Forecast Pro TRAC User's Guide

Eric A. Stellwagen and Sarah Darin
Business Forecast Systems, Inc.

Software Version 5.1
Manual Last Revised: October 30, 2018

Business Forecast Systems, Inc.
68 Leonard Street, Belmont, MA 02478 USA
Phone: (617) 484-5050 ◆ Fax: (617) 484-9219
Email: info@forecastpro.com ◆ Web: www.forecastpro.com

Contents

GETTING STARTED

TUTORIAL: General Operations

TUTORIAL: Forecasting Operations

REFERENCE

1

GETTING STARTED

Chapter 1

Before You Begin

Forecasters and Collaborators

Forecast Pro is available in two different types of licenses—Forecaster licenses and Collaborator licenses. The type of license you install will dictate the functionality you can access. A Forecaster has access to all of the functionality described in this manual. A Collaborator can view and manipulate a forecast project created by a Forecaster (including viewing graphs and reports, adding overrides and comments, saving output, etc.) but cannot generate new statistical forecasts.

Operationally, the two installations are very similar. However, a Collaborator does not have access to any of the functionality relating to reading historical data and generating statistical forecasts. Thus, there are sections of this manual which will not be relevant to Collaborators.

How to Use This Manual

This manual is divided into four parts.

Part 1 Getting Started

Chapter 1 describes the Forecast Pro system requirements and product support policy.

Chapter 2 describes the installation procedure.

Chapter 3 provides an overview of forecasting.

Part 2 Tutorial: General Operations

Lesson 1 walks you through a sample forecasting session using automatic model selection.

Lesson 2 teaches you some techniques to efficiently navigate the program and customize reports.

Lesson 3 teaches you how to customize the forecast override view. It covers working with external data, creating calculated rows as functions of other rows and specifying how rows are displayed.

Lesson 4 teaches you how to work with and customize Forecast Pro graphs

Lesson 5 teaches you how to work with multiple forecasts and specify which ones to use as the baseline forecast.

Lesson 6 teaches you how to work with conversions and alternative hierarchies.

Lesson 7 teaches you how to update your forecast project as new data become available.

Lesson 8 teaches you how to collaborate with colleagues to establish the final forecast. It covers both sharing stand-alone forecast projects and using the super project manager in Forecast Pro to break large forecasting jobs into

smaller pieces that can be worked on separately and then reassembled into a "super project."

Lesson 9 teaches you how to archive your forecasts and use waterfall reports to track your forecast accuracy.

Lesson 10 teaches you how to use exception reports to identify poor forecasts and other potential problems in an efficient manner.

Lesson 11 teaches you how to use Pareto analysis to assign ABC classifications to your forecast items and how to sort and filter reports.

Lesson 12 teaches you how to operate the program from the command line.

Part 3 Tutorial: Forecasting Operations

Lesson 13 teaches you how to use forecast modifiers to dictate model selection.

Lesson 14 teaches you how to build event models to capture promotional effects.

Lesson 15 teaches you how to use Forecast Pro's outlier detection and correction functionality.

Lesson 16 teaches you how to set up and forecast a multiple-level hierarchy.

Lesson 17 teaches you how to prepare forecasts using the custom component model.

Lesson 18 teaches you how to forecast new products prior to historic data being available and also teaches you how to use the forecast by analogy and Bass models.

Lesson 19 teaches you how to build and evaluate dynamic regression models.

Lesson 20 teaches you how to use weights to adjust for trading day impacts or to forecast new products based on similar products.

Lesson 21 teaches you how to assess forecasting performance using a holdout sample approach.

All of the lessons use sample data provided with the software.

Part 4 Reference

Chapter 1 explains how to select the appropriate data format and how to set up your historic demand database.

Chapter 2 describes the optional input data that can be imported into Forecast Pro.

Chapter 3 describes the various reports and output files that can be saved.

Chapter 4 describes how to use the Data Manager to define your data sources and generate forecasts.

Chapter 5 discusses the strategies that Forecast Pro uses to produce consistent forecasts for hierarchical data.

Chapter 6 describes all available model specification dialog boxes, forecast modifiers and their functions.

Chapter 7 is a complete reference to Forecast Pro menus, options and commands.

Chapter 8 describes how Forecast Pro reconciles forecast adjustments between different levels of the forecasting hierarchy.

The Appendixes describe how to speed up processing and provide other supplementary information.

Many users won't want to work their way through each and every part of this manual. However, we urge you to go through at least Chapters 1-3 of Getting Started and Lessons 1-7 of the Tutorial. This will give you a good overview of Forecast Pro operations and capabilities.

When you are ready to set up your own data, please consult Chapter 1 of Part 4 Reference.

Statistical Reference Manual (PDF only)

The *Forecast Pro Statistical Reference Manual* describes the statistical techniques, statistics and strategies that are implemented in Forecast Pro. You will be able to produce accurate forecasts with Forecast Pro without reading this manual.

The manual is provided in electronic format (pdf). When Forecast Pro is installed, the file is copied into the program directory. You can access the file directly or via the Help menu in Forecast Pro.

Chapter 1 is a detailed statistical reference. It explains the theory behind Forecast Pro's statistical models, diagnostics and methodology.

Chapter 2 is a discussion of some of the general considerations that apply to automatic batch forecasting.

Checking Your Forecast Pro Package

Your Forecast Pro package should contain the following items:

This User's Manual. This book walks you through the installation and use of Forecast Pro.

One installation CD.

A warranty registration card.

A maintenance and support contract.

If your package is missing any of the above items please contact Business Forecast Systems, Inc.

What You Need to Run Forecast Pro

A computer running a Windows operating system.

A minimum of 512 MB of random access memory (RAM). Additional memory will greatly enhance program performance. We recommend 2 Gigs if you are running a 32-bit operating system and as much RAM as your budget can afford if you are running a 64-bit operating system.

A hard disk drive with 50 MB of free space.

Registering Your Forecast Pro Package

Please take a moment NOW to fill out and mail the registration card you received with Forecast Pro or visit www.forecastpro.com and register on-line. Registering your software entitles you to the following benefits:

- Free maintenance and support service for one year. This service provides program updates and unlimited technical support for the first year you license the program. After the first year, this service is available on a subscription basis. Technical support is provided via the telephone (during regular business hours Eastern Standard Time), via email, mail and fax.
- Automatic notification of upgrades, revisions and new products.
- Special pricing on upgrades, revisions and new products.

Future Development of Forecast Pro

BFS has made many changes to Forecast Pro to accommodate special needs and circumstances in the corporate environment. BFS will continue to improve Forecast Pro in the years to come. If you encounter problems or have

any suggestions for improvements, please contact the BFS product development staff

Chapter 2

Installing Forecast Pro

Automatic Installation

Forecast Pro can be installed as a native 32-bit or native 64-bit application. The installation program, FPSetup.exe, will automatically detect the operating system you are running and install the appropriate version. Running FPSetup is the recommended way to install Forecast Pro.

If you would prefer to use an MSI to directly install the 64-bit version or if you are running a 64-bit operating system but would like to install the 32-bit version, you will find instructions for doing so in the next section.

To install Forecast Pro on your hard disk using FPSetup, perform the following steps.

Run the FPSetup program as an Administrator (right click on FPSetup and select “Run as administrator”). The program is located in the root directory of the CD (e.g., D:\).

Follow the directions that appear on the screen. The installation program will prompt you for your serial number and CD-Key. These numbers are printed on the CD label and sleeve. If you are installing from a download, the numbers should have been included with the email that provided the download link.

Setup will then create the program directory, data directories, copy the appropriate files, create a launch icon on the program menu and place a launch icon on the desktop. It will also install SQLite drivers. The driver installation launches a separate installation dialog which includes a separate click-through license agreement.

If you have any difficulty installing Forecast Pro, please do not hesitate to contact BFS technical support (phone: 617 484-5050, email: support@forecastpro.com).

If you wish to remove Forecast Pro from your computer, use the Add or Remove Programs utility in the Windows Control Panel.

Manually Installing the 32-bit or 64-bit Version

As mentioned in the previous section, the FPSetup program automatically detects whether you are running a 32-bit or 64-bit operating system and installs the matching version.

Alternatively you can run either the 32-bit or 64 bit setup program directly.

The 64-bit installation uses an MSI file called ForecastProTRAC.msi which is located in the directory d:\x64\en-us\. You *must* run the MSI as an administrator.

The 32-bit installation uses an InstallShield installation program called Setup.exe which is located in the directory d:\Win32\. You *must* run the Setup program as an administrator.

Chapter 3

A Quick Overview of Forecasting

What is Statistical Forecasting?

Everybody forecasts, whether they know it or not. Businesses have to forecast future events in order to plan production, schedule their work force, or prepare even the simplest business plan.

Most business forecasting is still judgmental and intuitive. Sometimes this is appropriate. People must integrate information from a large variety of sources—qualitative and quantitative—and this is probably best done by using the extraordinary pattern recognition capabilities of the human brain. Unfortunately, many companies also use judgmental forecasting where they should not.

Not everyone understands the concept of forecasting. It tends to get mixed up with goal setting. If a company asks its salespeople to forecast sales for their territories, these "forecasts" often become the yardsticks by which they are judged.

The main advantage of statistical forecasting is that it separates the process of forecasting from that of goal setting and makes it systematic and objective. Statistical forecasting can help almost any business improve planning and performance. There is, in other words, value added for a business.

The future is uncertain and this uncertainty can be represented quantitatively. Statistical forecasting represents uncertainty via a *probability distribution.* A probability distribution associates each possible outcome with a likelihood of it occurring. Two kinds of information are needed to describe the distribution: the *point forecasts*, which is essentially the "best guess" estimate, and the *confidence limit,* which captures how much uncertainty there is around the point forecast. The upper and lower confidence limits represent reasonable bounds for the forecast. You can be reasonably confident that the actual outcome will fall within the confidence limits.

Forecast Pro depicts this information graphically as well as numerically. In the graph below, the red line represents the point forecast, while the blue lines represent the upper and lower confidence limits.

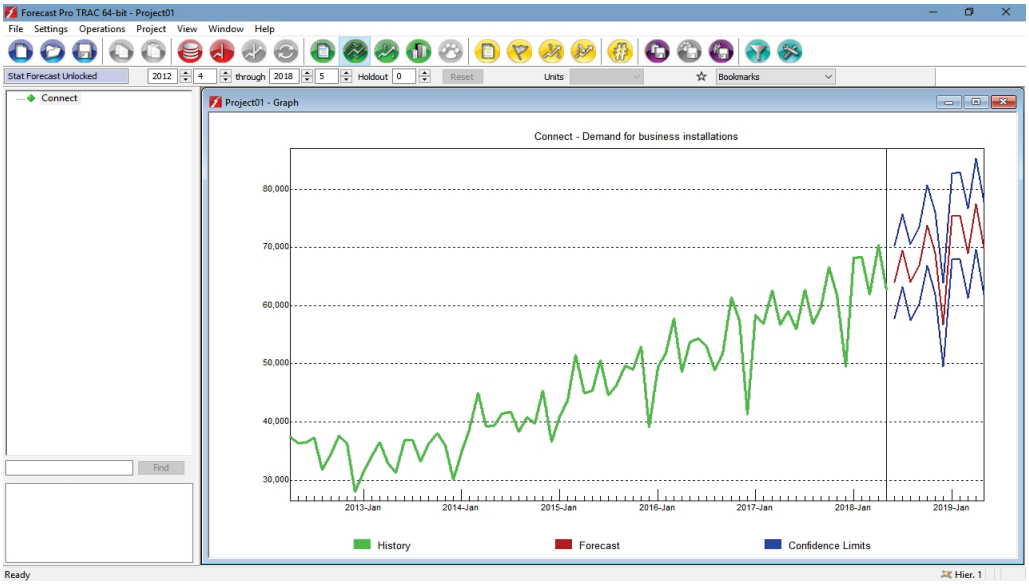

The upper confidence limit is often calibrated to the ninety-fifth percentile. This means that the actual value should fall at or below the upper confidence limit about 95% of the time. You can set the percentiles of both the upper and lower confidence limits. Sometimes, the upper confidence limit will be more useful for planning than the point forecast.

Let's illustrate this idea with an example. Suppose you were in charge of forecasting widget sales for your company. If you wanted to determine expected revenues for next month, you would be most interested in the point

forecast, since it is the mean value of the distribution. The point forecast gives you the minimum expected forecast error.

On the other hand, suppose you wanted to know how many widgets to produce. If you overproduce, warehousing costs will be excessive. But if you underproduce, you will probably lose sales. Since the cost of lost sales is usually greater than the cost of overstocking, you will be most interested in the upper confidence limit. The 95% upper confidence limit tells you how many widgets to produce to limit the chance of "stocking out" to less than 5%.

Forecasting Methodologies

A wide variety of statistical forecasting techniques are available, ranging from very simple to very sophisticated. All of them try to capture the statistical distribution that we have just discussed.

Forecast Pro offers the forecasting methodologies that have been proven to be the most appropriate for business forecasting—simple moving averages, discrete data models (Poisson or negative binomial), curve fitting, Croston's intermittent demand model, exponential smoothing, Box-Jenkins, Bass diffusion model, forecasting by analogy, dynamic regression, event models and multiple-level forecasting.

All of these methodologies forecast the future by fitting quantitative models to statistical patterns from the past. Therefore you must have historic records of your variables, preferably for several years. Forecast accuracy depends upon the degree to which statistical data patterns exist and their stability over time. The more regular the series, the more accurate the forecasts are.

Five of the methodologies are *univariate* techniques. They forecast the future entirely from statistical patterns in the past.

The *simple moving average* is widely used in business, mostly because it is so easy to implement. However, it is really only appropriate for very short or very irregular data sets, where statistical features like trend and seasonality cannot be meaningfully determined.

Discrete data models are used for data consisting of small whole numbers. These models are characteristically used to model a slow-moving item for

which most orders are for only one piece at a time. Forecasts are nontrended and nonseasonal.

Croston's intermittent demand model is not a widely known or used technique, but it can be extremely useful. It is usually used to model data in which a significant number of periods have zero demand but the non-zero orders may be substantial. This is characteristic of a slow-moving item which is ordered to restock a downstream inventory. Forecasts are nontrended and nonseasonal.

Exponential smoothing models are widely applicable. They are also widely used because of their simplicity, accuracy and ease of use. Their robustness makes them ideal even when the data are short and/or volatile. Exponential smoothing models estimate trend and seasonality and extrapolate them forward.

Box-Jenkins is a more elaborate statistical method than exponential smoothing. Box-Jenkins models estimate the historic correlations of the data and extrapolate them forward. It often outperforms exponential smoothing in cases when the data are fairly long and nonvolatile. However, it doesn't usually perform well when the data are unstable.

You can use Forecast Pro's *expert selection* to automatically choose the appropriate univariate forecasting technique for each item forecasted. Expert selection will determine which of the five univariate models is most appropriate for the data and select the modeling method that best fits the data. Alternatively, you can dictate that a specific method be used and, if desired, customize your models. Forecast Pro provides extensive diagnostics and statistical tests to help you make informed decisions.

Forecast Pro includes five additional forecasting techniques that are not considered in expert selection—event models, custom component models, forecast by analogy and the Bass diffusion model.

Event Models are extensions of exponential smoothing models that allow you to capture responses to promotions, business interruptions and other aperiodic events. These models allow you to assign each period into logical categories and incorporate an adjustment for each category. For example if you establish a category for promoted periods then your model would include an adjustment for promoted periods. If you ran three different types of promotions you could

establish three categories and have a different adjustment for each type of promotion.

Custom Component Models are also extensions of exponential smoothing models. The method generates statistical forecasts for the different components found in an exponential smoothing model (sales level, trend, seasonal pattern and events) and then allows you to customize any of the estimated components. The model is very useful in circumstances where not all of the components can be accurately estimated from the demand history. Examples include short data sets where the seasonal pattern cannot be reliably estimated and you wish to use a seasonal pattern from a similar product, forecasting the impact of future events that have not occurred historically, tempering the trend for longer-term forecasts, etc.

Forecast by Analogy is a new product forecasting technique that allows you to create a forecast that "looks like" a different product's demand pattern or a launch profile that you create.

Bass Diffusion Model is a new product forecasting technique designed to forecast the spread of a new product based on the adoption rates of two types of users—innovators who are driven by their desire to try new products and imitators who are primarily influenced by the behavior of their peers.

Dynamic regression produces a forecast based on the forecasted item's history (like univariate methods) and that of the explanatory variables (e.g., product promotion, advertising, demographic variables, or macroeconomic indicators). You must provide historic values for the variable to be forecast and the explanatory variables.

Forecast Pro can provide automated selection for the dynamic terms (e.g., lagged dependent variables and Cochrane-Orcutt terms), but the user needs to specify which explanatory variables to include in the model. Forecast Pro provides several test batteries and diagnostics for variable selection and gives you specific advice on how to improve the model. Building a dynamic regression model thus consists of deciding which variables to consider, and following the program's advice, step-by-step, to identify your final model.

Dynamic regression can outperform exponential smoothing and Box-Jenkins in cases where strong explanatory variables exist, and you have reasonably accurate forecasts for them. Unfortunately, this is not always the case, so the

forecasts may not be as accurate as those from univariate methods. Nevertheless, the method has considerable appeal, because it forces you to create a causal model for your data, and thus to improve your understanding.

If you are new to forecasting and these techniques seem a little intimidating, don't worry. We designed Forecast Pro to guide you completely through the forecasting process. Just follow the program's advice, and you will soon be generating accurate forecasts and adding value to your business.

Some Forecasting Tips

Forecast Pro uses your data history to forecast the future. Thus it is extremely important that your data be as accurate and complete as possible. Keep in mind the rule, “Garbage in, garbage out!”

You will also want to give some thought to what data you should forecast. If you want to forecast demand for your product, you should probably input and forecast incoming orders rather than shipments, which are subject to production delays, warehousing effects, labor scheduling, etc. Many corporations are making large investments to obtain data as close to true demand as possible.

The more data you can supply the program, the better. The program can work with as few as five data points, but the forecasts from very short series are simplistic and less accurate. Although collecting additional data may require some effort, it is usually worth it.

If your data are seasonal, it is particularly important that you have adequate data length. The automatic model selection algorithms in Forecast Pro will not consider seasonal models unless you have at least two years’ worth of data. This is because you need at least two samples for each month or quarter to distinguish seasonality from one-time irregular patterns. Ideally, you should use three or more years of data to build a seasonal model.

2

TUTORIAL: General Operations

Lesson 1

The Basics

Before you can use Forecast Pro effectively, you must learn some basic concepts. This tutorial will guide you through these fundamentals. The tutorial is divided into staged lessons. You should complete them all to get a sound introduction to Forecast Pro.

This first lesson provides an overview of Forecast Pro and guides you through a sample forecasting session.

Before you can use Forecast Pro, you must install it. If you have not already done so, follow the installation instructions discussed in Chapter 2.

The Main Menu

To start Forecast Pro double-click the Forecast Pro icon on your desktop or click the Start button, select All Programs and click the Forecast Pro icon. After the program is loaded, you will see a display like the one below.

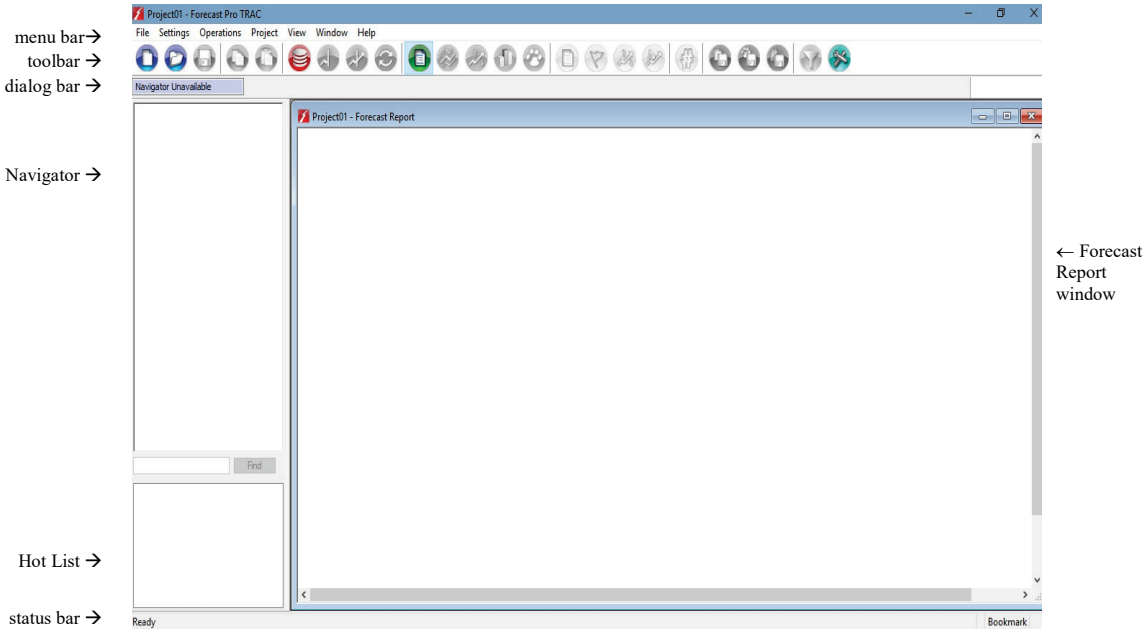

The main window, entitled *Forecast Pro* consists of a menu bar, a toolbar, a dialog bar, a Navigator (currently unavailable), a Hot List area, a Forecast Report window and a status bar.

The menu bar provides access to all procedures. The most commonly used procedures can also be accessed instantly via the icons on the toolbar. The dialog bar is currently blank. After data have been read in, the dialog bar will include the span of the data and the specified holdout sample. The Navigator and Hot List are currently unavailable. After data have been read in, the Navigator and Hot List can be used to select individual items that you wish to view.

The Forecast Report window is used to display the forecasting model, within-sample statistics, historic data and forecasts. The Forecast Report window is one of five context-sensitive (item-specific) *views*. A graph view will become available after data has been read in. An Override view will become available after forecasts have been generated. An Advanced Diagnostics view shows fitted errors, autocorrelation functions and a cumulative graph. A Tracking Report view will become available if there are archived forecasts (forecasts generated in prior periods). Context-sensitive views display information for the item currently selected on the Navigator. When more than one view is

active, the program will tile them. The context-sensitive views are toggled on or off using the green icons or the View menu.

Four additional noncontext-sensitive views are also available—the Item Report view, the Override Report view, the Outlier Report view and the Numeric Output view. These views display information for all items forecasted and are toggled on or off using the yellow icons or the View menu.

An Overview

Running Forecast Pro consists of executing the following four steps:

1. Prepare the database. This is performed outside of the program and entails selecting the type of data format to use and preparing the files. Creating your data files is discussed in the first chapter of the Reference section. For this lesson, you will use sample data files provided with the package.

2. Specify the default options. This is accomplished by selecting the appropriate entries from the Options Dialog box. All options set by the user will be used for the current forecast project. You can also save your selections as the default for future projects.

3. Define the historic data and other inputs. This is accomplished by specifying the input data in the Data Manager.

4. Read in the data and prepare the forecasts. This is accomplished by clicking the *Read & Forecast* button in the Data Manager. Forecast Pro will then read the data from disk into RAM and display the starting and ending dates for the forecast run on the dialog bar. The software will then prepare the forecasts. Once the forecasts have been created you may view them graphically, adjust them and save them to disk.

The remainder of this lesson will guide you through steps 2-4.

Setting the Defaults

To set the default options click the light blue Options icon ( ) which will open the Options dialog box shown below.

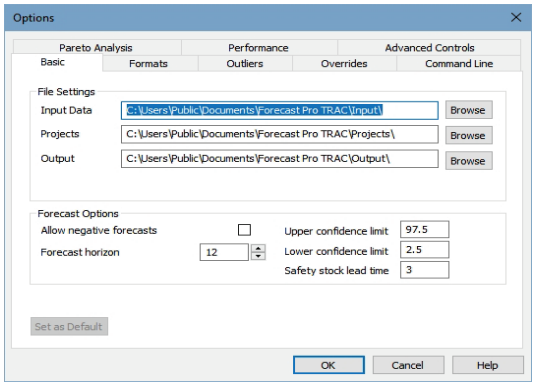

There are eight pages of options on the dialog box accessible via the tabs.

