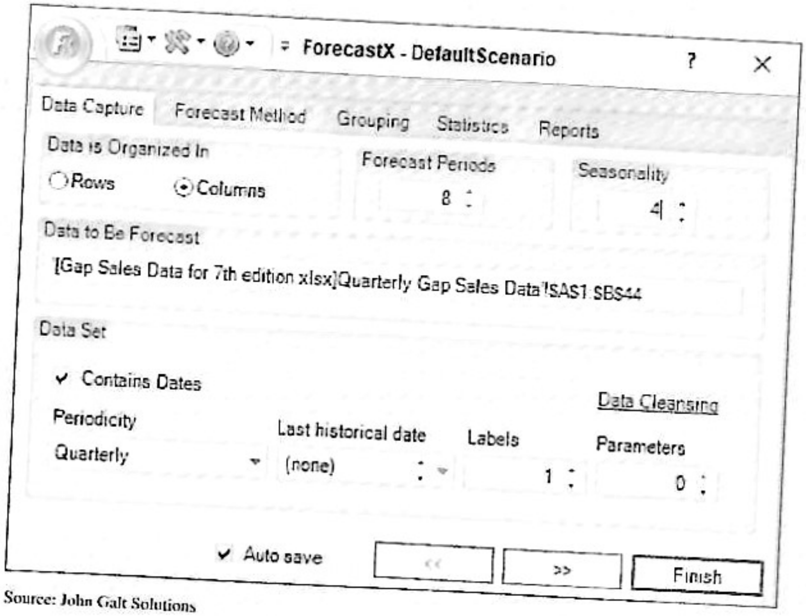

Source: John Galt Solutions

Then click the **Forecast Method** tab. In the **Method Selection** dialog box, click the down arrow in the **Forecasting Technique** box and select **ProCast™**. Click the down arrow in the **Error Term** box and select **Mean Absolute Error** (or another error term you want to use).

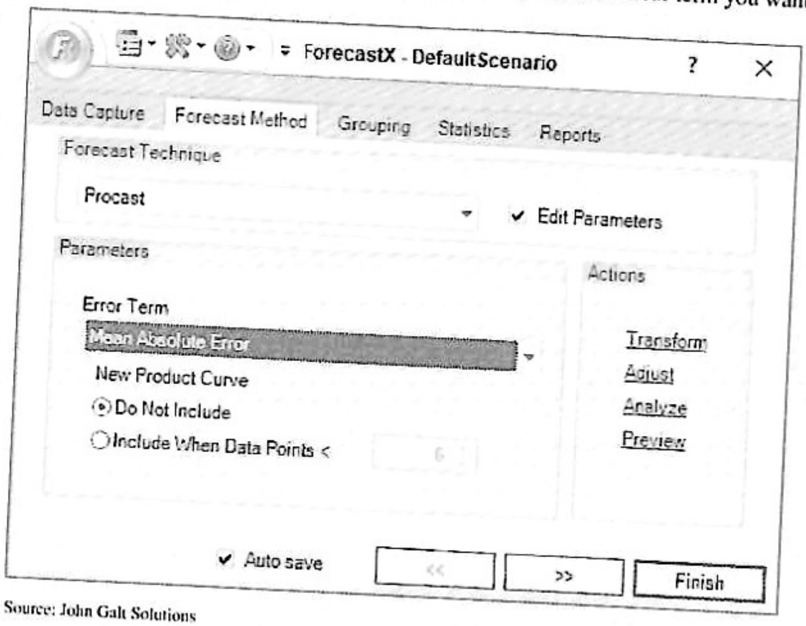

Source: John Galt Solutions

Then click the **Statistics** tab. In this dialog box, select the statistics that you desire. Remember that there are more statistics choices if you click the **More Statistics** button near the bottom right of the dialog box.