- **Basic**: controls the more commonly changed options. This tab is shown above.

- **Formats:** used to set details of the data file interface and control the precision in the on-screen displays. The data file interface defaults have been carefully chosen and should only need to be changed by users using comma-delimited data files, by users who do not follow the US conventions for the decimal separator, or by users who need to adjust the way in which zeros or blanks are handled by the program.

- **Outliers**: allows you to control the settings for outlier detection and correction. This will be covered in a tutorial.

- **Overrides**: allows you to customize the override facility. This will be covered in a tutorial.

- **Command Line:** allows you to control the settings for command-line operation. This will be covered in a tutorial.

- **Pareto Analysis:** allows you to dictate how to calculate the ABC classifications. This will be covered in a tutorial.

- **Performance:** allows you to exclude several forecasting methods from consideration in expert selection mode, use manual recalculation mode to control the timing of when the forecasts are reconciled and set policies for the loading and retention of archived forecasts. All of these settings can be used to improve Forecast Pro performance when working with very large projects.

- **Advanced Controls**: allows you to specify how integer forecasts should be rounded, how dates should be displayed and how often Forecast Pro checks for product updates.

Let's take a closer look at the Basic tab. Forecast Pro reads in the historical data from the Input Data directory. Project files are saved to the Projects directory. Forecast files and reports are written to the Output directory.

You can set the percentiles for the confidence limits, the lead time for safety stock calculations, the length of the forecasts (forecast horizon), and you can indicate if Forecast Pro should allow negative forecasts.

The data for this lesson were copied into your Forecast Pro input folder when you installed the program (normally Forecast Pro TRAC\Input in your Public Documents folder). Make sure that the Input Data directory is set correctly and that the other options match those shown above. Click the OK button to accept the defaults and return to the main menu.

If you are using a Forecaster license: proceed to the next section— Defining the Data.

If you are using a Collaborator license: the next section (Defining the Data) describes operations that are not supported. The first step in a Collaborator session is to open a forecast project that was created by someone using a Forecaster license. Select **File>Open** and open the forecast project *Tutorial - Basics.* Skip the next section and proceed to the Viewing the Forecasts section.

Defining the Data

The next step is to define the historic demand data you wish to forecast and any other inputs you would like to use.

In this example we will prepare sales forecasts for 123 Bakery. The data consist of monthly sales of Cakes and Muffins.

Click the red Data Manager icon ( ) to call up the Data Manager shown below.

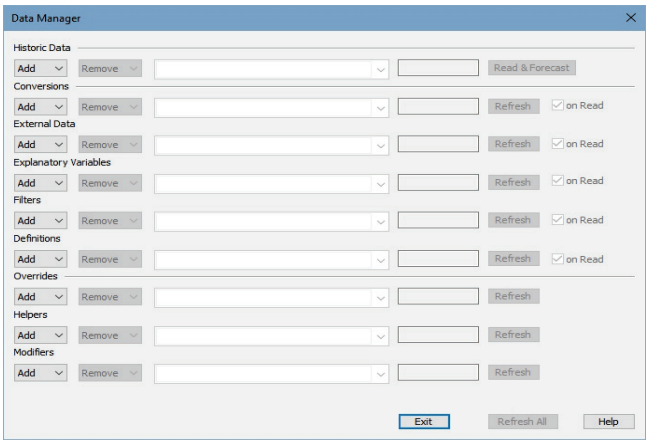

The Historic Data row on the Data Manager is used to specify the demand histories for the items you wish to forecast. The remaining rows are used to define other (optional) inputs. We will be exploring these options in future lessons. For all rows, the Add drop down is used to specify the file(s) or ODBC table(s) containing the relevant inputs.

Click the Add drop down on the Historic Data row, select Excel and select *123 Bakery – Historic Data.xlsx* to add it to the Historic Data row. See "Setting Up Your Historic Data" in the Reference section for complete details on file setup.

When you list a data file in the Historic Data row, the program will forecast *every series in the data file.*

Notice that the Historic Data row includes a button labeled *Read & Forecast*. As you probably guessed, clicking the button will read in the data and generate the forecasts. Before we perform this action, let's define some other inputs to use.

Click the Add drop down on the Conversions row, select Excel and select *123 Bakery – Conversions.xlsx* to add it to the Conversions row. The historic

demand in our example is in cases. *123 Bakery – Conversions.xlsx* contains conversion factors that allow us to view the forecasts in dollars, kilos and pallets.

Click the Add drop down on the Filters row, select Excel and select *123 Bakery – Filters.xls* to add it to the Filters row. *123 Bakery – Filters.xlsx* contains filter fields for the Sales Rep and Corporate Pareto Code associated with each end item. Filter fields are often useful when customizing reports and data displays.

Notice that with the exception of the Historic Data row, all of the rows have a Refresh button and five of the rows have an "on Read" checkbox. After the historic data has been read in and forecasted, the Refresh button can be used to load the specified, optional data files. If the "on Read" checkbox is selected for a file, that file will be automatically read in whenever you hit the Read and Forecast button. The on Read option is useful when the associated data will be changing each forecast period. We will discuss how each of these files is used by Forecast Pro in future lessons.

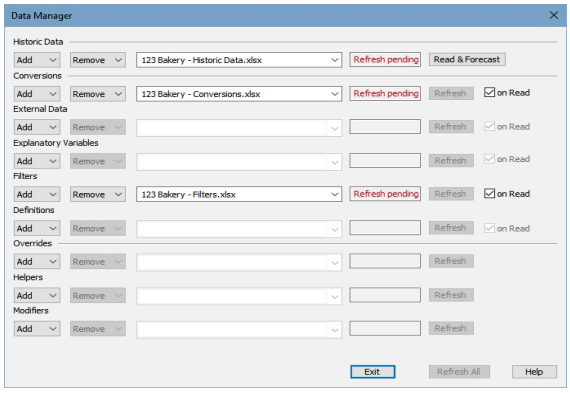

Your Data Manager should look like the one above.

Click the Read & Forecast button to read in the data and generate the forecasts. Click the Exit button to exit the Data Manager.

Viewing the Forecasts

Click the green View Graph icon ( ) to open the Graph view and display a graph for *Total*.

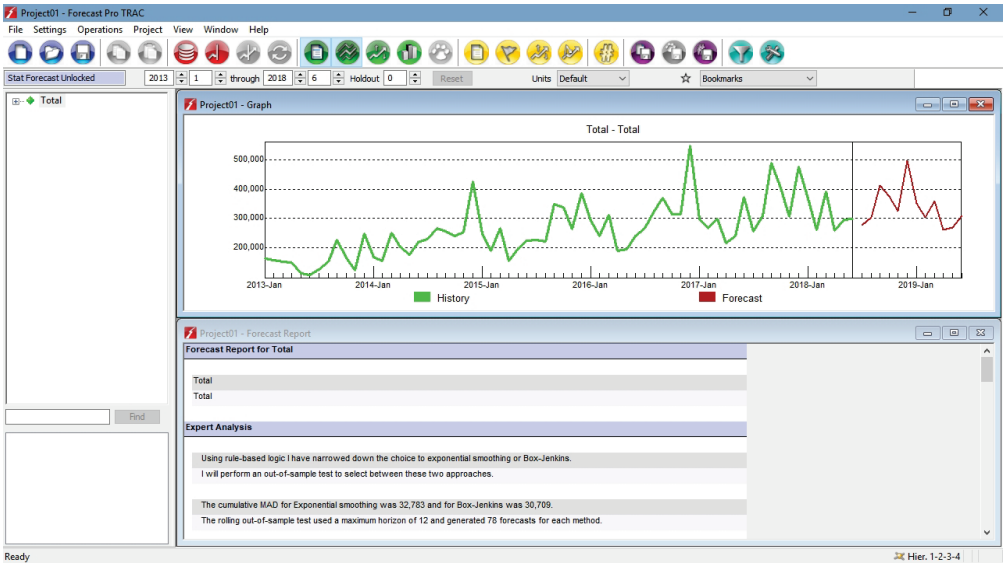

The green line in the Graph view is the historic demand and the red line is the forecast. Below, the Forecast Report view includes the forecasts and information about how they were generated—we will explore this shortly. Click on the green Forecast Report icon ( ) to close the Forecast Report.

Move the cursor to a point on the graph, hold the mouse very still and click the left mouse button. Notice that the display now includes a vertical red line and a box listing the corresponding date and numeric values of the graph variables.

Click January 2019 (the months are marked along the x-axis). Notice that the red line and box have moved to the new location.

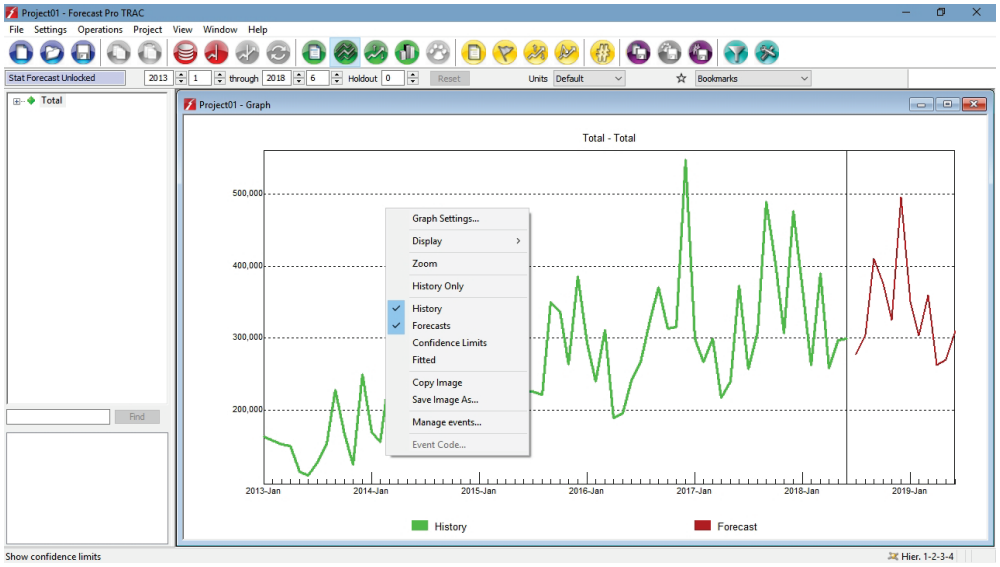

Click the right mouse button to invoke the Graph view's context menu. Notice that multiple graph components (History, Forecast, Confidence Limits and Fitted Values) can be toggled on or off. Select **History Only** to turn off the forecast. If other forecast-related components (e.g., confidence limits, fitted values, etc.) were displayed, selecting History Only would turn off all forecast-related components. Select **History Only** a second time to turn the forecast back on.

The Graph view can be customized using **Graph Settings** on the context menu to adjust the data display (e.g. colors, graph type, titles), edit the data shown on the graph, or display multiple items. We will explore Forecast Pro's graphing functionality in more detail in another lesson.

Double-click *Total* on the Navigator to expand the Navigator's tree. You'll see that the total sales breakdown into Cakes and Muffins. Select *Total>Cakes* on the Navigator. Notice that the Graph view is immediately updated to match our current selection.

If you expand *Total>Cakes* on the Navigator, you'll discover a customer-level breakdown and finally a SKU level breakdown. Experiment with the Navigator until you are comfortable with its operation. In the next lesson, we'll show you some shortcuts for navigating complex hierarchies.

After you have explored the data, click on the green Forecast Report icon ( ) to re-open that report and arrange your display so that it matches the one shown below.

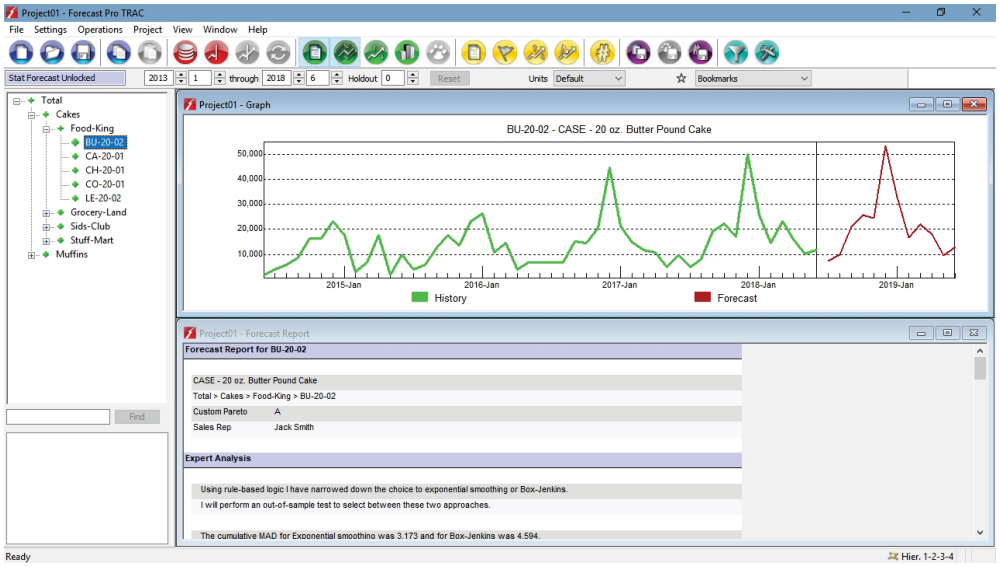

Notice that after the data was read in, the dialog bar was updated to include the time span of the historic data and the holdout sample.

The dialog bar displays the earliest and latest period found for the entire historic data set. The ending date is very important. If an individual item has data on the ending date, by default the item is active and will be forecasted. If an item's history ends prior to the ending date, the item is inactive and will not be forecasted. If you do not adjust the starting or ending dates Forecast Pro will prepare forecasts using all available data for each active series. You can use the Formats tab on the Options dialog box to adjust how blanks and zeros are handled and when items should be considered inactive. For this lesson, we will use the default settings.

For this example, we used expert selection to create the forecasts. Expert selection is the default method. Expert selection can be thought of as an automatic pilot. It analyzes your data statistically to determine the best forecasting technique to use and then proceeds to build the forecasts.

Expert selection begins by running a series of statistical tests on the data. The results of these tests are run through a rule-based logic system. At times this

logic system will directly select an appropriate forecasting method. At other times the rule-based logic will narrow down the potential forecasting methods that could be used without producing a clear winner. In these instances, expert selection will use an out-of-sample test to select between the competing models.

The Expert Analysis section of the Forecast Report documents how Forecast Pro makes its decision. In our example, the rule-based logic for *Total>Cakes>Food-King >BU-20-02* was inconclusive, but narrowed down the choices to either an exponential smoothing model or a Box-Jenkins model. An out-of-sample test was then used to choose between these two candidates.

Notice that the expert system acts in much the same way as an expert statistical analyst. It performs a series of statistical tests, interprets the results, performs more tests if necessary and finally makes recommendations based upon the results.

Expert selection works well for most applications. However, Forecast Pro also allows you to customize your forecasting approach by adding *modifiers* to the Navigator. A modifier on a particular item gives Forecast Pro instructions on how to create the forecasts for that item. We will explore the use of modifiers in future lessons.

Viewing the Forecast Report

Let's take a closer look at Forecast Report view. Turn off the Graph view by clicking its green icon ( ).

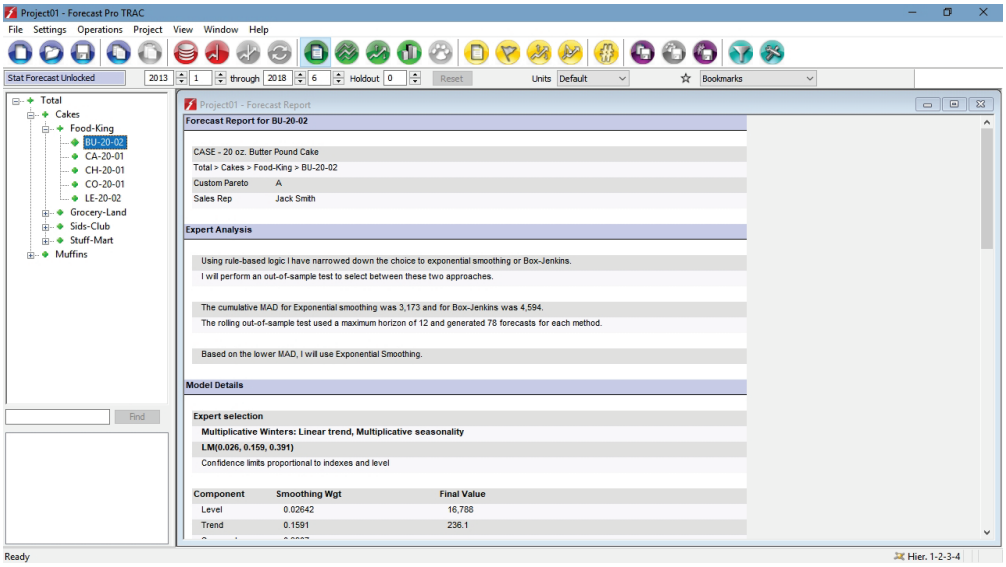

You are currently looking at a standard format Forecast Report. The Forecast Report can be customized and can include up to fifteen sections. They are:

Expert Analysis documents the logic behind the expert selection's decision. This section will be omitted if a user-specified forecasting model is used.

Model Details documents the specific forecasting model that was used to generate the forecasts.

Within-Sample Statistics provides a set of standardized model statistics that can be used to diagnose the current model as well as to compare and contrast alternative forecasting approaches.

Variable Specification Test Battery (not currently active) provides test statistics for explanatory variables to add to a regression model. This is only available when dynamic regression is used.

Dynamics Test Battery (not currently active) provides test statistics for dynamic terms to add to a regression model. This is only available when dynamic regression is used.

Correlation Matrix (not currently active) shows the correlation coefficient between each set of explanatory variables in a regression model. This is only available when Dynamic Regression is used.

Covariance Matrix (not currently active) shows the covariance between each set of explanatory variables in a regression model. This is only available when Dynamic Regression is used.

Historic Data (not currently active) lists the history and fitted values along with some aggregated totals and summary statistics.

Forecast Data lists the forecasts and confidence limits along with some aggregated totals and summary statistics.

Converted Forecasts (not currently active) lists the forecasts in the converted units defined by the conversion factors that were imported into the project.

External Rows (not currently active) lists the values of any external data rows that were imported into the project.

Safety stocks lists the safety stocks for different lead times along with the demand during lead time and reorder point.

Out-of-Sample Rolling Evaluation and *Out-of-Sample Static Evaluation* (not currently active) list the out-of-sample statistics that are generated when a holdout analysis is performed.

Outliers (not currently active) lists any detected and/or corrected outliers.

Overrides displays any overrides and/or comments that were entered.

Right click on the Forecast Report view to invoke its context menu. Select **Forecast Report Design**. Notice that this dialog box allows you to select a standardized report style or create a custom one.

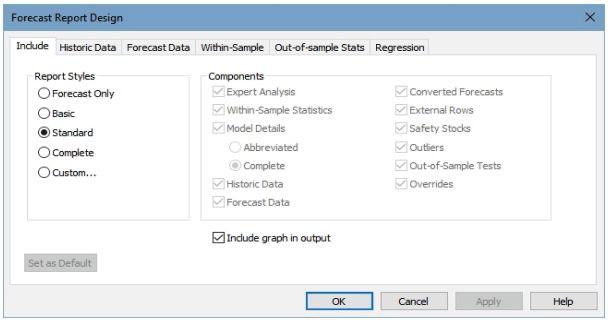

Select Forecast Only and click OK. Notice that this report is much simpler that the Standard one. Return to the Forecast Report Design dialog box and experiment with the settings. After you are comfortable with their operation edit the settings to display the Standard report style and click OK.

Overriding the Forecasts

There will undoubtedly be times when you want to judgmentally override the statistical forecasts. For instance, you may know of a large incoming order and need to change a number or two. Forecast Pro allows you to make these changes quickly and easily.

Select *Total>Cakes>Food-King* on the Navigator. Click the green Override Forecasts icon ( ) or select **View>Overrides** from the menu. Close the Forecast Report view by clicking the Forecast Report icon ( ) and open the Graph view by clicking on the Graph icon ( ). Your display should look like the one below.

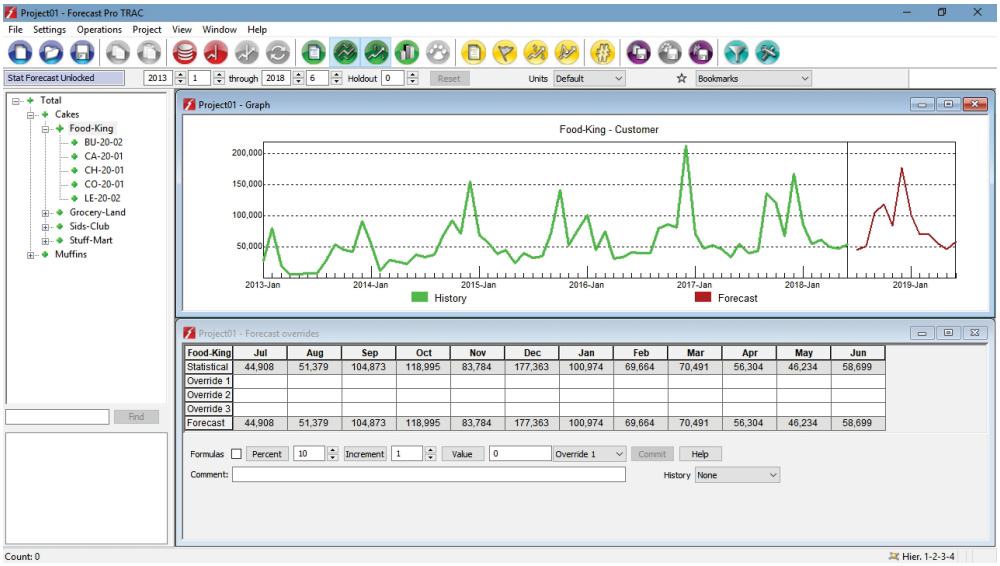

The Override Forecasts view provides several different ways to override the statistical forecasts. You can enter values in individual override cells, you can select a range of cells and use the Percent, Increment or Value buttons to write the desired overrides to the target override row, or you can use the mouse to

drag points on the graph. If a time period has overrides in both Override 1 and Override 2, Override 2 will take precedence.

Highlight the Statistical values for August and September of 2018. Set the Percent box to 12% and click the Percent button. Notice that the proposed overrides are written to the target row indicated to the immediate left of the Commit button—in this example, Override 1. Notice also that the values displayed in the Forecast row have not changed—this is because we have not yet committed to the overrides. Uncommitted overrides are displayed in red.

Click the Commit button to accept the overrides. Notice that the formerly uncommitted values in the override row change to black, the Forecast row is updated and the graph is updated.

Notice also that the icon for *Total>Cakes>Food-King* on the Navigator now appears in red and the icons for *Total>Cakes* and *Total* in yellow. A red Navigator icon signifies that an override has been made to the item. If a group icon appears in yellow it signifies that an override has been made to a variable lower down that branch of the hierarchy.

When you make an adjustment to an item in a hierarchy, Forecast Pro will automatically reconcile the entire hierarchy to reflect the change. Thus when we made our override to *Total>Cakes>Food-King* this impacted the groups above our selection and all groups and items below our selection. The Reference section of this manual describes how the reconciliation is performed.

Important: *If you restrict your overrides to a single level of the hierarchy, the reconciliation process is simple and intuitive. We strongly urge you to limit your overrides to a single level if at all possible.*

We have illustrated the basics of the override facility. There is a considerable amount of additional override functionality, including the ability to:

- Associate comments with override cells
- Add and remove rows
- Name override rows

- Make overrides on the graph using your mouse
- Enter overrides as formulas
- Make incremental adjustments to the Statistical Forecast row rather than overriding the values.

You also have the ability to customize the override grid to include rows to display history, alternative units of measure, imported information (e.g., externally generated forecasts, current orders, etc.) and calculations. Some of this advanced functionality is explored in future lessons, and all of the functionality is described in the Reference section of this manual.

Turn off the Override view by clicking its green icon ( ).

Saving the Forecasts

Forecast Pro provides two different ways to output your forecasts. Each is described below.

The *Numeric Output file*(s) is used to save the forecasts and other information in a concise format. If you plan to import the forecasts into a database or a planning system, this is the file you should use. You have a great deal of control over the content and format of this file and can save it in text, spreadsheet, ODBC or XML formats.

The *Forecast Report* is an Excel file containing a separate forecast worksheet for each item forecasted. Each report worksheet will contain the currently displayed Graph view for the item and the currently displayed Forecast Report view for the item. Thus, in addition to allowing you to view information on screen, the Graph view and the Forecast Report view also are used to design the Excel-based forecast report.

Arrange your display so that it includes both the Graph view and the Forecast Report view.

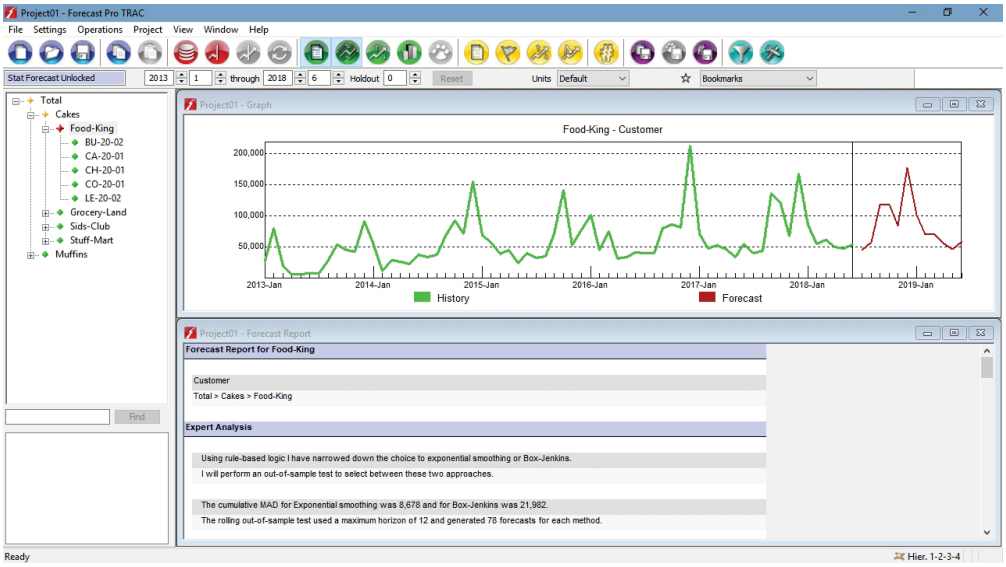

The purple buttons on the tool bar are used for saving outputs. Click on the Save Formatted Forecast Report icon (). A dialog box will appear allowing you to name the Excel file. Name the file *Forecast Report* and save it.

Examine *Forecast Report.xlsx* in Excel and verify that it contains a worksheet for each item in the Navigator and that the contents and format matches the currently displayed Graph and Forecast Report views.

Close the Graph window and the Forecast Report window. Open the Numeric Output Preview window by clicking the yellow Preview Numeric Output icon ().

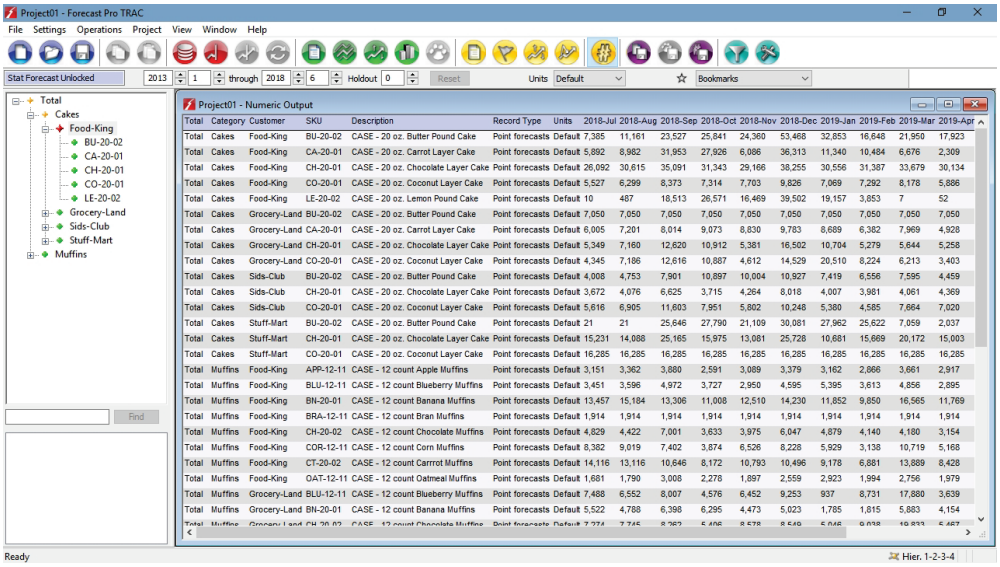

The Numeric Output Preview window displays the contents and format of the currently specified Numeric Output file. Click **Settings>Numeric Output Design**.

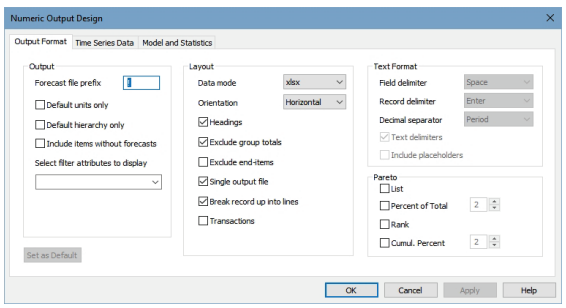

This dialog box allows you to configure the Numeric Output. The first tab allows you to specify whether you wish the output to use the currently displayed units and hierarchy or always output the default units and hierarchy. It also allows you to specify the file type and control the layout. The Time Series Data tab and the Model and Statistics tabs allow you to specify the contents.

Explore the options on this dialog box. If you are unsure of the function of any selection click the Help button for a description.

The settings you select will be used for your current project. If you click the Set as Default button on the displayed tab, the current settings will also be saved as the default to use for all new projects. Exit the dialog box.

Click the purple Save Numeric Output icon ( ). A dialog box will appear allowing you to name the file. Name the file *Test Numeric Output* and save it. Examine the file and verify that its contents and format matches those specified in the Numeric Output Design dialog box.

Saving the Project

Our final step will be to save our current forecasting session as a *project*. Saving a forecast project allows you to return later and pick up where you left off or to share the session with others.

Select **File>Save As** and save the project using the name *test project*.

Be aware that a forecast project saves the data, forecasts and overrides that are currently in memory. If you update your input data files after saving a project, when you reopen the project you'll have the option of either opening the project as it existed when you saved it or reading in the updated data and revising the forecasts.

Exit the program.

This concludes *The Basics* lesson.

Lesson 2

Advanced Navigation

In this lesson we will be looking at three tools designed to allow you to use Forecast Pro more efficiently. First we will explore using Hot Lists which allow you to define and work with subsets of your data. Then we'll examine some sorting and filtering capabilities that allow you to customize multi-item report views. Finally we will see how defining bookmarks allow you to save and return to customized reports and views.

If you are using a Forecaster license: proceed to the next section—Generating the Forecasts.

If you are using a Collaborator license: the next section describes operations that are not supported. The first step in a Collaborator session is to open a forecast project that was created by someone using a Forecaster license. Select **File>Open** and open the forecast project *"Tutorial - Advanced Navigation"* then skip the next section and proceed to the Working with a Hot List section.

Generating the Forecasts

Start Forecast Pro.

Click on the red Data Manager icon ( ) to call up the Data Manager.

Click the Add drop down on the Historic Data row, select Excel and select *123 Bakery – Historic Data.xlsx* to add it to the Historic Data row. Your Data Manager should look like the one below.

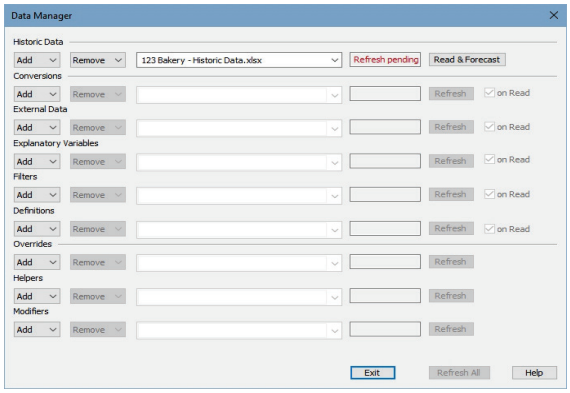

Click on the Read & Forecast button and then on the Exit button to close the Data Manager.

Turn on the Graph view by clicking its green icon ( ), and turn off the Forecast Report view by clicking the Forecast Report icon ( ).

Right click on the Navigator to invoke the Navigator's context menu. Select **Expand All** to fully expand the Navigator.

Drag the top of the Hot List section (right below the Navigator) up a bit to increase the window size. Your display should now look like the one below.

Hot List →

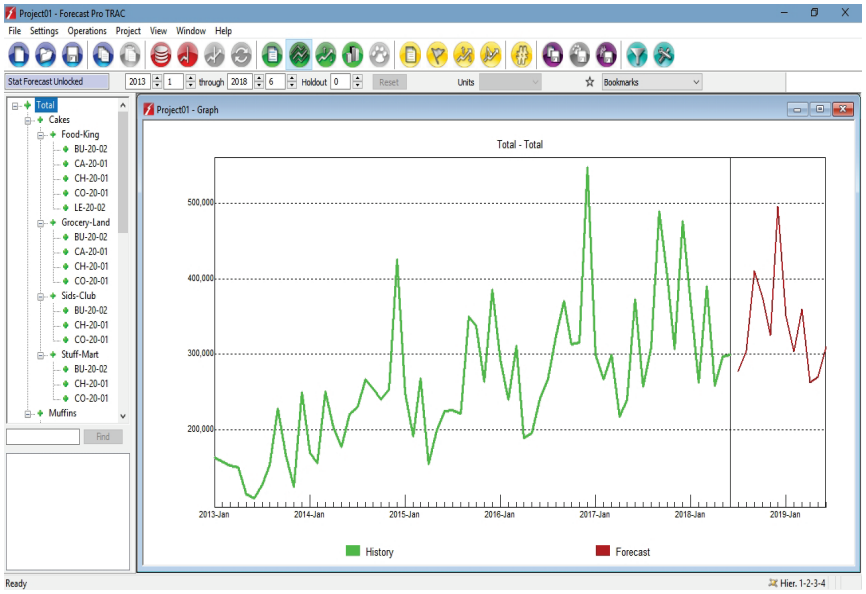

Working With a Hot List

The *Hot List* is the (currently empty) box, located in the bottom left-hand corner of the screen, below the Navigator. Placing items on the Hot List allows you to efficiently navigate and work with a subset of your items.

Drag *Total>Cakes>Food-King* from the Navigator into the Hot List area. Your display should now look like the one below.

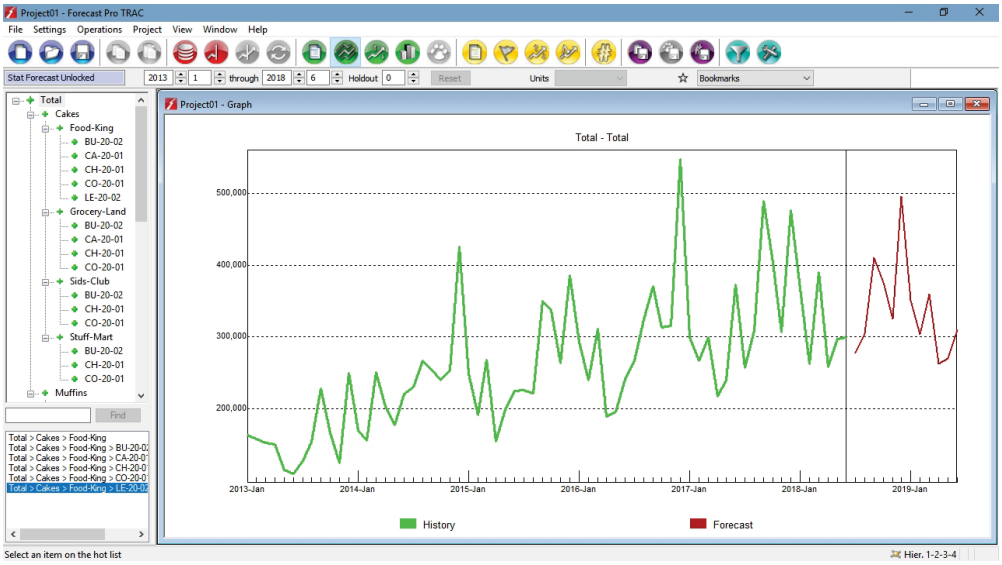

Click on the different items in the Hot List and notice that as you do so, the Navigator selection and views are immediately updated to match your selection.

Right click in the Hot List area to call up the Hot List's context menu. Notice that there are a full range of options to control how Hot List items are displayed, to import and export Hot Lists, and to save outputs for only Hot List items. Select **Clear** to clear the Hot List.

Now select *Total>Cakes>Food-King* on the Navigator and right click to bring up the Navigator's context menu and select **Add to Hot List**. Notice that there is a full range of options to add items to the Hot List. *Children, Parents* and *Siblings* refer to one level down on the current branch, one level up on the current branch and the same level on the current branch, respectively. In the "Analysis" menu selection, *Overrides* refer to items with direct overrides and *Affected* refers to items with indirect overrides (i.e., items where a direct override elsewhere in the hierarchy changed the items' forecasts). The other selections available in the "Analysis" menu selection are described in the Reference section in this manual.

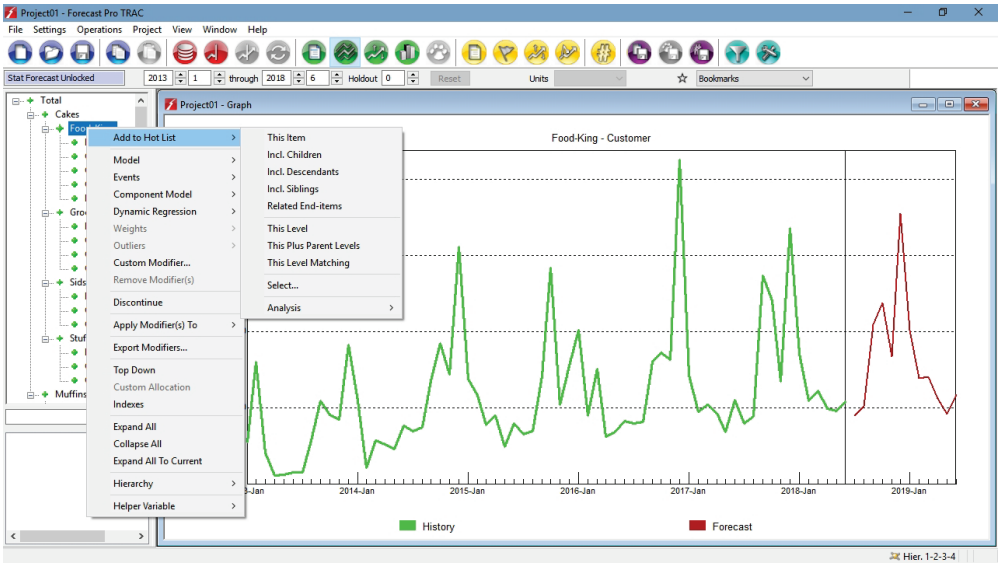

Experiment with adding and removing items from the Hot List until you are comfortable with its operation.

Clear the Hot List by selecting **Clear** on Hot List context menu. Select *Total>Cakes>Food-King* on the Navigator, invoke the Navigator's context menu and select **Add to Hot List>This Level**. Your screen should now match the display below.

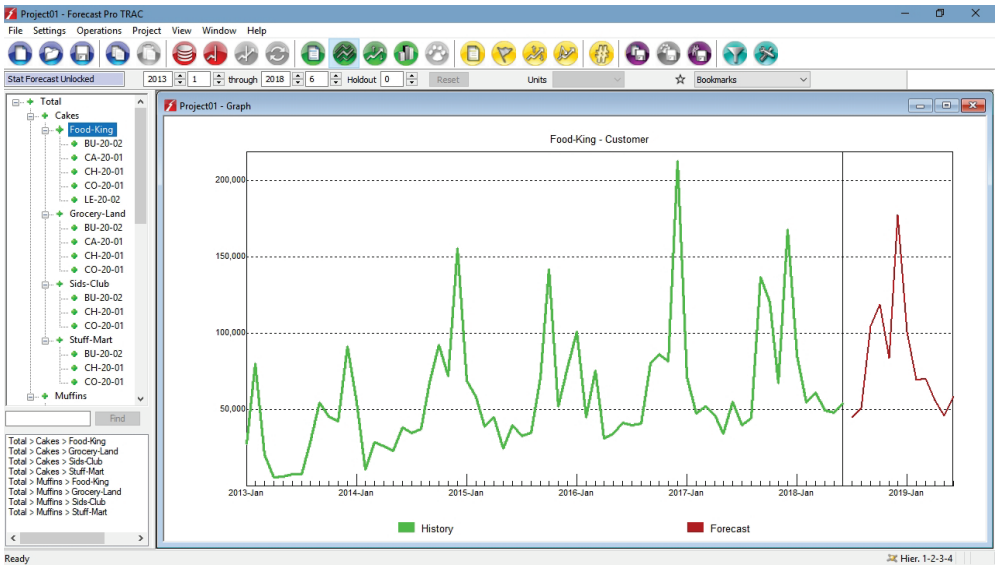

Filtering Multi-item Reports

The yellow icons on the toolbar all represent non-context specific (i.e., global) views that may include all forecasted items or a static subset of items (e.g., a hot list). The contents of these views do not depend on what item is selected in the Navigator. Most will be covered in more detail in subsequent tutorial lessons and all are documented in the Reference section of this manual.

- The Item Report displays model details, model statistics, forecasts and other time series data in a concise display.
- The Exception Reports compare the forecasts (or history) to data that you specify (e.g. forecast vs. actual, forecast vs. history, etc.) allowing you to monitor performance and identify potential issues.
- The Override Report allows you to display all overrides that are present in the current project.
- The Outlier Report lists information pertaining to all detected outliers in the current project.
- The Numeric Output view is very similar to the item report, but it is primarily intended to be used for defining output for use with other applications.

Click on the green Graph icon ( ) to close the Graph view, and click on the yellow Item Report icon ( ) to open the Item Report. Your screen should look like the one below.

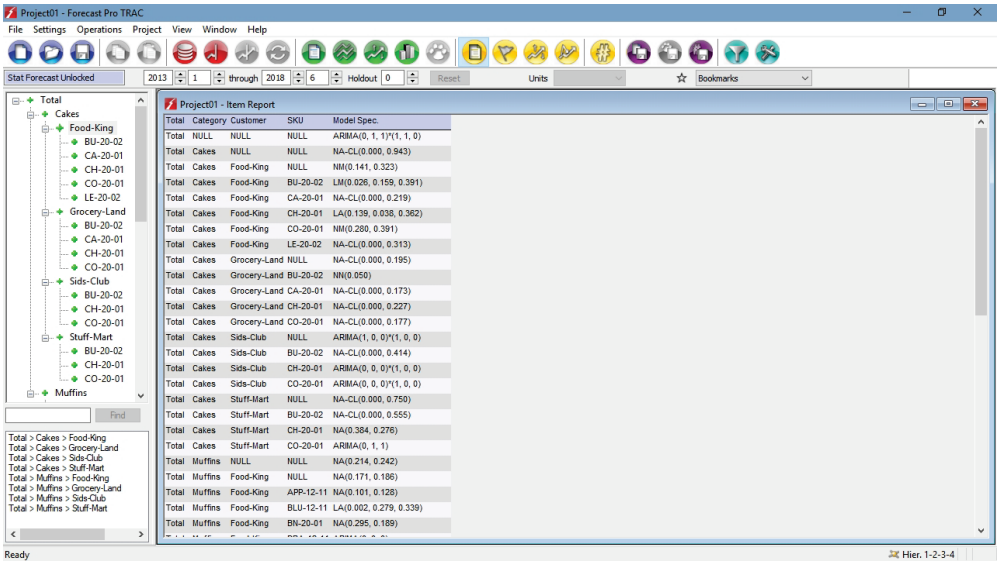

Notice that each row represents an item on the Navigator. By default, the report shows the Model Specification for all items and levels in the hierarchy. Right click on the Item Report to invoke its' context menu and select **Item Report Design**.

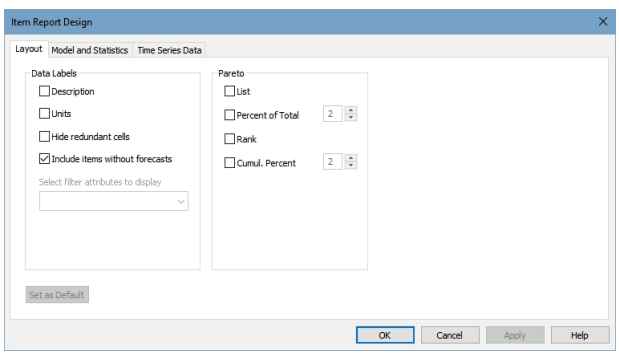

The options in the Item Report Design dialog box allow you to add and remove columns from the Item Report view. The tabs organize the available options into logical groupings.

The Layout tab allows you to add or remove descriptive fields, filters (if a filters file was read in the Data Manger) and Pareto statistics. It also gives you control over how attribute fields are displayed and whether or not to include rows for items that were not forecasted.

Select the “Model and Statistics” tab. This tab allows you to add or remove fields pertaining to model descriptions, model statistics and safety stocks. The spinners to the right of each item are used to specify numeric precision. Select Historic Mean and use the spinner to decrease the numeric precision to 0. Check MAPE, leaving the numeric decision at the default of 2 and click Apply. Notice that the report has been updated to include the additional columns.

Select the Time Series Data tab. This tab allows you to add or remove various time series such as the forecasts, external data, etc. Select *Forecasts*, select *Upper limits* and click the Apply button to update the Item Report view.

Notice that there are now two rows for each listed item. The first displaying the forecasts and the second displaying the upper confidence limits. For this example we don’t want to include the upper confidence limits, so remove the checkmark for *Upper limits* and click OK to exit the dialog box and update the Item Report view. Your screen should now match the one below.