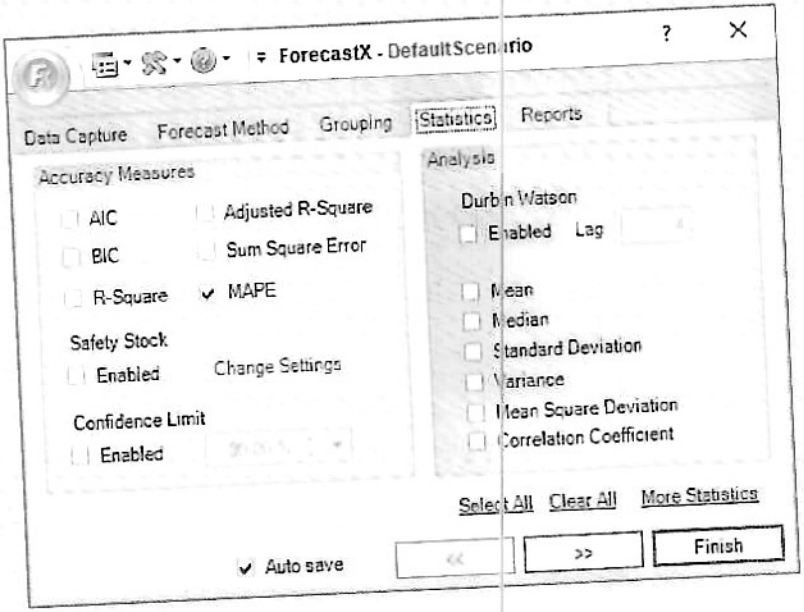

Source: John Galt Solutions

After selecting the statistics you want to see, click the **Reports** tab.

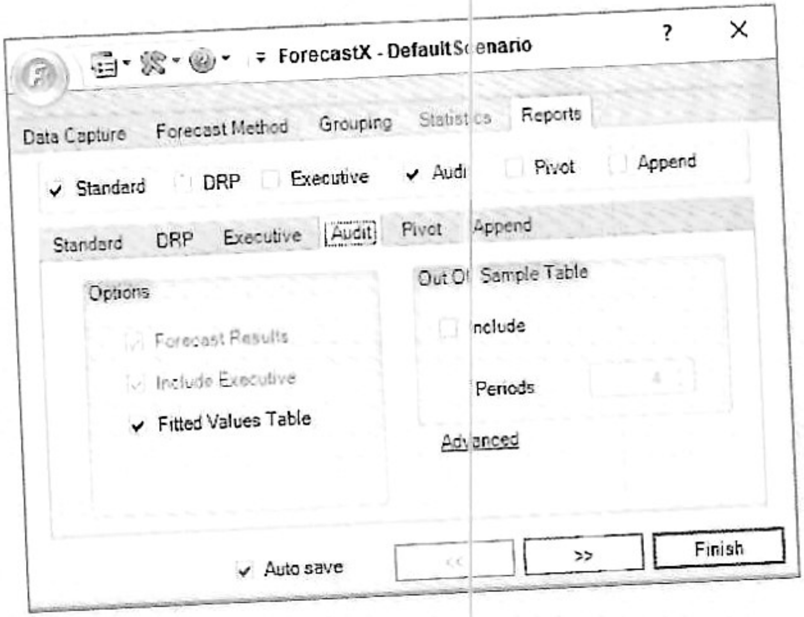

Source: John Galt Solutions

In the **Reports** box, select those you want. Typical selections might be those shown here. When you click the **Standard** tab, select the **Show Chart** and **Classic.** In the **Audit Trail** tab (the active tab shown here), click the **Fitted Values Table.**

Then click the **Finish** button. In the Audit Trail output, you will find the method that ProCast™ used to make the requested forecast.

Using an automated forecasting method such as ProCast™ is all right if you understand the selected method well enough to evaluate whether it is truly a logical choice. It is wise to exercise some caution when allowing any software to select a method automatically. By using a software package, such as ForecastX™, over a period of time, you may develop confidence in the selections it makes. Then using an automated process may provide considerable time savings—such as in situations where there are hundreds or thousands of items that must be forecast frequently.

Suggested Readings

Armstrong, J. Scott. "Research Needs in Forecasting." *International Journal of Forecasting* 4, no. 3 (1988), pp. 449–65.

Chase, Charles W., Jr. "Business Forecasting: A Process Not an Application." *Journal of Business Forecasting* 11, no. 3 (Fall 1992), pp. 12–13.

Fisher, Marshall; and Kumar Rajaram. "Accurate Retail Testing of Fashion Merchandise: Methodology and Application." *Marketing Science* 19, no. 3 (Summer 2000), pp. 266–78.

Harrington, Lisa H. "Retail Collaboration: How to Solve the Puzzle." *Transportation and Distribution*, May 2003, pp. 33–37.

Keating, Barry; and J. Holton Wilson. "Forecasting Practices and Teachings." *Journal of Business Forecasting* 7, no. 4 (Winter 1987–88), pp. 10–13, 16.