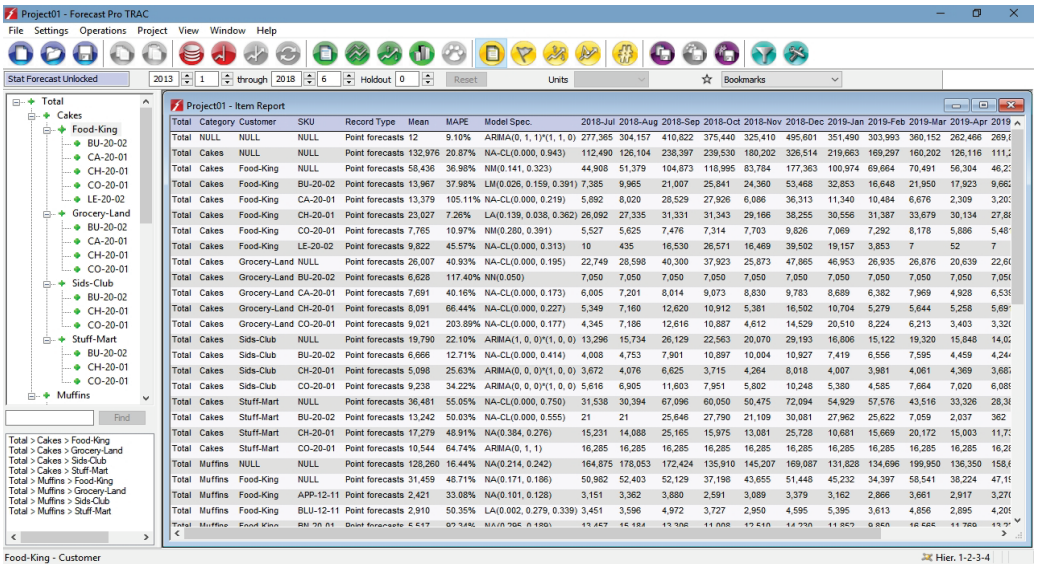

Now let’s explore changing the items included in the Item Report. Right click on the Item Report to invoke its’ context menu. Notice that **All Items** is currently selected. Below All Items are two alternative item displays—**Hot List** and **Selected Item**. Select **Hot List**. Notice that the display is updated to show only items on the current Hot List. Now select **Selected Item**. The report is updated to display only the item that is currently selected in the

Navigator (*Total>Cakes>Food-King*). Turn All Items back on. This functionality is available in all Non-context specific (Global) reports (anything accessible with a yellow icon on the toolbar).

Filtering is used to customize the items shown on the report. Click the light blue Filter icon ( ) on the toolbar to activate the filtering and sorting mode. Note that a drop down icon has been added to each column in the Item Report.

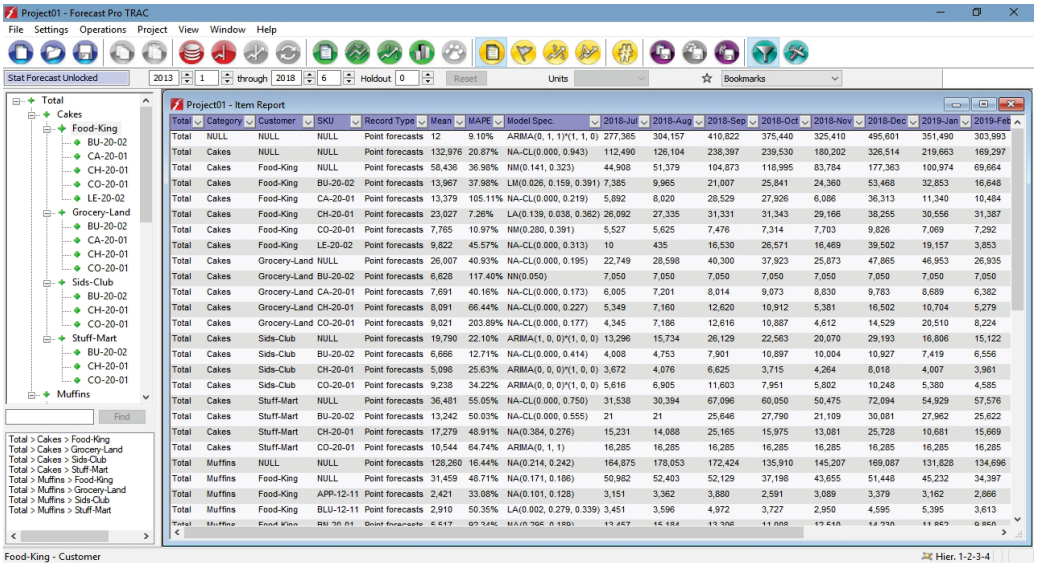

Click on the drop down for the MAPE column. Select “Sort Largest to Smallest” and press OK.

Suppose we want to now filter out low volume items. Click on the Mean filter. Select the “Numeric” drop down and select **Numeric.** The Numeric Analysis dialog box will appear. In the top drop down in the Numeric Analysis dialog box select “is greater than or equal to” and enter 10000 in the textbox to the right.

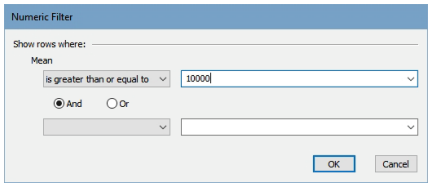

Click OK on the Numeric Analysis dialog box and then click OK on the Mean dialog box. Your screen should now show only items with a mean greater than 10,000, sorted by MAPE.

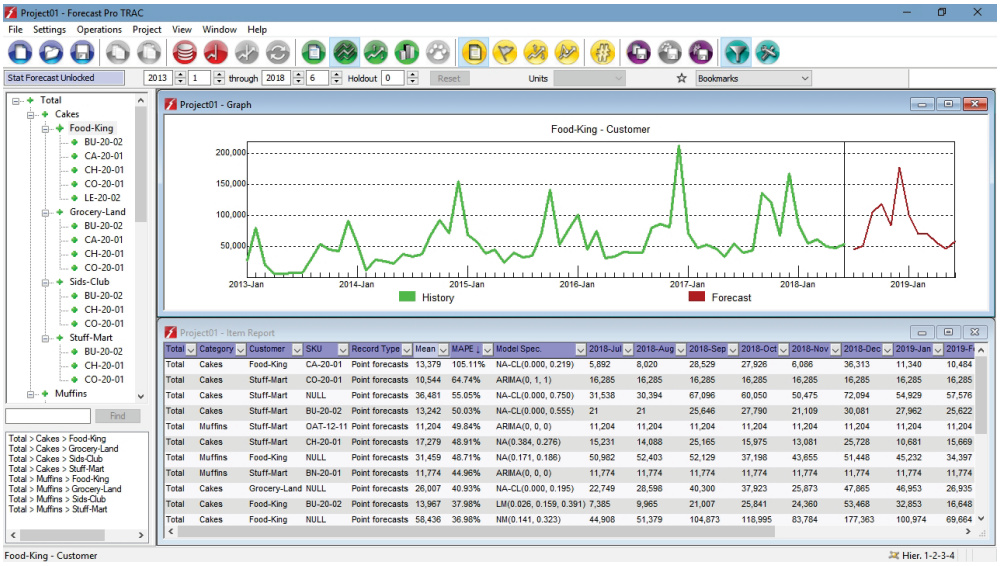

Click the green Graph icon ( ) to open the Graph view. Point to the first item on the item report and double-click. Point to a different item on the report and double-click. Notice that when you double-click an item on the item report, the Navigator jumps directly to that item. Thus, you can use the item report as the equivalent of a Hot List to navigate through the listed items. All of the global report views (yellow-icon views) support this kind of navigation.

Working with Bookmarks

In this section we will explore using bookmarks to save and return to customized reports and views.

Our current display includes a filtered item report and a graph. To add a bookmark for this display, click on the Bookmark star (☆) and select *Add* on the context menu. The Add Bookmark dialog box will appear.

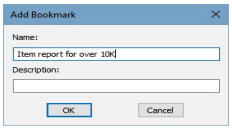

Name the Bookmark “Item report for over 10K” and click OK.

Click on the yellow Item Report icon ( ) to close the Item Report view. Click the green Override Forecasts icon ( ) to open the Override view. This is a useful display for reviewing the forecasts, and it is likely that we will want to return to it in the future—so let’s save a bookmark for it.

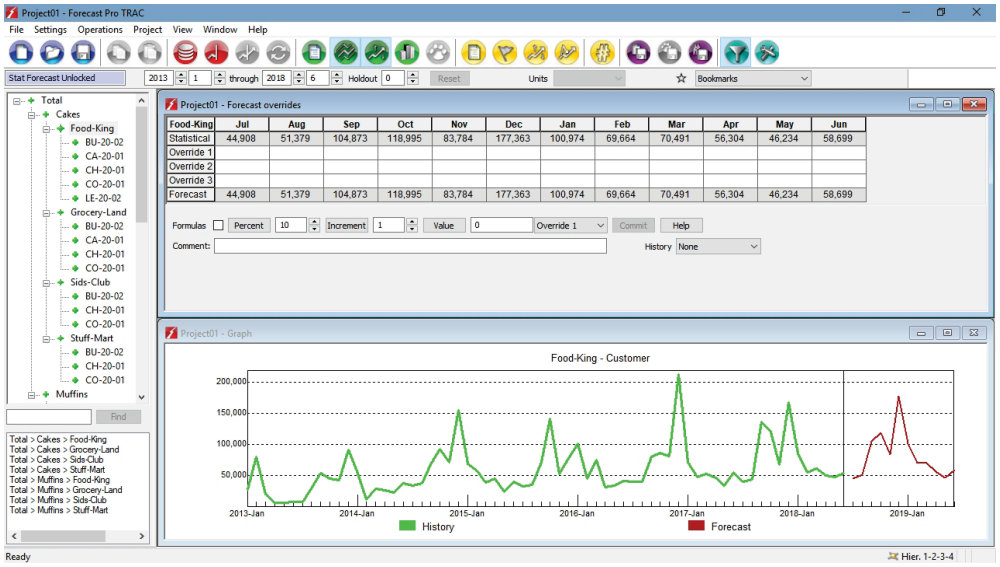

To add a bookmark for this view, click on the Bookmark star (☆) and select *Add* on the context menu. Name the bookmark “Forecast View” and click OK.

Notice that the Bookmark star is now blue and the drop down menu to the right reads “Forecast View.” The blue star indicates that a saved bookmark is active, while the drop down box shows which bookmark is active. Right click on the graph to bring up the context menu and select Confidence Limits. Notice that the star has turned from blue to clear. The clear star indicates that the current display is no longer consistent with a saved bookmark.

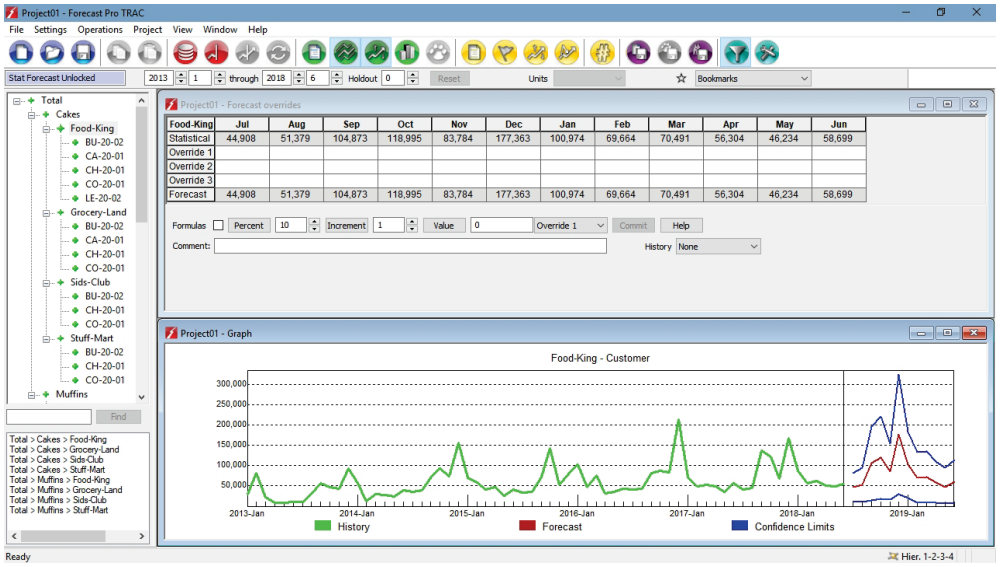

Let's update the "Forecast View" bookmark to reflect our changes. Click on the bookmark star icon and select *Refresh* on the Bookmark menu. Select "Forecast View" in the Refresh Bookmark dialog box and click OK. The star turns to blue and the Bookmark drop down displays "Forecast View" to reflect that the current display is consistent with the revised "Forecast View" bookmark.

Open the bookmark drop down list and select *Item report for over 10K*. Notice that this action restores the Item report for over 10K display and marks it as the active bookmark.

Open the bookmark drop down list again and select *Forecast View*. Notice that this action restores the Forecast View display and marks it as the active bookmark.

Bookmarks have functionality well beyond what we have explored here. You can customize the settings and display for each view using *Manage* on the Bookmark menu, delete bookmarks using the *Remove* on the Bookmark menu or import bookmarks from a different project using *Import* on the Bookmark menu. This functionality is documented in the Reference section of this manual.

Exit the program. This concludes the *Advanced Navigation* lesson.

Lesson 3

Customizing the Forecast Override Worksheet

In the first lesson we used the Override Forecasts view to add judgmental overrides to the statistical forecasts. In this lesson we will explore this view more thoroughly and illustrate how to create a customized worksheet that can include external information (such as alternative forecasts, current orders, etc.) and calculated rows.

Start Forecast Pro.

If you are using a Forecaster license: proceed to the next section—Generating the Forecasts.

If you are using a Collaborator license: the next section describes operations that are not supported. The first step in a Collaborator session is to open a forecast project that was created by someone using a Forecaster license. Select **File>Open** and open the forecast project *"Tutorial – Customizing Overrides"* then skip the next section and proceed to the Formatting the Display section.

Generating the Forecasts

We'll begin this lesson by generating a forecast for the cakes and muffins data that we explored in *The Basics*.

Click the red Data Manager icon ( ) to call up the Data Manager.

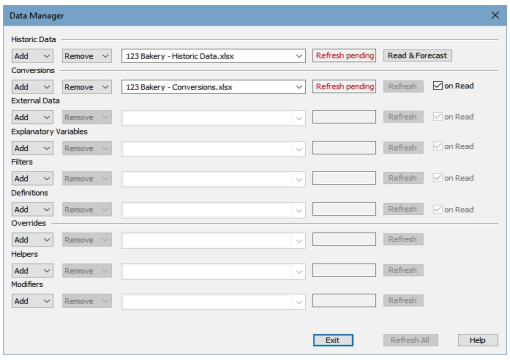

Click the Add drop down on the Historic Data row, select Excel and select *123 Bakery – Historic Data.xlsx* to add it to the Historic Data row.

Click the Add drop down on the Conversions row, select Excel and select *123 Bakery – Conversions.xlsx* to add it to the Conversions row. Your Data Manager should now match the one above.

Click the Read & Forecast button to read in the historic demand data and generate the statistical forecasts. Since the Conversion row's *on Read* option is active, notice that they conversions were read in as well.

Click the Exit button to exit the Data Manager.

Formatting the Display

Display the Graph view by clicking its green icon ( ). Display the Override Forecasts view by clicking its green icon ( ). Turn off the Forecast Report view by clicking its green icon ( ).

Right click on the Navigator to invoke the Navigator's context menu. Select **Expand All** to fully expand the Navigator. Select *Total>Cakes>Food-King>Bu-20-02* on the Navigator. Your display should now look like the one below.

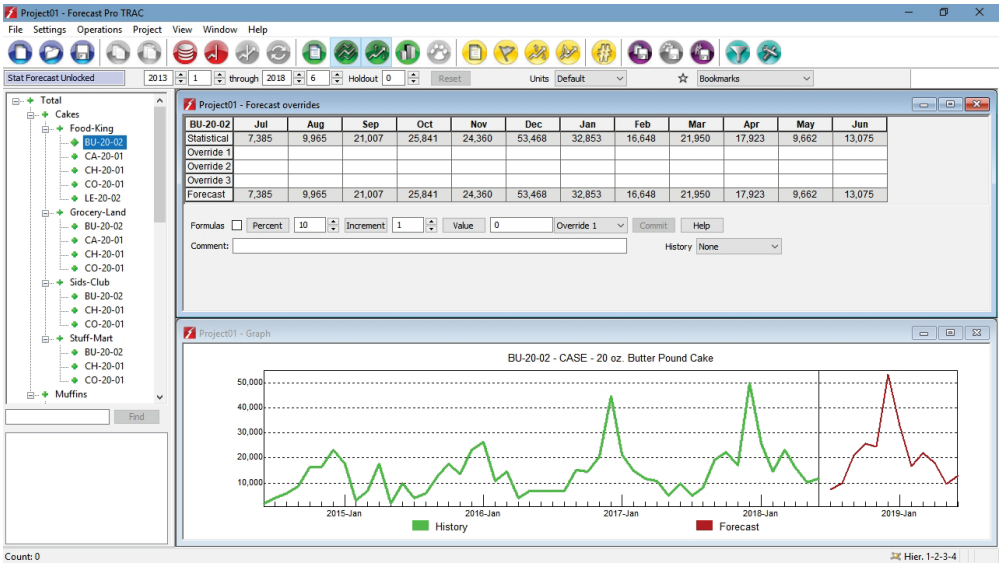

The default Forecast Override view display contains a row for the statistical forecast, three override rows and a row for the current forecast. Forecast Pro allows you to include up to 10 override rows. Dedicating different override rows to different types of adjustments and labeling them appropriately is an excellent practice.

Right click the row labeled Override 1 to invoke the context menu. Notice that this menu allows you to add or remove override rows, edit the row labels, set the precision of the numbers displayed and control whether a thousands separator should be used. Select **Edit Row Label** and rename Override 1 “Demand Planning.” Right click the row labeled Override 2 and use **Edit Row Label** to rename it “Management.” Right click the row labeled Override 3 and use **Remove Row** to remove it from the display. Your screen should now look like the one below.

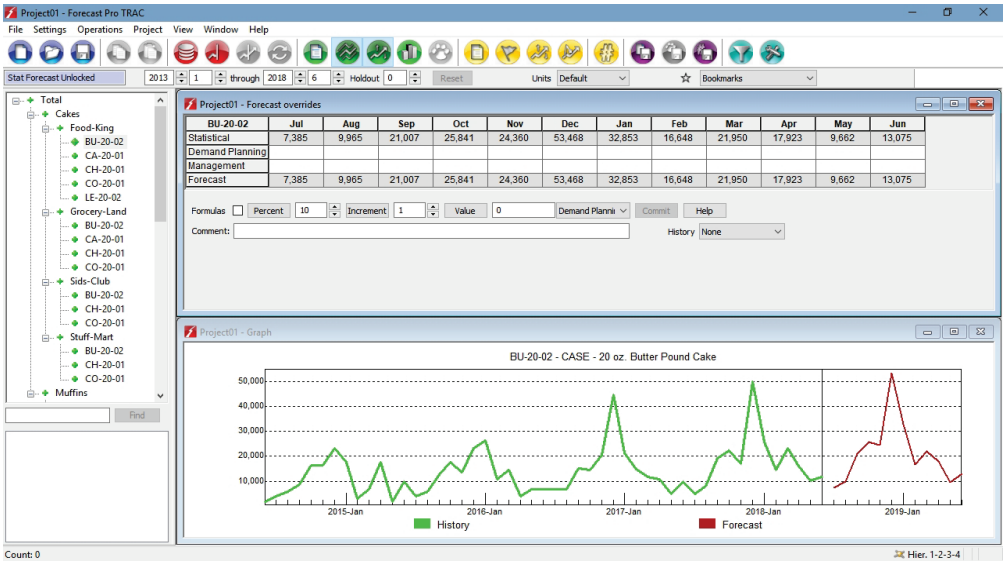

All of the rows currently being displayed are part of the override section of the grid. The first row (currently labeled “Statistical,” which is the default baseline forecast) displays the baseline forecast. This is followed by the override rows and finally the Forecast row. There is a logical order for these rows to appear in and they always share a consistent unit of measure.

You can add additional rows underneath the override section of the grid to display conversions, external data, calculations etc.

Right click on the grid to invoke the context menu and select **Overrides Design.**

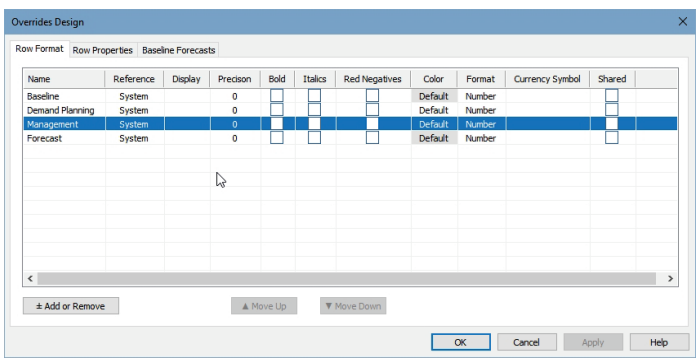

The Row Format tab of the Overrides Design dialog box allows you to control how rows are displayed.

Click on the Add or Remove Button to invoke the Add or Remove dialog box. The Add or Remove dialog box is used to define the optional rows that you wish to work with in the Row Format tab.

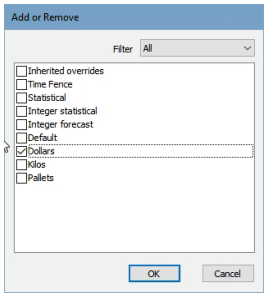

Select Dollars and click OK to return to the Overrides Design dialog box. Notice that there is now a Dollars row displayed on the Row Format tab. Make sure that the Display box for the Dollars row is checked, set the Dollars row Format to Currency and click OK. Notice that our Overrides Grid display now includes a Dollars conversion row.

Importing External Data

At times it may be useful to have visibility to external information while working with the forecasts. In this example, we'll bring in some forecasts that were generated outside of Forecast Pro.

Click the red Data Manager icon ( ) to call up the Data Manager.

Click the Add drop down on the External Data row, select Excel and select *123 Bakery - External Data - SKU-Level Forecasts.xlsx* to add it to the External Data row.

Click the Refresh button on the External Data row to read in the external data.

Click the Exit button to exit the Data Manager.

Right click on the grid to invoke the context menu and select **Overrides Design.**

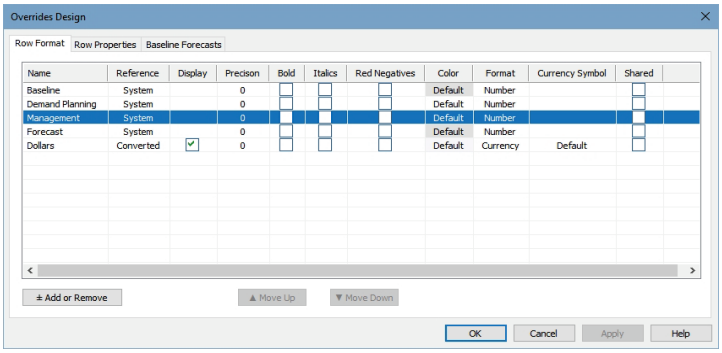

Click on the Add or Remove button to bring up the Add or Remove dialog box. Click on the Filter drop down on the top of the dialog box and select *External Data*. The three new rows that we just imported from Excel are listed.

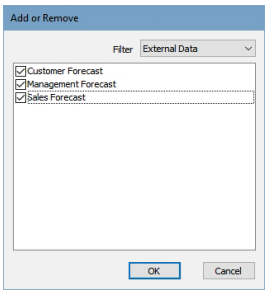

Select all 3 rows and click OK to return to the Overrides Grid dialog box. Notice that the Reference column lists our three new rows as External and all three rows have checkmarks in the Display box. Click Apply to add these rows to the Overrides Grid display without leaving the design box. Remove the checkmarks from the Display column for the Sales Forecast and Customer Forecast rows and click OK to update the display and exit the dialog box.

Select *Total>Cakes>Food-King>Bu-20-02* on the Navigator.

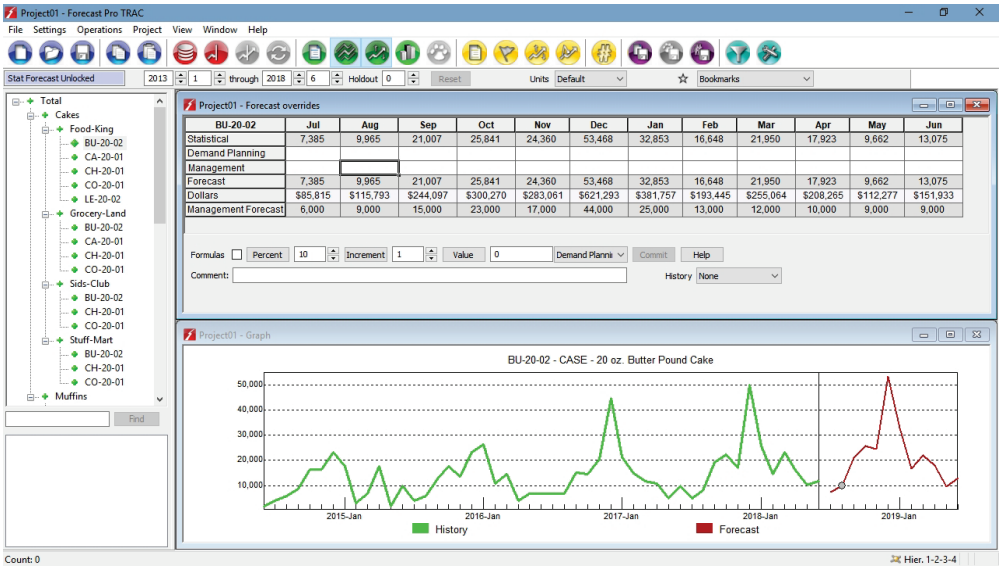

Creating Calculated Rows

We will now add a calculated row to show the difference between our current Forecast Pro forecast and the Management Forecast.

Highlight the Management Forecast row, right click to invoke the context menu and select **Add Row**. A small dialog appears to allow you to name the row. Name the row “The GAP.”

The GAP is a calculated row. You can enter formulas into the cells of calculated rows and they will be used for all items on the Navigator. Select the first cell in The GAP row (July) and enter the formula:

=forecast – {management forecast}

Formulas are not case sensitive. You must use the curly brackets when referencing an external row or a calculated row. “Forecast” is a system row, so the name is a *token* (i.e., pre-defined term) and does not require brackets. System rows are defined within Forecast Pro and are consistently named across projects. A complete listing of supported operations and tokens can be found in *The Override Forecasts View* section of the Command Reference in this manual.

Use the copy and paste commands to copy the formula into the other cells in The GAP row. Click the Commit button to confirm the changes.

Let's add another row to display The GAP as a percentage. Highlight the Management Forecast row, right click to invoke the context menu and select **Add Row**. Name this row "Percent GAP."

Select the first cell in the Percent GAP row (July) and enter the formula:

= {the gap} / forecast

Notice that the value is being rounded to zero. We will reformat the row to display as a percentage in a moment. Use the copy and paste commands to copy the formula into the other cells in the Percent GAP row. Click the Commit button to confirm the changes.

Right click on the override grid to invoke the context menu and select **Overrides Design.**

On the Management Forecast row in the dialog box, select Custom on the Color drop down and set the background color to white in the Color dialog box. Click the Apply button to update the display without leaving the design box.

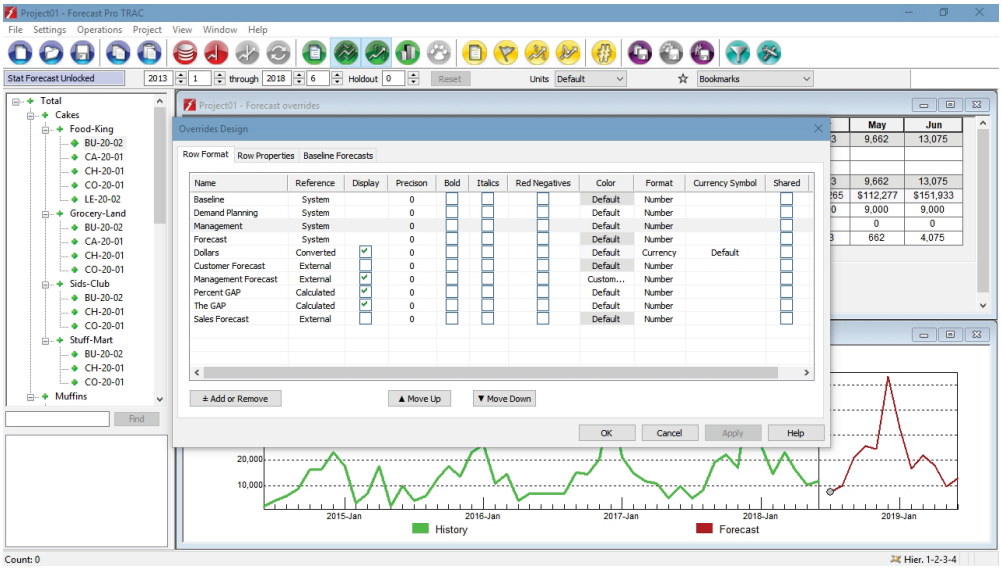

Select the Percent GAP row, set the Format to Percent and check Red Negatives. Select The GAP row, check Bold, check Red Negatives and set the background color to yellow by selecting Custom in the Color drop down. Click the Apply button to update the display without leaving the design box.

Select the Row Properties tab.

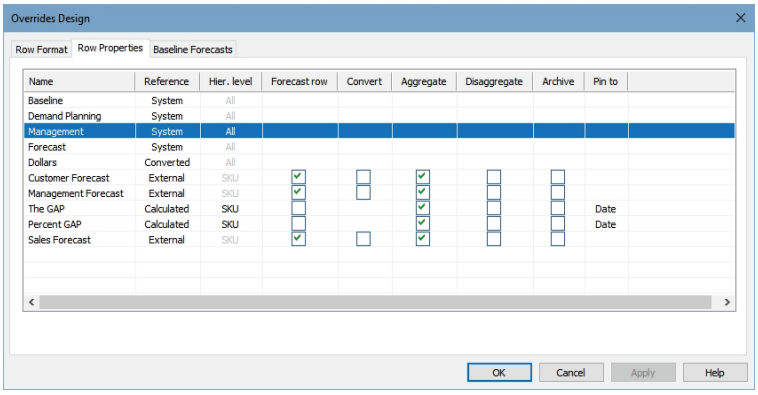

This tab allows you to control how rows are calculated, whether or not they appear at different levels of the hierarchy and how they are stored in the Forecast Pro database.

When defining a calculated row, you can elect either (1) perform the calculations at a single level of the hierarchy and then (optionally) display aggregated and/or disaggregated values at other levels, or (2) calculate the values at all levels. Depending on the formulas you define, these two approaches can result in very different values being displayed.

In our current example we have two calculated rows, The Gap and Percent Gap. Both are currently set to calculate the values at the lowest level (i.e. SKU) and display aggregated values at higher levels. The Gap is simply the difference between two forecast rows and both approaches will yield the same results. The Percent Gap formula defines a percentage and aggregating percentages is very different from calculating them at all levels.

Change the Hier. Level setting for Percent Gap to *All* so that the formula is calculated for all hierarchy levels. Click OK to update the display and exit the dialog box.

The last change that we'll make to the worksheet is to insert a blank row between the Dollars row and the Management Forecast row to make the display a bit easier to read. Highlight the Dollars row, right click to invoke the context menu and select **Add Row**. Delete the default row name "Calculated 3" so that the box is empty and click OK.

Let's now modify the graph to display a bar graph showing the Forecast and the Management Forecast.

Right click on the graph to invoke its context menu and select **Graph Settings**. On the Layout tab, turn off the historic values by unchecking the History rows' Display checkbox. Click the Add or Remove button to open the Add or Remove dialog box. Select Management Forecast and click OK. Make sure that the Management Forecast rows' Display checkbox on the Layout tab is checked. Click the Apply button to update the display without leaving the design box.

Let's create a bar graph. On the Layout tab, use the Forecast rows' Style drop down to change the style from *Line* to *Bar*. Now change the Management Forecast rows' style to *Bar* as well. Click OK to update the display and exit the dialog box. Your display should look like the one shown below.

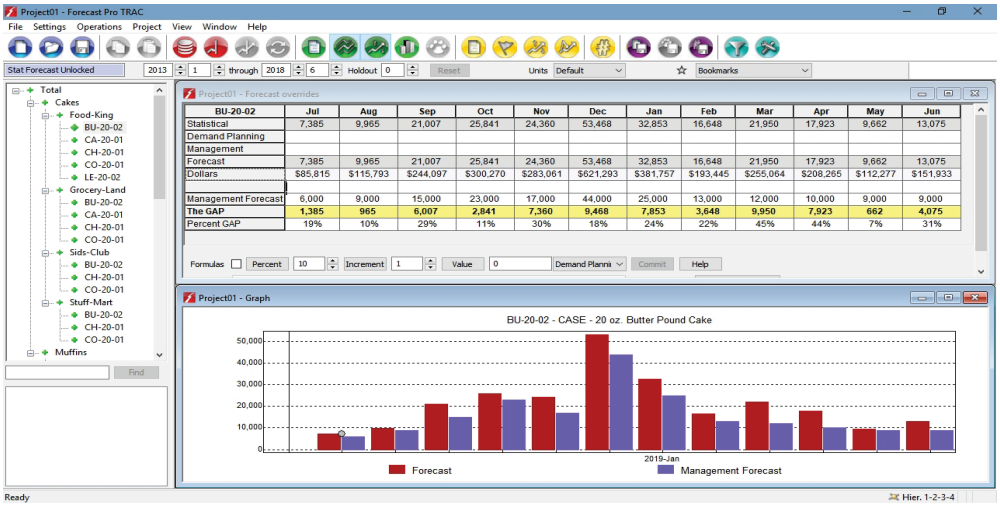

BU-20-02	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun
Statistical	7,385	9,965	21,007	25,841	24,360	53,468	32,853	16,648	21,950	17,923	9,662	13,075
Demand Planning												
Management												
Forecast	7,385	9,965	21,007	25,841	24,360	53,468	32,853	16,648	21,950	17,923	9,662	13,075
Dollars	$85,815	$115,793	$244,097	$300,270	$283,061	$621,293	$381,757	$193,445	$255,064	$208,265	$112,277	$151,933
Management Forecast	6,000	9,000	15,000	23,000	17,000	44,000	25,000	13,000	12,000	10,000	9,000	9,000
The GAP	1,385	965	6,007	2,841	7,360	9,468	7,853	3,648	9,950	7,923	662	4,075
Percent GAP	19%	10%	29%	11%	30%	18%	24%	22%	45%	44%	7%	31%

Continue to experiment with the forecast override and graphical options until you are comfortable with their operation. When you are finished, exit the program. This concludes the lesson.

Lesson 4

Customizing the Graph View

In *The Basics* lesson, we learned how to visually review item-level histories and forecasts with the Graph view. The Graph view can be customized using the Graph Settings dialog box to adjust the data display (e.g. colors, graph type, titles, etc.), edit the data displayed or display multiple items in one graph. In this lesson, we will look more closely at the Graph view functionality.

If you are using a Forecaster license: proceed to the next section—Generating the Forecasts.

If you are using a Collaborator license: the next section describes operations that are not supported. The first step in a Collaborator session is to open a forecast project that was created by someone using a Forecaster license. Select **File>Open,** open the forecast project *"Tutorial – Customizing the Graph"* and proceed to the Customizing the Graph Components section.

Generating the Forecasts

We'll begin this lesson by generating a forecast for the cakes and muffins data that we have explored in prior lessons.

Click the red Data Manager icon ( ) to call up the Data Manager.

Click the Add drop down on the Historic Data row, select Excel and select *123 Bakery – Historic Data.xlsx* to add it to the Historic Data row.

Click the Add drop down on the Conversions row, select Excel and select *123 Bakery – Conversions.xlsx* to add it to the Conversions row.

Click the Add drop down on the External Data row, select Excel and select *123 Bakery - External Data - SKU-Level Forecasts.xlsx* to add it to the External Data row.

Your Data Manager should now match the one below.

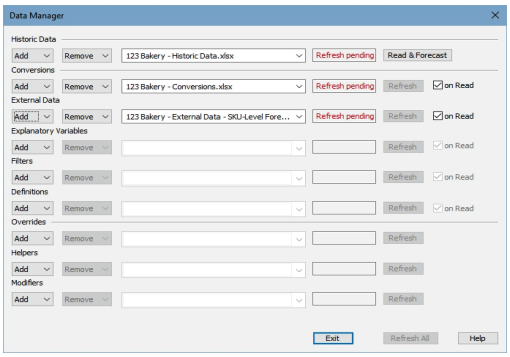

Click the Read & Forecast button to read in the historic demand data and generate the statistical forecasts.

Click the Exit button to exit the Data Manager.

Customizing the Graph Components

Click the green View Graph icon ( ) to open the Graph view and display a graph for *Total*.

Click the green Forecast Report icon ( ) to close the Forecast Report.

Go to the Navigator and right click to bring up the Navigator context menu. Select **Expand All**.

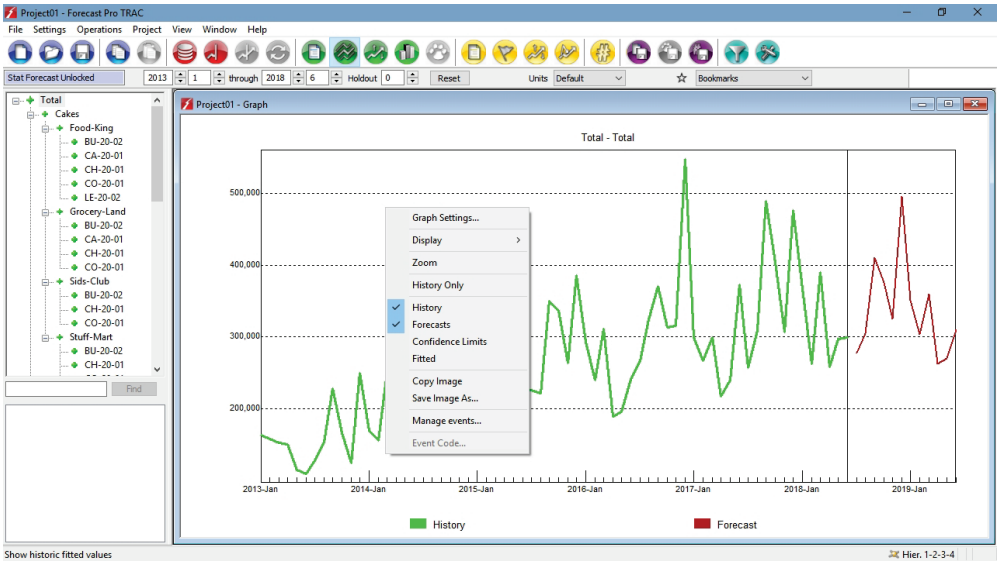

Right click on the graph to display the context menu. Your screen should match the display above.

Notice that various forecast components (History, Forecast, Confidence Limits and Fitted Values) can be toggled on or off. The Graph view will display all forecast components for the active item. Make sure that History and Forecasts are both checked and that the other forecast components are not selected.

Select **Graph Settings** on the Graph context menu. The Graph Settings dialog box includes three tabs. The Layout tab allows you to specify how each series should be displayed, the Components tab specifies global display options and the Items tab allows you to display multiple items on the same graph.

Select the Components tab.

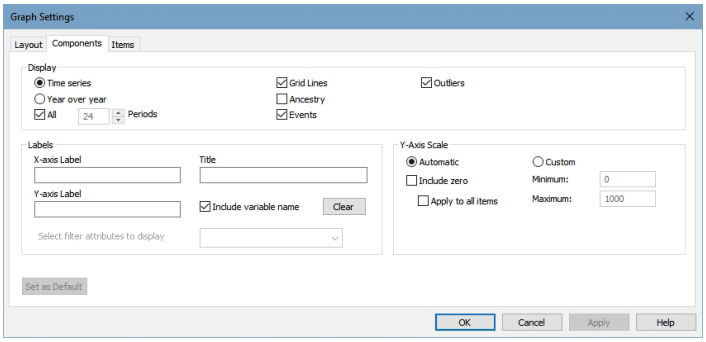

Type “123 Bakery, Inc.” in the Title field and click the Apply button. Notice that the title has been added to the graph. Clicking the Apply button implements your current selection without leaving the Graph Settings dialog box, while clicking OK implements your current selection and exits the dialog box.

Forecast Pro supports two types of graphs—time series and year over year. Select *Year over Year*. Notice that the display is redrawn as a year over year graph and that there is a spinner enabling us to specify how many years to include. Because time series graphs and year over year graphs are distinctly different, you may want to format them differently (e.g., make the year over year a bar graph and the time series a line graph). If you plan on using different formatting options for different types of graph, keep in mind that bookmarks provide an easy way to switch between different formatted displays.

Select *Time series*. When you are viewing a time series graph, you have the option of displaying all of the data in a non-scrollable display or “zooming in” and graphing a subset of the data. Remove the check from the “All” option and click OK. Notice that the graph now displays 36 data points and is scrollable. Use the horizontal scrollbar to scroll through the display.

Right click on the graph to display the context menu. Notice that since you are currently zoomed in the Zoom toggle is checked. Toggle Zoom off to graph the full data set.

Selecting the Data to Include

Select **Graph Settings** on the Graph context menu. Select the Layout tab on the Graph Settings dialog box. Your display should look like below

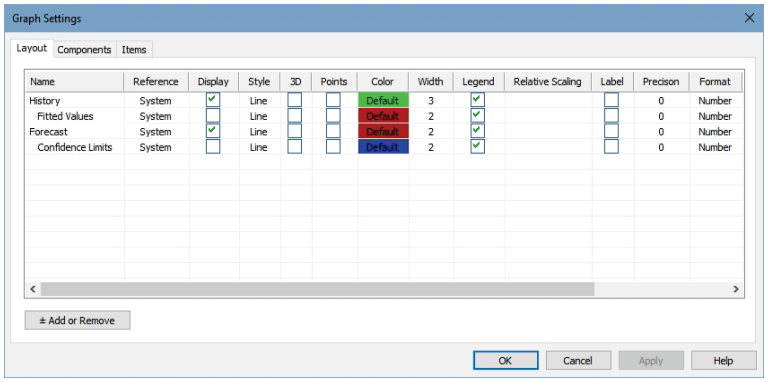

History, Fitted Values, Forecast and Confidence Limits are included in the default layout. Each listed row is displayed on the graph if the Display checkbox is selected. Note that Display is checked for both History and Forecasts, while Fitted Values and Confidence Limits are not checked.

You can add additional rows to the Layout tab using the Add or Remove button. Click the Add or Remove button to invoke the Add or Remove Dialog box. You will see all rows that can be added to the Layout tab (and thus graphed) listed in the Add or Remove dialog box. On the top you can filter the list by the type of row (System, Calculated, External, etc.) Select *External Data* on the Filter drop down selection box.

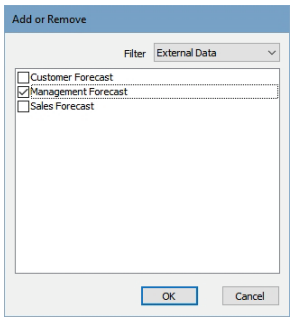

Check “Management Forecast” and click OK to return to the Layout tab of the Graph Settings dialog box. Notice that the Management Forecast row is now

listed and that its' Display field is checked. Click OK. Your display should now look like the one below.

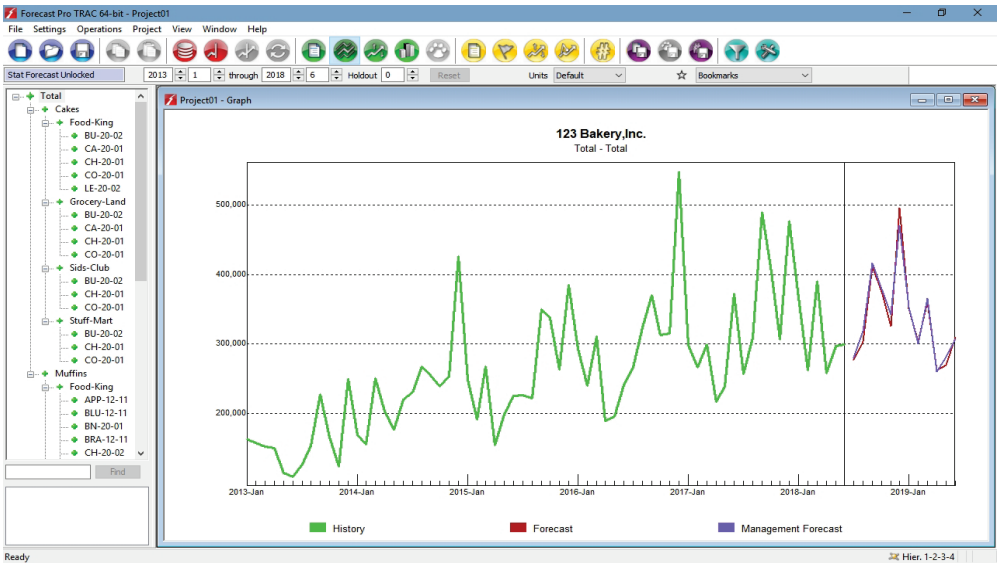

Customizing the Layout

Open the Graph context menu by right clicking on the Graph display. Select **Graph Settings** and select the Layout tab.

The Layout tab allows you to select the graph style for each row, choose custom line colors and widths, specify if legends and/or labels should be used and link color selections to the Forecast Overrides view.

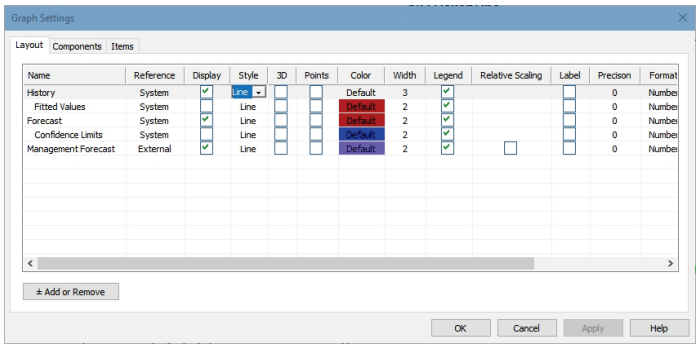

Change the History row's style to bar by clicking on the current style selection (Line) to invoke the drop down selection control and then select *Bar*. Repeat the procedure to change the Forecast row's Style to bar. Click the Apply button to update the Graph view without leaving the Graph Setting dialog box.

Change the color for Management Forecasts by clicking on the current color selection to invoke the drop down selection control and choose *Custom*. A Color dialog box will come up. Select the color of your choice and click OK to close the Color dialog box. Click the Apply button to update the Graph view without leaving the Graph Setting dialog box.

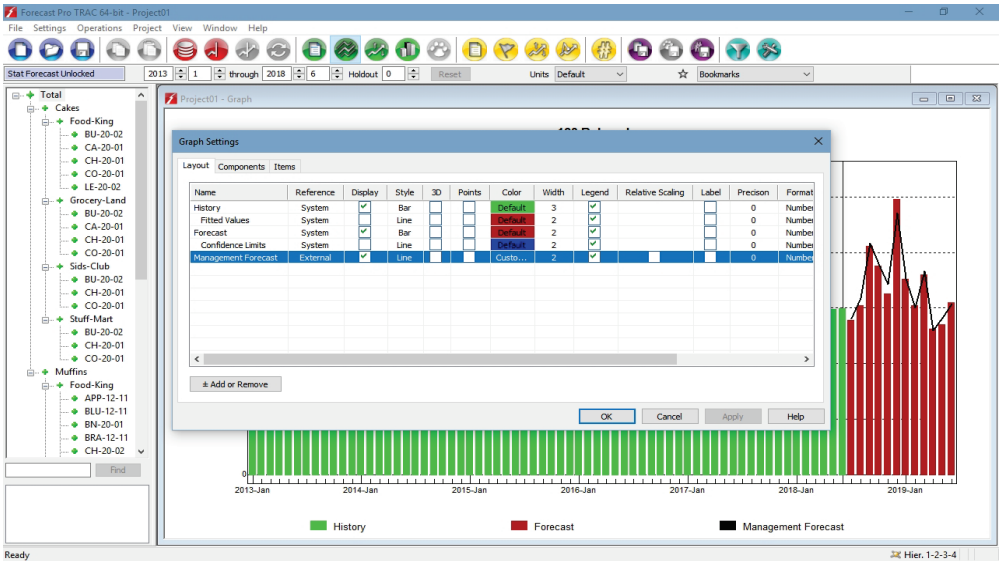

Notice that a Relative Scaling checkbox is available for the Management Forecast row. If this option is selected, Forecast Pro will re-scale the selected row's display to maximize use of the vertical space available. Thus when relative scaling is selected, *the specified row will be shown on a different scale than the other rows*.