Larréché, Jean-Claude; and Reza Moinpour. "Managerial Judgement in Marketing: The Concept of Expertise." *Journal of Marketing Research* 20, no. 2 (May 1983), pp. 110–21.

Lawless, Mark J. "Effective Sales Forecasting: A Management Tool." *Journal of Business Forecasting* 9, no. 1 (Spring 1990), pp. 2–11.

———. "Ten Prescriptions for Forecasting Success." *Journal of Business Forecasting* 11, no. 4 (Spring 1997), pp. 3–5.

LeLee, Gary S. "The Key to Understanding the Forecasting Process." *Journal of Business Forecasting* 11, no. 4 (Winter 1992–93), pp. 12–16.

Lynn, Gary S.; Steven P. Schnaars; and Richard B. Skov. "Survey of New Product Forecasting Practices in Industrial High Technology and Low Technology Businesses." *Industrial Marketing Management* 28 (November 1999), pp. 565–71.

Mentzer, John T.; and Kenneth B. Kahn. "State of Sales Forecasting Systems in Corporate America." *Journal of Business Forecasting* 11, no. 4 (Spring 1997), pp. 6–13.

Moon, Mark A.; and John T. Mentzer. "Improving Salesforce Forecasting." *Journal of Business Forecasting* 18, no. 2 (Summer 1999), pp. 7–12.

Moon, Mark A.; John T. Mentzer; Carlo D. Smith; and Michael S. Garver. "Seven Keys to Better Forecasting." *Business Horizons* (September–October 1998), pp. 44–52.

Pammer, Scott E.; Duncan K. H. Fong; and Steven F. Arnold. "Forecasting the Penetration of a New Product—A Bayesian Approach." *Journal of Business & Economic Statistics* 18, no. 4 (October 2000), pp. 428–35.

Raghunathan, Srinivasan. "Interorganizational Collaborative Forecasting and Replenishment Systems and Supply Chain Implications." *Decision Sciences* 30, no. 4 (Fall 1999), pp. 1053–71.

Reese, Sean. "The Human Aspects of Collaborative Forecasting." *Journal of Business Forecasting* 19, no. 4 (Winter 2000–2001), pp. 3–9.

Reyes, Luis. "The Forecasting Function: Critical Yet Misunderstood." *Journal of Business Forecasting* 14, no. 4 (Winter 1995–96), pp. 8–9.

Szmania, Joe; and John Surgent. "An Application of an Expert System Approach to Business Forecasting." *Journal of Business Forecasting* 8, no. 1 (Spring 1989), pp. 10–12.

Tkacz, Greg. "Neural Network Forecasting of Canadian GDP Growth." *International Journal of Forecasting* 17, no. 1 (January–March 2001), pp. 57–69.

Weitz, Rob R. "NOSTRADAMUS—A Knowledge-Based Forecast Advisor." *International Journal of Forecasting* 2, no. 1 (1986), pp. 273–83.

Wilson, J. Holton; and Hugh G. Daubek. "Marketing Managers Evaluate Forecasting Models." *Journal of Business Forecasting* 8, no. 1 (Spring 1989), pp. 19–22.

Exercises

1. Discuss how forecasting has moved from purely judgmental methods to highly complex methods.
2. Explain the process of going from raw data to actions based on a forecast.
3. What two groups must communicate well in order for the forecast process to be effective? Explain why.
4. Describe the nine-step forecast process presented in the chapter.
5. What are the two main things to consider when selecting a forecast method? Why?
6. The forecast process begins with a need to make one or more decisions that depend on the future value of some variable. Think about this as it relates to the daily weather forecast you hear, and write a list of five decisions that might depend on such a forecast.
7. The availability and form of data to be used in preparing a forecast are often seen as especially critical areas. Summarize, in your own words, the database considerations in the forecasting process (step 4).
8. Suppose that you have been asked to recommend a forecasting technique that would be appropriate to prepare a forecast, given the following situational characteristics:
   a. You have 10 years of quarterly data.
   b. There is an upward trend to the data.
   c. There is a significant increase in sales prior to Christmas each year.
   d. A one-year forecast is needed.
   e. You need to have the forecast done and the presentation ready in just a few days.
   f. What method(s) would you consider using and why?
9. Write an outline of what you would like to see in a forecast presentation from the perspective of a manager who needs to use the forecast.
10. Explain how the predictions made using analytics are somewhat different than those from traditional forecasting models. Are the data used differently? Are the types of predictions different?