Since History, Forecasts, Fitted Values, and Confidence Limits should always be displayed on the same scale the relative scaling option is not available for these series.

Check the Relative Scaling box for the Management Forecast row and click the Apply button. Notice that Management Forecast's display was rescaled. In our example displaying the Management Forecast on a different scale than the

current system forecast doesn't make sense (it hinders the comparison). Uncheck the Management Forecast row's Relative Scaling box and click Apply to redraw the Management Forecast using the same scale as the system forecast.

At times it is useful to coordinate the formatting between the Graph and Forecast override views. Scroll the column display to the right using the horizontal slide bar. At the far right is a column labeled Shared. Selecting Shared for a given row will link your selections for color, precision, format, currency symbol, and shared for the Graph and Forecast override views. This allows you to make changes to the shared row options in either view and have them automatically appear in the other view as well.

Check the Shared checkbox for Forecast row, click OK to close the Graph Settings dialog box and update the graph. Open the Forecast Overrides view by clicking on the green icon ( ). Notice that the color you selected for Forecasts is now used in both the Graph and the Forecast Overrides view.

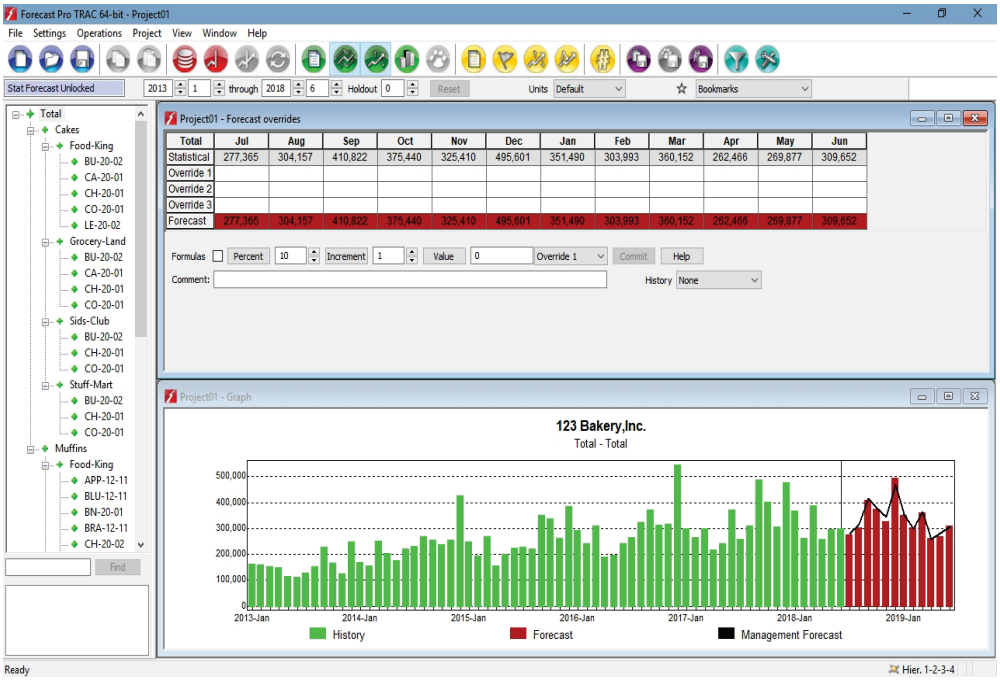

Close the Forecast Overrides view by clicking on the green icon ( ).

Experiment with the graph settings until you are comfortable with their operation.

Viewing Multiple Items

Normally the Graph view is “context sensitive” meaning that when you change your selection on the Navigator, the Graph view is immediately updated to display your current selection.

There may be times when you wish to graph multiple Navigator items on a single graph. In our current example, it might be interesting to view a single graph showing customer-level sales of cakes for our four customers so we can get a feel for the relative sales volumes.

The Items tab on the Graph Settings dialog box allows you to select up to five items to graph together. Because these graphs would be very confusing if they contained lots of information for each item, the graphs are limited to line formats of the history and forecasts for each selection.

Open the Graph Settings dialog box by right clicking on the Graph view to bring up the Graph context menu and selecting **Graph Settings**. Select the Items tab.

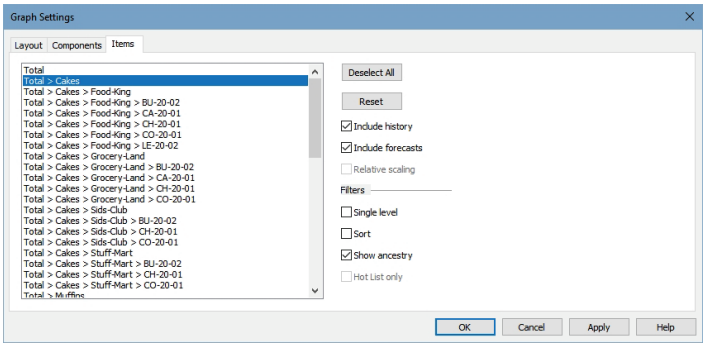

Click the Deselect All button to clear existing selections. Click on each of the following items to highlight them:

Total>Cakes>Food-King
Total>Cakes>Grocery-Land
Total>Cakes>Sids-Club
Total>Cakes>Stuff-Mart

Remove the checkmark from the *Include history* control to limit our graph to the forecasts and make sure Relative Scaling is turned off. Click OK. Your screen should look like the one below.

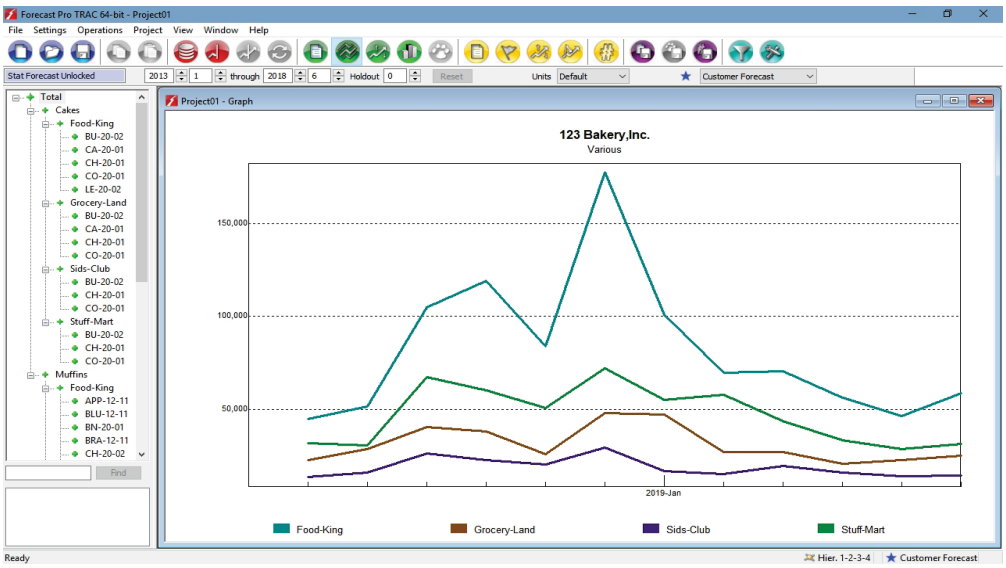

Notice that Food-King is clearly our largest customer, followed by Stuff-Mart, Grocery-Land and Sids-Club.

Once we select a new item on the Navigator, the Graph view will return to the item-specific view. Let's create a bookmark so that we can easily restore this multiple-item graph view. Click on the Bookmark star (☆) to bring up the Bookmark menu and select **Add**. Enter "Customer Forecast" in the Name textbox on the Add Bookmark dialog.

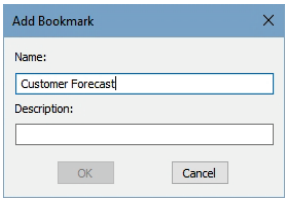

Click OK. Notice that the Bookmark star icon has turned blue and Customer Forecast is displayed in the Bookmark drop down to the right of the star. This indicates that the Customer Forecast bookmark is currently active.

Go to the Navigator Panel and select *Total>Muffins*. The graph is now updated to show *Total>Muffins*, using the Graph Settings specified on the Layout and Components tab. The Bookmark drop down is updated to show *Bookmarks* to reflect that the Customer Forecast bookmark is no longer active. Select *Customer Forecast* on the Bookmark drop down to restore the multiple graph view we created earlier.

Exit the program without saving changes.

This concludes the *Customizing the Graph View* lesson.

Lesson 5

Specifying Alternative Baseline Forecasts

By default, Forecast Pro generates a statistical forecast and then allows you to add overrides to this "baseline" forecast to create the final forecast. There may be times where you wish to use an alternative forecast generated outside of Forecast Pro as the baseline forecast rather than the statistical forecast generated by Forecast Pro.

In this lesson you will learn how to specify alternative baseline forecasts both on an item-by-item basis and for groups of items.

If you are using a Collaborator license: specifying baseline forecasts is not supported. Please proceed to the next lesson.

Specifying the Baseline on an Item-by-item Basis

Start Forecast Pro, select **File>Open** and open the project, *Tutorial - Baselines*. This project contains forecasts for the cakes and muffins data that we worked with in previous lessons.

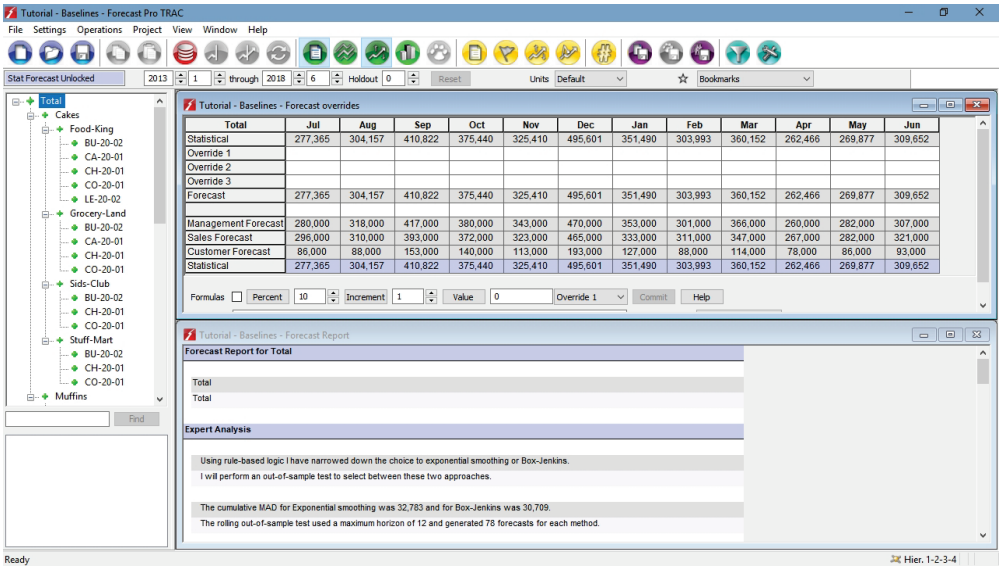

The Forecast Override view includes rows displaying a Management Forecast, a Sales Forecast and a Customer Forecast. These external forecasts were generated outside of Forecast Pro, imported using the Data Manager and added to the Override Grid using the Add or Remove button on the Overrides Design dialog box. Notice there are also two rows displaying the Statistical forecast. One is being used for the baseline and the other was added to the display using the Overrides Design dialog box.

Select *Total>Cakes>Food-King>BU-20-02* on the Navigator. Let's say that for this item rather than using the Statistical forecast as our baseline, we want to use the forecast supplied by our customer. To accomplish this, select the Customer Forecast row by clicking on the row label, then right mouse click to invoke the context menu and select **Set as Baseline**. Notice that the first Statistical forecast row has been replaced by the Customer Forecast row and the Forecast row reflects the change in the baseline.

Select *Total> Cakes>Food-King >CA-20-01* on the Navigator. Let's say that for this item rather than using the Statistical forecast as our baseline, we want to use the forecast supplied by management. To accomplish this, select the Management Forecast row by clicking on the row label, then right mouse click to invoke the context menu and select **Set as Baseline**. Notice that the first Statistical forecast row has been replaced by the Management Forecast row and the Forecast row reflects the change in the baseline.

Select *Total>Cakes>Grocery-Land*. Notice that the baseline row (first row on the current display) is labeled “Statistical.” Select *Total>Cake>Food-King*. Notice that the baseline row (first row on the current display) is labeled “Baseline” due to the fact that we’ve specified alternative baseline forecasts for some of the SKUs that comprise this group.

Forecast Pro allows you to specify alternative baseline forecasts at the end-item level (i.e., lowest level of your hierarchy). When you do so, any higher-level groups’ baseline forecasts will be the aggregate of the lower-level baselines.

Specifying the Baseline for Groups of Items

In addition to specifying the baseline forecasts on an item-by-item basis, Forecast Pro provides the capability to assign baseline forecasts to groups of items all at once.

Right click on the override grid to invoke the context menu and select **Overrides Design**. Select the Baseline Forecasts tab.

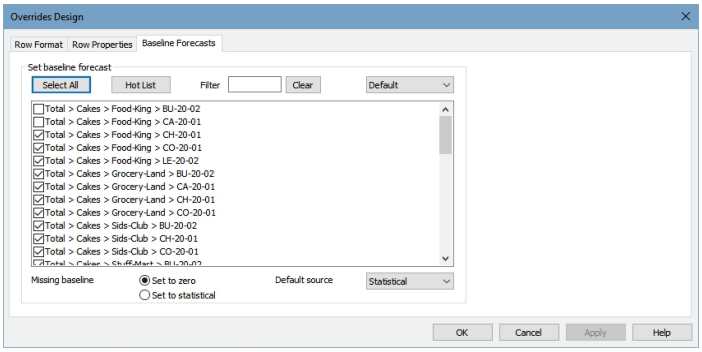

The tab contains a list of all end items. The baseline selection drop down box is currently set to Default and all end items whose baseline is currently set to the Default appear with a checkmark.

Click the Select All button and then click the Apply button. This will set all baseline forecasts to the default source which is currently set to Statistical.

Open the baseline selection drop down box and select Customer Forecast. Notice that none of the end items are checked. Enter the word “Food” in the Filter box. Notice that the list of end items has been filtered and now only includes items whose name or ancestry includes “Food.” In our example, this has filtered the list to all of our Food-King end items.

Click the Select All button and then click the Apply button. This will set all baseline forecasts (i.e., all Food-King SKUs) for our current filtered list to the Customer Forecast.

Click the OK button to exit the dialog box and then examine different items on the Navigator to confirm that the baselines were assigned properly.

Continue to experiment with the baseline forecast assignment options until you are comfortable with their operation. When you are finished, exit the program without saving changes to the *Tutorial - Baselines* project.

This concludes the *Specifying Alternative Baseline Forecasts* lesson.

Lesson 6

Working with Conversions and Alternative Hierarchies

In this lesson we will explore how Forecast Pro allows you to view and work with your forecasts in different units of measure and different hierarchical structures.

The Default State

When you set up your data for Forecast Pro, the units of measure you use for the historic demand and the hierarchy that you define are referred to as the *default* units of measure and the *default* hierarchy respectively. Together they comprise the *default state*.

Let's illustrate this with an example. Start Forecast Pro. Select **File>Open** and open the project *Tutorial – Conversions part 1*.

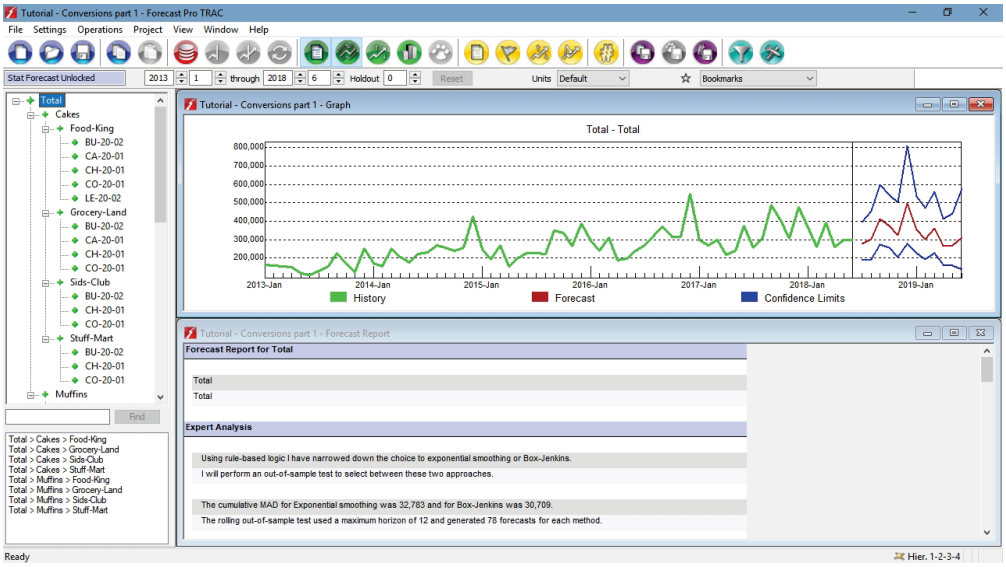

This project uses the 123 Bakery data that we have explored in prior lessons. As you may recall, the data were read in from an Excel file called *123 Bakery – Historic Demand.xlsx*. The spreadsheet contains sales for the different products measured in cases, making cases our default unit of measure. The spreadsheet also defines the default hierarchy, which in this example has four levels—a total level, a category level (consisting of Cakes and Muffins), a customer level and a SKU level.

The combination of our default units of measure displayed and our default hierarchy is referred to as our default state. Notice that "Default" is currently displayed on the dialog bar's Units drop down box and that in the lower right hand corner of the status bar the default hierarchy is listed as "1-2-3-4".

Locked and Unlocked Modes

Changing the display units and/or altering the hierarchy will "lock" the statistical forecasts. When the statistical forecasts are locked, you are no longer able to change the statistical forecasting methods used to generate the statistical forecasts and the sections of the Forecast Report that pertain to the statistical modeling (i.e., Expert Analysis, Model Details, Within-Sample Statistics, Out-of-Sample Tests and Outliers) are omitted. If you subsequently return to the default state (i.e., default units and default hierarchy) the statistical forecasts will be "unlocked."

Working with Alternative Units of Measure

Forecast Pro allows you to define item-level conversion factors (i.e., multipliers) which can be used to display the history and forecasts in alternative units.

You create a conversions file to define the item-level conversion factors and utilize the Conversions row of the data manager to specify and read in the file. The creation of conversion files is fully documented in the Reference section of this manual.

In our current example we defined three alternative units of measure in the conversion file, *dollars, kilos* and *pallets.*

Turn off the Forecast Report view and turn on the Forecast Overrides view. Right click on the override grid to invoke its context menu and select **Overrides Design**.

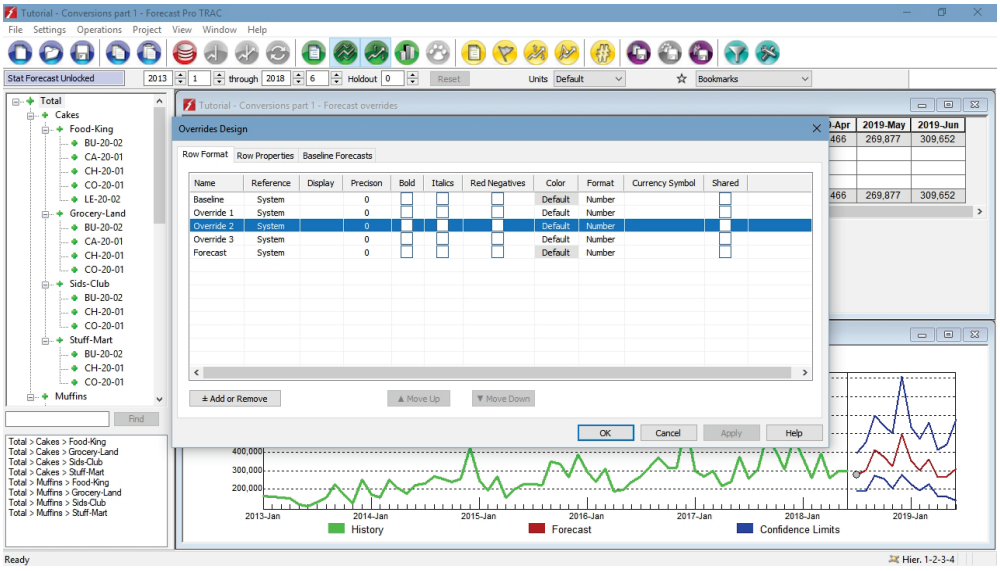

The Overrides Design box gives you a great deal of control over the Forecast overrides display.

Click on the Add or Remove button. Notice that the Add or Remove dialog box includes Default, Dollars, Kilos and Pallets rows.

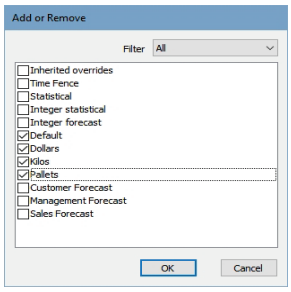

Select all four measures and click OK. For each measure, make sure that the Display box is selected and set the format for Dollars to Currency. Click the OK button.

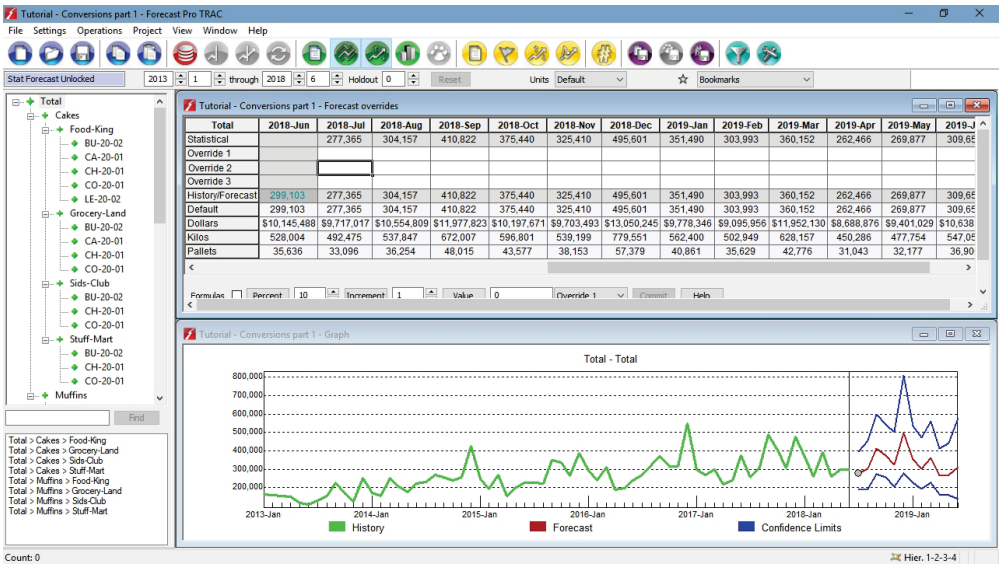

Your display should now look like the one above.

Notice that the conversion rows display the forecasts in the alternative units of measure. This allows you to enter changes in the currently displayed units (in this example the default—cases) and immediately see the impact on the other units of measure.

Open the Units drop down box on the dialog bar and select Dollars.

A warning message informs you that this action will lock the statistical forecasts and move all current overrides into the inherited override row. Changing the Units back to Default will unlock the statistical forecasts and restore overrides. Click OK to continue.

Notice that the history and forecasts are now displayed in dollars and the status box above the Navigator reads “Stat Forecast Locked”.

You are now working with dollars. Any overrides that you make to the forecasts, reports that you generate, etc. will be in dollars.

Changing the Hierarchy

In our current example we have a 4-level hierarchy. As we commented previously, because we are currently viewing the default hierarchy (i.e., the one defined in the historic data file) the status bar lists the hierarchy as 1-2-3-4.

Select *Total>Cakes>Food-King* on the Navigator. This displays the total cake sales for Food-King. The “children” (i.e. component series) of this group are the SKU-level cake sales to Food King.

What if we wanted to view or adjust total sales for a given SKU *across* customers? The current hierarchy is not set up to do this. To accomplish this we need to reorganize our hierarchy so that the SKUs appear on level 3 of the hierarchy and the customers appear on level 4.

Right click *Total> Cakes>Food-King* on the Navigator to invoke the context menu. Select **Hierarchy>Demote**. Fully expand the Navigator and select BU-20-02. Your display should look like the one below.

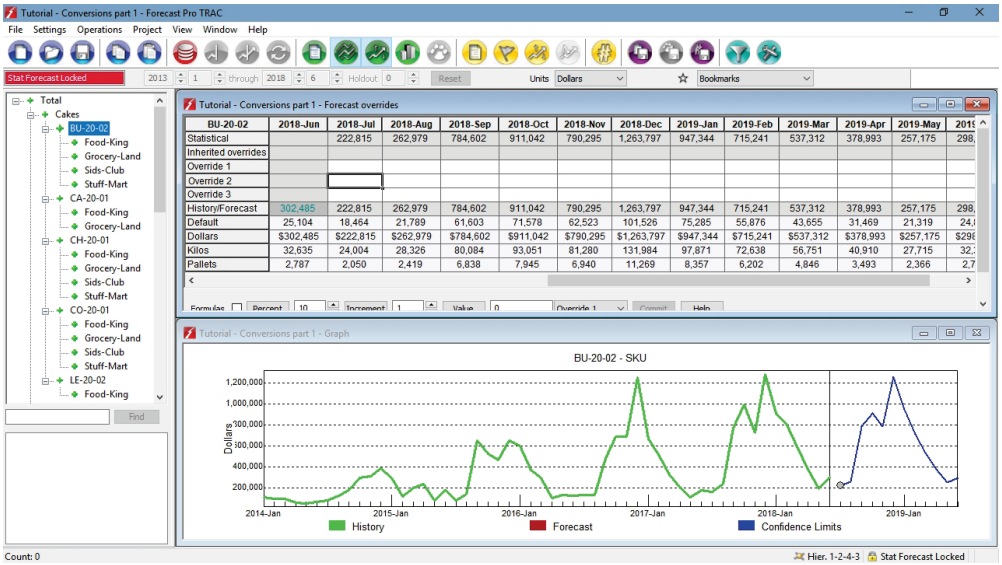

Notice that by demoting level 3 of the default hierarchy we have created a new hierarchy that allows us to view and adjust total sales by SKU across customers. The status bar lists this new hierarchy as 1-2-4-3 to reflect the fact that we demoted the original level 3.

Experiment with the hierarchy adjustment and conversion options until you are comfortable with their operation. Then select **File>Close** to close the project. When prompted, do not save the changes to the *Tutorial – Conversions - part 1* project.

Working with Override Sets

In this section we will explore how Forecast Pro maintains your overrides when you change units or modify the hierarchy.

Select **File>Open** and open the project *Tutorial – Conversions part 2*.

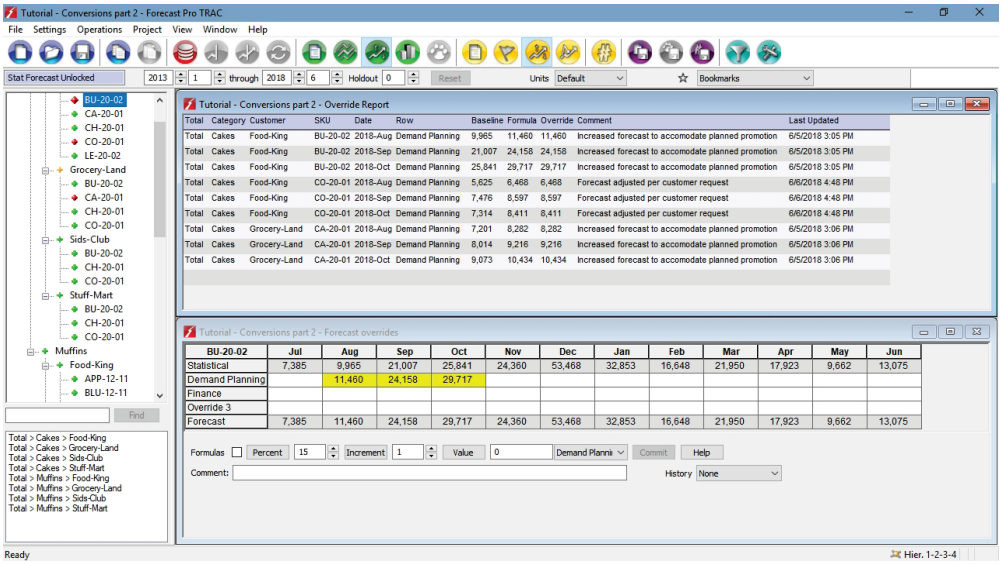

The top window is displaying the Override Report view. This view shows all current overrides in the project. The Override Report view is turned on and off using the yellow View Override Report icon ( ).

Notice that this project is in the default state (i.e., default units and default hierarchy), the statistical forecasts are unlocked and overrides exist for three of the SKUs (as indicated by the red icons next to the SKUs).

The override view allows you to enter overrides and comments in specific cells on the override grid. We refer to these overrides as appearing *in-place*.

When you change units and/or change the hierarchy, it is not always possible to maintain the in-place overrides. This is due to factors such as the item or group that had an in-place override may no longer exist after you change the hierarchy or that changing the units would result in a different forecast allocation across the hierarchy.

To avoid these problems and to insure that the final forecasts will remain consistent regardless of the hierarchy or units you are viewing, Forecast Pro will move the in-place overrides into a special override row labeled *Inherited overrides* prior to performing hierarchy changes and/or conversions.

Select the Units drop-down box on the dialog bar and select Dollars.

A warning message informs you that this action will lock the statistical forecasts and move all current overrides into the inherited override. Click OK to continue.

Notice that the in-place overrides have disappeared, the Override Report is now empty and there is now a row labeled *Inherited overrides* that contains the changes.

Right click on the Override Report to invoke its context menu and select **Override Report Design**.

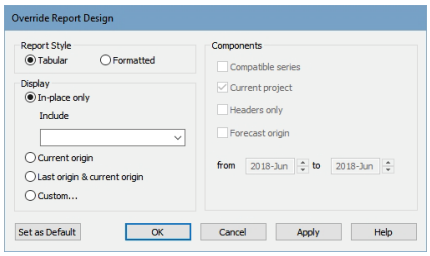

Notice that the Override Report is currently set to display in-place overrides only. Since there are no current in-place overrides the display is empty.

Select “Current origin” and click OK. This will display all overrides associated with the current project during the current forecast period.

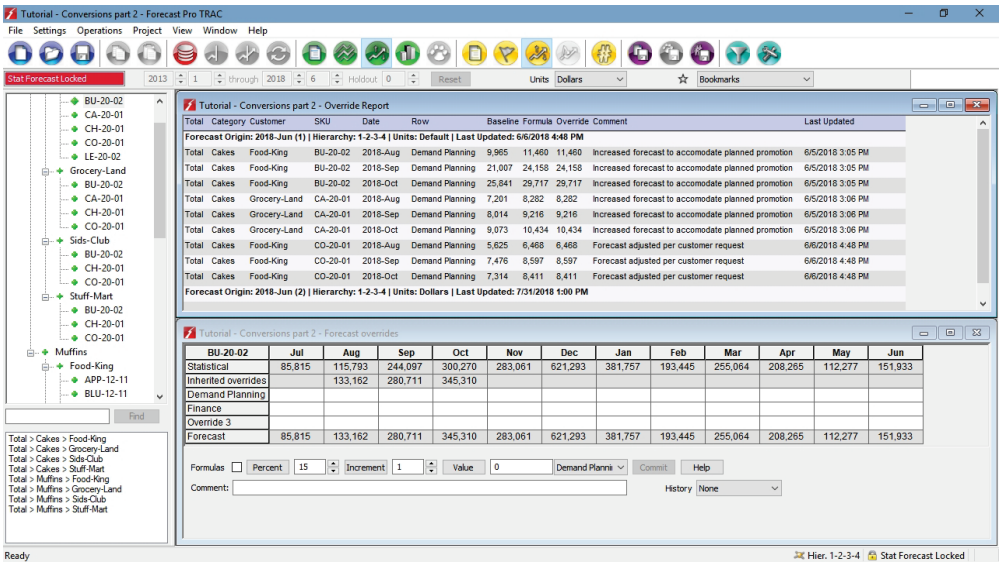

The report above displays two *override sets*. The first one captures the overrides made to the initial default state. Its header row identifies the forecast origin (i.e., last historic data point) as June 2018, the hierarchy as 1-2-3-4 and the units as Default. The second override set's header row identifies the forecast origin as June 2018, the hierarchy as 1-2-3-4 and the units as Dollars.

Select *Total* on the Navigator. Enter the value 14,000,000 into the September override cell on the Finance override row and click Commit. Notice that the Override Report is updated to display your latest change.

The overrides and comments entered in every state (i.e., hierarchy/unit combination) are saved to the database as an "override set." Override sets can be viewed in the Override Report, imported and exported to projects and loaded from the database's archive when you update your project.

If possible, it is a good practice to limit your overrides to a single state. This allows you to view all overrides in-place and vastly simplifies loading previously entered overrides when you update your project.

Exit Forecast Pro without saving changes to the *Tutorial – Conversions part 2* project.

This concludes the *Working with Conversions and Alternative Hierarchies* lesson.

Lesson 7

Updating a Forecast Project

For most organizations updating the forecast is a routine operation that occurs every planning period. For instance, if the forecasts are prepared using monthly data, then each month the data are updated to include the latest observation and the forecasts are regenerated.

Forecast projects enable you to quickly update last period's forecast and give you the option of loading or not loading your previously specified forecasting models, Hot Lists and forecast overrides. In this lesson we will explore how this is accomplished.

If you are using a Collaborator license: updating forecasts is not supported. Please proceed to the next lesson.

Defining the Forecasting Process

An important key to implementing a successful forecasting process using Forecast Pro is to define the process upfront and make sure that everyone involved understands their role in establishing the final numbers.

Some of the more important questions to answer before you start include:

What output is required and what format should it be in?

What naming conventions should be used for data files, output files and forecast projects, where should these files be kept and how will they be backed up?

Who is authorized to enter forecast overrides, at what level(s) are they entered and in what units are they entered?

What are the steps required to update our forecasts each planning period?

Depending on the complexity of your forecasting process and the number of individuals involved, it may be useful to document the decisions made and the steps required to update the forecasts each planning period.

Step One: Updating Your Data Files

Consider the following situation. We work for the 123 Bakery. We have defined an extremely simple forecasting process where the forecasts are generated and adjusted by one individual. All adjustments are made to the SKU-level data in base units and in the base hierarchy. The projects are saved with the statistical forecast unlocked (i.e., we never change states).

It is June 2018 and we have historic demand data that starts in 1/2012 and ends in 5/2018. Thus, our first forecast period will be June 2018. We use Forecast Pro to create our forecasts and then we save a forecast project named *Tutorial – Updating - June 2018*.

Now imagine that a month passes by. We now have our sales figures for June and we update our historic data files.

This is accomplished externally to Forecast Pro. It may entail running a data extraction routine to generate the new files, updating your spreadsheets by hand, or some other process to update the historical data files.

Step Two: Opening Your Forecast Project

Start Forecast Pro, select **File>Open** and select the project, *Tutorial - Updating - June 2018*. The dialog box below appears.

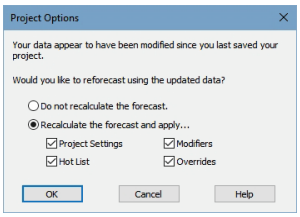

Forecast Pro has noticed that the forecast project *Tutorial – Updating - June 2018* and the data file *123 Bakery.xlsx* are out-of-sync. It is giving you two options.

If you select "Do not recalculate the forecast", Forecast Pro will open the project and restore the forecasts session as it existed when it was saved last month. *It will not read the new data. It will not revise the forecasts.*

Select "Do not recalculate the forecast" and click OK to verify this. Notice that the historic data ends in May of 2018 and the first forecast period is June 2018.

Now select **File>Open** and select the project, *Tutorial – Updating - June 2018* again. The data file and project are still out of sync so the dialog box reappears.

If we now select "Recalculate the forecast and apply", Forecast Pro will read in the revised data, generate new forecasts and (if you request them) restore the project settings, Hot List, forecast modifiers and overrides.

Select "Recalculate the forecast and apply" and click OK. Expand the Navigator and select *Total>Muffins>Food-King>CA-20-01*. Display the Graph view by clicking its green icon (), and display the Override Forecasts view by clicking its green icon (). Click the green Forecast Report icon () to close the Forecast Report.

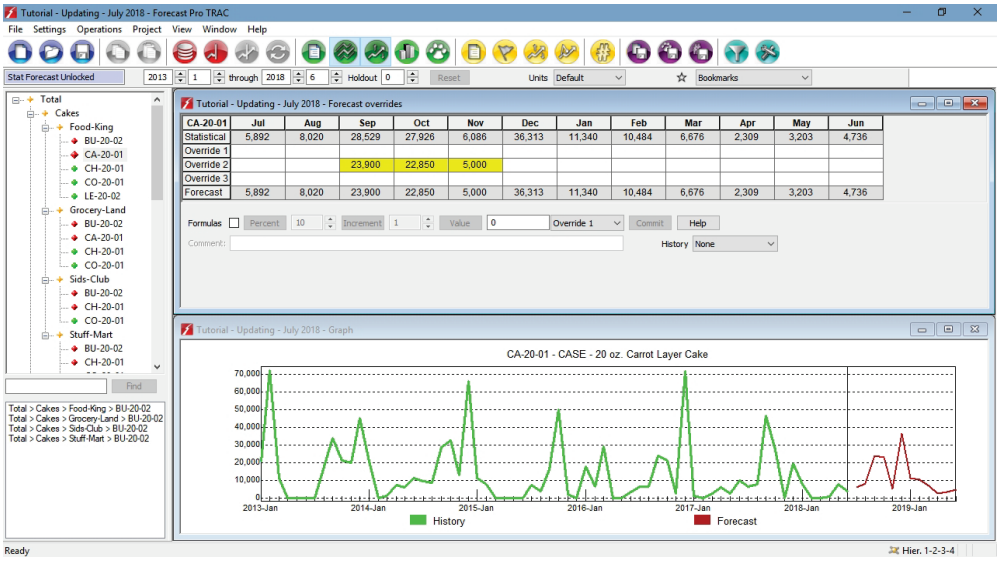

Notice that the historic data now ends in June 2018 and the first forecast period is July 2018. Notice also that Forecast Pro has restored the Hot List, forecast modifiers and overrides and that the overrides correctly reflect the new dates.

At this point we have generated a new forecast. Let's save this project using a different name so that we can retain both our *Tutorial – Updating - June 2018* project and our new July forecasts. Select **File>Save As** and name the project *Tutoria l- Updating - July 2018*.

Step Three: Saving Your Work

If this was a real forecasting session we would now review the new forecasts, add any needed overrides, adjust the forecasting models as appropriate and share the forecast project with others, etc. until we establish the final July forecast.

Once we've established the final forecast we would save our final project file for the July forecast and save any desired output files.

Exit Forecast Pro using **File>Exit**.

Working with Override Sets

In our last example we updated a very simple project that did not include inherited overrides. This means that the person who created the project only entered overrides in base state (i.e., they did not enter overrides after changing units nor changing the hierarchy). When you update a project that does not include inherited overrides, Forecast Pro allows you to automatically load the overrides.

If you update a project that contains inherited overrides then you will have to explicitly load any override sets that you wish to apply using **Project>Import>Archived Overrides**. Let's see how this works.

The project we will use is very similar to the project we looked at previously; the only difference is that it contains overrides both in the base state and after converting to dollars. Thus, the project included inherited overrides.

Start Forecast Pro, select **File>Open** and select the project *Tutorial –Updating - part 2*. The dialog box below appears.

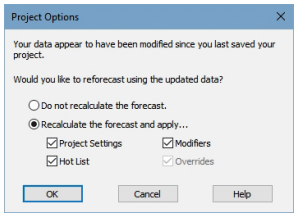

Notice that the option to automatically load the overrides is greyed out. This is because the project contained overrides that were made in a non-base state (e.g., in a converted unit of measure and/or in a shuffled hierarchy). Select "Recalculate the forecast and apply" and click OK. Forecast Pro will read in the revised data and generate new forecasts.

If you are using a Collaborator license, the above dialog box will not appear because reforecasting is not supported. Please proceed with the rest of this lesson.

Close the Forecast Report window and open up the Override Grid by clicking on the override grid icon.

Click on the yellow View Override Report icon ( ) to open the override report view. Right click the override report view to bring up its context menu and select **Override Report Design**.

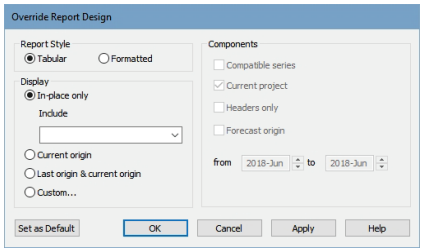

Select "Last origin & current origin" and click OK. This will display the override sets associated with the current forecast period and the previous forecast period. This is a useful display to view when you are loading override sets.

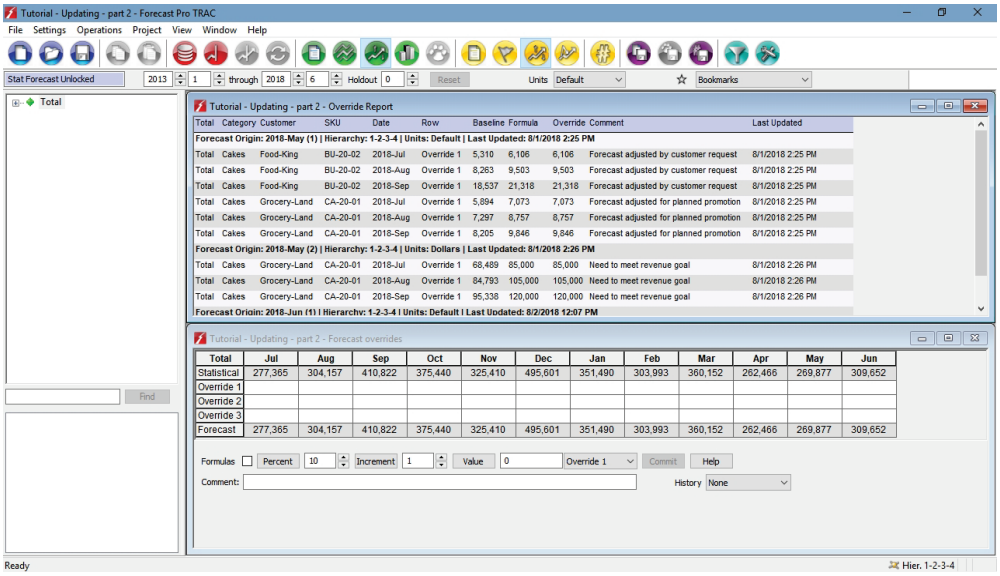

Notice that there are three override sets displayed. The first set's origin is May (i.e., the last historic data point was May, first forecast was June) and the overrides were made in base hierarchy (1-2-3-4) and base units (default). The second set's origin is also May and the overrides were also made in base hierarchy, but the units were dollars.

The third set's origin is June. It is currently empty since no overrides have been entered nor have any override sets been loaded.

Select **Project>Import>Archived Overrides**.

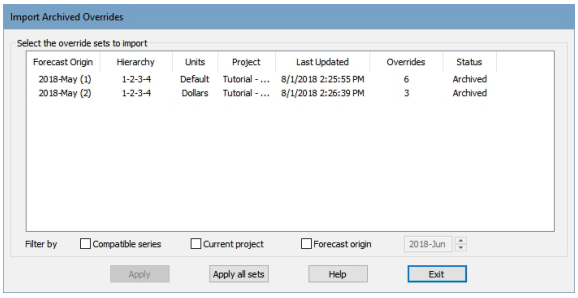

The Import Archived Overrides dialog box allows you to load any override set that is in the Forecast Pro database. In our example, the two May-origin override sets are the only entries in the database.

The "Apply all sets" button will apply all override sets currently being displayed. The Apply button can be used to load specific override sets one at a time.

In our example, "Apply all sets" would load the first set (which was made in base state) then convert to dollars and load the second set. Since converting to dollars requires moving overrides made in base state into the inherited overrides row, if you select "Apply all sets" in our current example your project will contain an inherited overrides row after the overrides are loaded.

Let's say that our forecasting process calls for the forecaster to update the statistical forecasts, enter overrides in base state and then save the revised project which then goes to finance, who converts to dollars and makes additional changes.

If this was our process and we were the forecaster, we'd probably want to begin by restoring the override set we saved last forecast period.

Click on the first override set to select it. Click Apply to load the override set. Click Exit to exit the Import Archived Overrides dialog box. Expand the Navigator and select *Total>Muffins>Food-King>CA-20-01*. Notice that the

overrides have been loaded and that since these overrides were made in base state, there is not an inherited row.

As you can see, updating your forecasts becomes considerably more complex when there are multiple override sets involved. For this reason, we recommend that you create a forecast process that uses as few override sets as possible. Ideally, you'd limit your overrides to a single state. This allows you to view all overrides in-place and vastly simplifies loading previously entered overrides when you update your project.

Exit Forecast Pro without saving changes to the *Tutorial – Updating - part 2* project.

This concludes the *Updating a Forecast Project* lesson.

Lesson 8

Collaborating With Colleagues

Corporate forecasting often requires collaboration among colleagues. Perhaps your manager wishes to review your forecasts and make a few changes. Perhaps you want to run your forecasts past the sales force who may know about pending orders. Perhaps your business has distinctly different markets that require specific individuals to work on specific markets. Whatever the reason, it is often useful to be able to work with others to establish the final forecasts.

In this lesson we will explore how Forecast Pro supports collaborative forecasting.

Stand-alone Project vs. Super Project

There are two different strategies for working with others using Forecast Pro—sharing a single stand-alone project and consolidating multiple projects into a "super project" using the Super Project Manager. There are pros and cons to each approach.

Sharing a stand-alone project is by far the simplest way to work with others to establish the final forecasts. Using this approach the forecaster begins by creating forecasts for all items that need to be forecasted, enters any desired

overrides and saves a forecast project—often to a network drive. The project can then be opened by anyone using Forecast Pro (Forecaster or Collaborator licenses) to review, make changes, etc. If there is a third individual who needs to review the project, that person would do so after the second individual has updated the project.

Alternatively, the super project approach allows you to break large forecasting jobs into smaller pieces that can be worked on separately and then consolidated. To implement this approach you first create, forecast and save separate component projects and then use the Super Project Manager to combine the component projects into a new super project that allows you to view and manipulate the consolidated forecast. A Super Project can only be created by someone with a Forecaster license, but it can be used by anyone using Forecast Pro (Forecaster or Collaborator licenses).

A forecasting process that uses a stand-alone project is *much* simpler to implement and maintain than one that uses a super project. If it is not essential that your total hierarchy be broken out into smaller pieces, then sequential review of stand-alone projects is the preferred approach.

Sharing a Stand-alone Project

In this lesson we will open a forecast project that was created by someone else, review their work and make some overrides.

Providing a colleague access to your forecast project is a simple way to share your work. As long as your colleague has Forecast Pro (Forecaster license or Collaborator license), he or she can open the project and continue the session.

To illustrate how to share a project we'll open up a project that is included with Forecast Pro. Start Forecast Pro, select **File>Open** and open the project, *Tutorial – Collaborating - part 1*. The data in this project represent sales of different types of cameras and film.

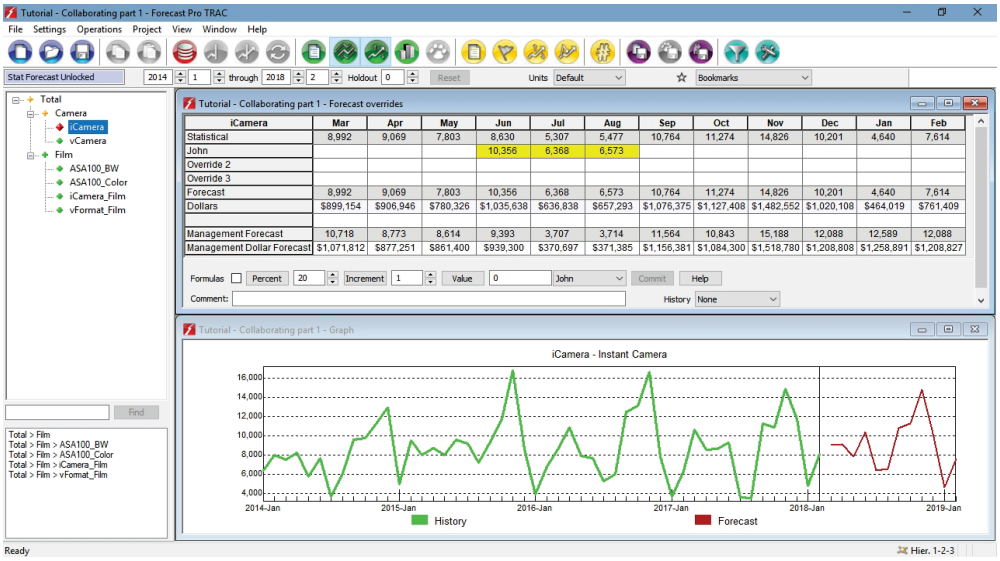

iCamera	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb
Statistical	8,992	9,069	7,803	8,630	5,307	5,477	10,764	11,274	14,826	10,201	4,640	7,614
John				10,356	6,368	6,573						
Override 2												
Override 3												
Forecast	8,992	9,069	7,803	10,356	6,368	6,573	10,764	11,274	14,826	10,201	4,640	7,614
Dollars	$899,154	$906,946	$780,326	$1,035,638	$636,838	$657,293	$1,076,375	$1,127,408	$1,482,552	$1,020,108	$464,019	$761,409
Management Forecast	10,718	8,773	8,614	9,393	3,707	3,714	11,564	10,843	15,188	12,088	12,589	12,088
Management Dollar Forecast	$1,071,812	$877,251	$861,400	$939,300	$370,697	$371,385	$1,156,381	$1,084,300	$1,518,780	$1,208,808	$1,258,891	$1,208,827

Notice that we did not define the historic data in the Data Manager nor did we use the Read & Forecast button. When we opened the forecast project, Forecast Pro restored the historic data, forecasts, overrides and Hot List that were present when the forecast project was saved. The person restoring the forecast project does not even need to have access to the files containing the historical data that were used to generate the forecasts—the only thing he or she needs is the forecast project file.

Notice that the person who created the forecasts entered some overrides in June, July and August. For convenience, let's refer to the person who created this file as "John". Notice also that the background color for these cells is yellow. This indicates that the cells contain comments. Click on the June override cell to display the comment.

The comment informs us that John increased the forecast for the summer months by 20% in anticipation of a planned promotion. Let's say that we feel that a 20% lift is overly optimistic for August (the final month of the promotion) and we want to lower it to 15%. Set the target row box (located to the immediate left of the Commit button) which currently reads *John* to *Override 2*. Highlight the cell for August on the Statistical forecast row. Set the percent box to 15 and click the Percent button. Let's now add a comment. Highlight the cell for the override you just made, click the comment field and type, "John, I lowered the promotional lift for summer madness to 15%" and click the Commit button to accept the overrides and comment.

Now let's edit the row label *Override 2*. Click on the row label *Override 2* and right mouse click to call up the context menu. Select **Edit Row Label**. Enter *Demand Planning* and click OK. Notice that the row label has changed.

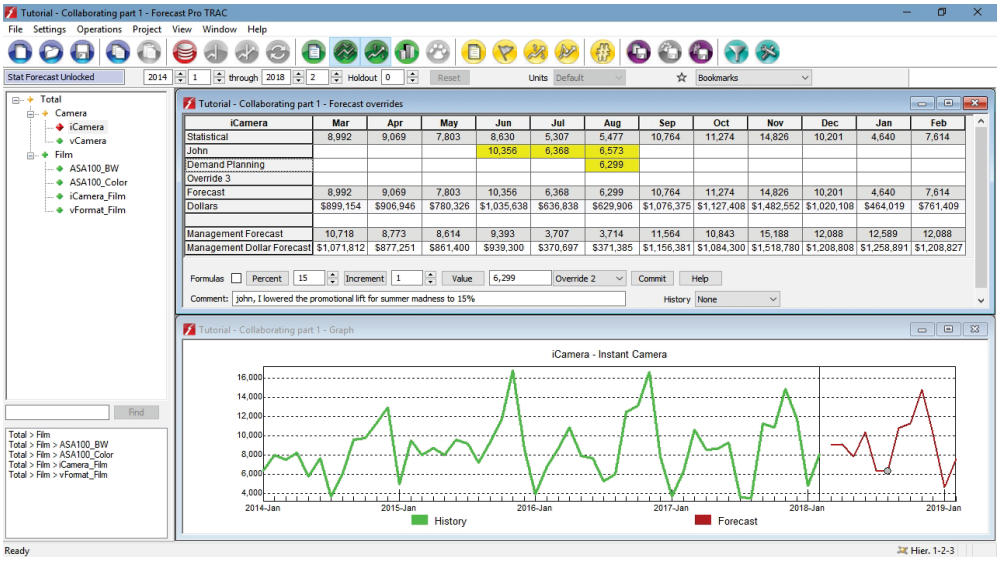

iCamera	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb
Statistical	8,992	9,069	7,803	8,630	5,307	5,477	10,764	11,274	14,826	10,201	4,640	7,614
John				10,356	6,368	6,573						
Demand Planning						6,299						
Override 3												
Forecast	8,992	9,069	7,803	10,356	6,368	6,299	10,764	11,274	14,826	10,201	4,640	7,614
Dollars	$899,154	$906,946	$780,326	$1,035,638	$636,838	$629,906	$1,076,375	$1,127,408	$1,482,552	$1,020,108	$464,019	$761,409
Management Forecast	10,718	8,773	8,614	9,393	3,707	3,714	11,564	10,843	15,188	12,088	12,589	12,088
Management Dollar Forecast	$1,071,812	$877,251	$861,400	$939,300	$370,697	$371,385	$1,156,381	$1,084,300	$1,518,780	$1,208,808	$1,258,891	$1,208,827

At this point we have several options on how we might save our work. For instance, we could:

Use **File>Save as** to save a new forecast project to send back to John.

Click the purple Save Numeric Output icon ( ) to save an output file to submit to production, management or some other destination.

Click the purple Save Forecast Report icon ( ) to save formatted reports to Excel.

Exit Forecast Pro without saving changes to the *Tutorial – Collaborating - part 1* project.

Working with Super Projects

Let's consider the following example. A company sells their products in three countries—the United States, Canada and Mexico. They employ three different demand planners, each of whom is responsible for forecasting a

specific country. After the demand planners have completed their forecasts, they meet with management to review the forecasts and make any required adjustments.

Accomplishing this using a single stand-alone project would require a sequential review process. For example, assume the demand planner for the United States goes first. He or she would create a forecast project which includes all three countries and then review and make adjustments to the United States data and save the project. Next, the planner for Canada would open the project, review and adjust the Canadian data and save the revised project. Next, the planner for Mexico would open the project, review and adjust the Mexican data and save the revised project. At this point, the demand planners' forecasts are complete and the project is ready for the management review meeting.

Clearly, there are some downsides to the above procedure. One is that the demand planners cannot work on their country-level forecasts simultaneously—they have to wait until it is their turn to do so. Another problem is that even though they are working on a single country, the project contains data and forecasts for all three countries, which makes it harder to manage and more prone to error.

An alternative is to use the super project approach. This allows you to create a separate forecasting project for each country and then consolidate them into a new super project which includes all three countries. Using this approach, the three demand planners can work on their forecasts simultaneously using a project that only contains their country. When all three are done, the Super Project Manager is used create a super project containing the consolidated forecast.

Start Forecast Pro, select **File>Open** and open the project, *Tutorial – Collaborating - USA*.

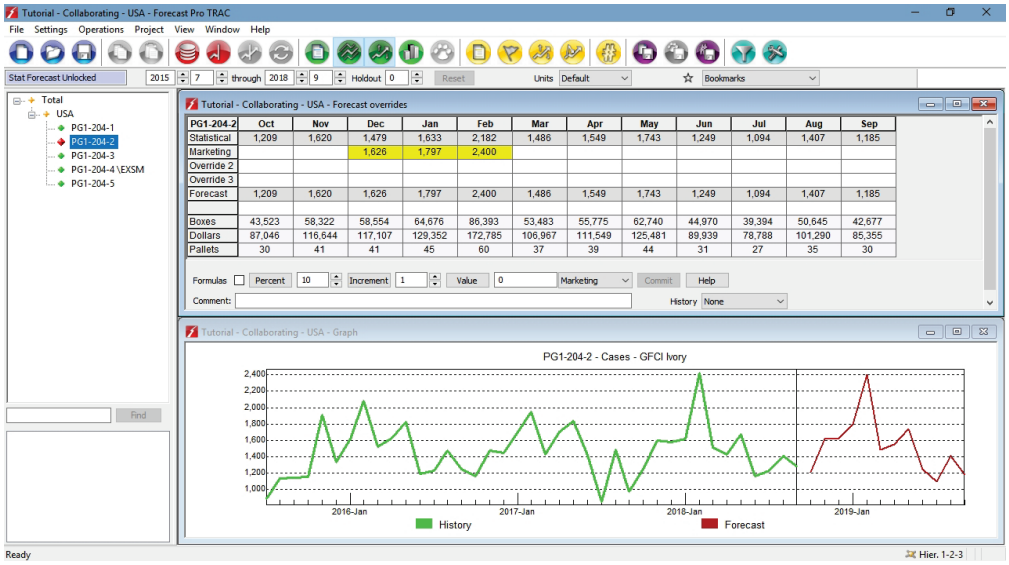

This is a component project containing data for the United States. Notice that it includes forecast overrides for *Total>USA>PG1-204-2* and a forecast modifier for *Total>USA>PG1-204-4*.

Select **File>Open** and open the second component project— *Tutorial – Collaborating - Canada*. Open the project *Tutorial – Collaborating - Mexico* to view the third component project. As you have observed, the component projects have all been forecasted and include some overrides and forecast modifiers.

If you are using a Collaborator license, completing the remainder of this lesson requires that you have a Forecaster license. Please proceed to the next lesson.

If you are using a Forecaster license, close the current project by selecting **File>Close** and then open a new blank project by selecting **File>New**.

Click the red Data Manager icon ( ) to call up the Data Manager.

Click the Add drop down on the Historic Data row, select Excel and select the following three files:

Hardware USA – Historic Data
Hardware Canada – Historic Data

Hardware Mexico – Historic Data

These were the data files used to create the three component projects. Click the Read & Forecast button to read in the historic demand data and generate the statistical forecasts. Exit the data manager, expand the Navigator and open the graph window. Your project should now look like the one below.

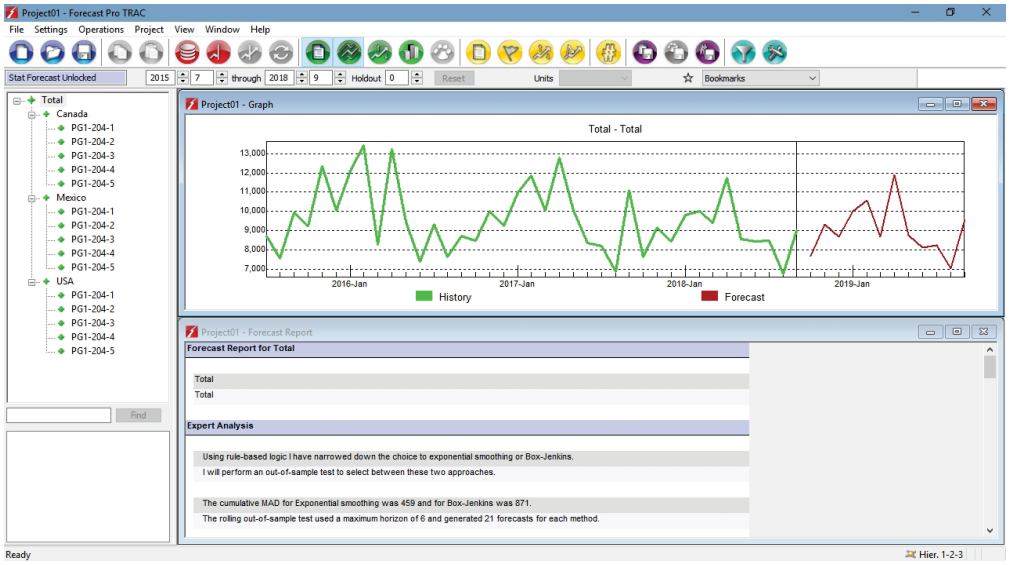

Notice that by reading in all three data files we've created a project that contains all three countries. If we wanted to recreate the forecasts from the individual component projects into a consolidated project, one way to do so would be to import the forecast modifiers and overrides from each component project into our current project using the Data Manager. This approach would work, but it would be a bit cumbersome.

The Super Project Manager in Forecast Pro is designed to streamline this type of consolidation. It allows you to create a new project which combines the component projects' data sources and then imports the component projects' modifiers and overrides. The end result is a "super project" that includes all of the forecasts created in the component projects.

We will now use the Super Project Manager to create a super project for our current example.

Close the current project by selecting **File>Close** and then open a new blank project by selecting **File>New**.

Select **Project>Super Project** to open the Super Project Manager.

The first step is to specify the component projects that you would like to combine. Click the Add button, select *Tutorial – Collaborating - USA*, *Tutorial – Collaborating - Canada* and *Tutorial – Collaborating - Mexico* and click Open.

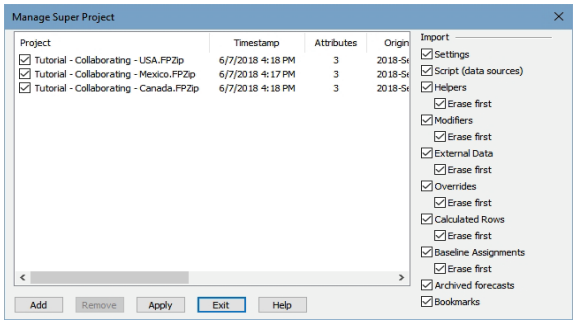

The Import options allow you to specify what you wish to import from the component projects. When first creating a super project you will generally import everything. When updating a super project you may want to pick and choose what to import.

The "Erase first" options are only relevant if you are updating a super project. If selected, they will erase any data corresponding to the specified component prior to loading in the new data.

Activate the checkbox to the left of each project. Make sure that your selections match the ones above and click Apply. The program will prompt you to save the super project. Name the project *Super Project Example* and click Save.

After the super project is created, open the Override view, fully expand the Navigator and select *Total>Canada>PG1-204-4*.

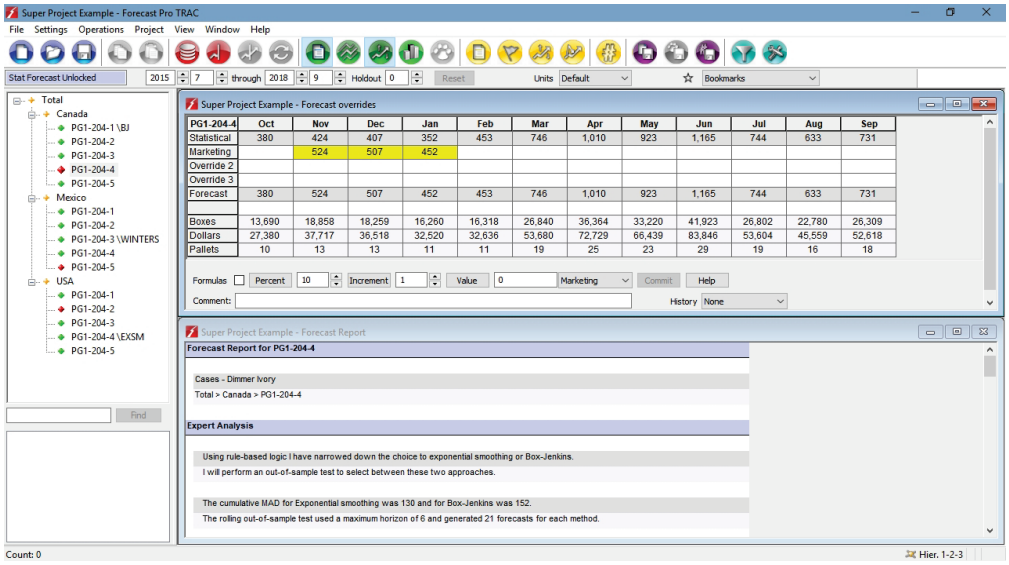

Notice that our new super project combines all three countries and includes the modifiers and forecast overrides from the component projects.

This exercise was designed to provide an introduction to super projects. Setting up and maintaining super projects is reasonably complex, so if you are planning on using super projects you should read the *Working with Super Projects* chapter of the Reference section of this manual carefully. Business Forecast Systems and Forecast Pro distributors also offer consulting services to help you establish your forecasting process.

Exit Forecast Pro.

This concludes the *Collaborating With Colleagues* lesson

Lesson 9

Tracking Forecast Accuracy

When you save a forecast project, Forecast Pro stores the statistical forecasts, the baseline forecasts and the final (i.e., adjusted) forecasts in the project's database. This means that if you have been updating your project and forecasts every planning period (for example, every month), that in addition to the current month's forecast, your project's database will contain forecasts that were generated in previous months. This record of previously generated forecasts is referred to as a *forecast archive*.

Maintaining a forecast archive allows you to track your forecast accuracy by comparing the forecasts that you generated previously to what actually happened. In this lesson we will explore Forecast Pro's tracking report and tracking graph views which are the primary vehicles for making this comparison.

Reading the Tracking Report

Start Forecast Pro, select **File>Open** and open the project, *Tutorial - Accuracy*. This project contains forecasts for the cakes and muffins data that we worked with in previous lessons.

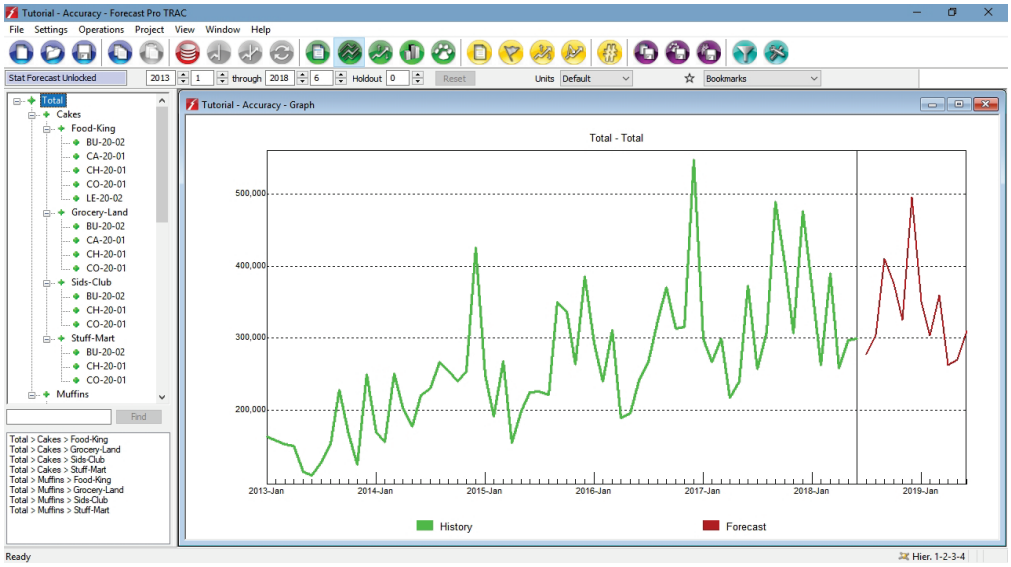

Notice that for the current forecast period, the historic data begins in January 2013 and ends in June 2018 and that the first forecast period is July 2018. Forecast Pro refers to the most recent historic data period used to generate a forecast as the *forecast origin*. For the current forecast, the forecast origin is June 2018.

In addition to the current period's forecast, this forecast project contains archived forecasts for the last 12 months.

Turn off the Graph view ( ). Click on the green View Tracking Report and Graph icon ( ) to open the tracking report and tracking graph. Your display should now look like the one below.

Close the Tracking Report Graph window by clicking the top right corner. Now, we can view the whole Tracking Report.

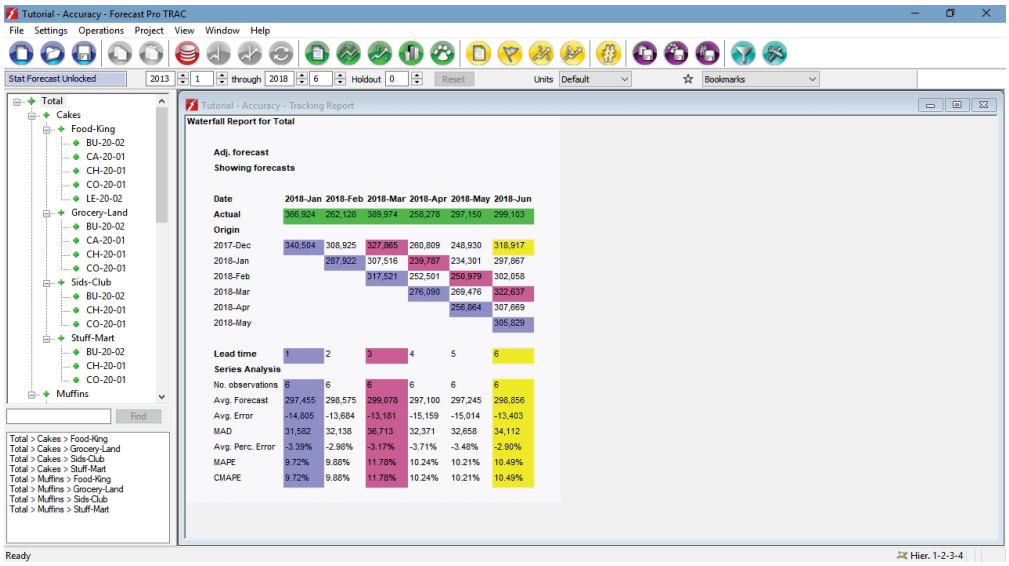

Due to its cascading-like appearance, the tracking report is sometimes referred to as a *waterfall* report. The report compares what we forecasted to what actually happened, therefore it is based on two key elements—the actual demand history and archived forecasts for the periods being analyzed.

The actual demand history for the most current six months is shown in the first row with the green shading. The next 6 rows display the forecasts generated for these periods from different forecast origins. Thus, the row labeled 2017-Dec displays the forecasts generated six months ago when the forecast origin was December 2017 and the first forecast period was January 2018. The row labeled 2018-May displays the forecast generated last month when the forecast origin was May 2018 and the first forecast period was June 2018.

The waterfall report allows you to highlight different *lead* times. A lead time refers to the number of periods ahead of the forecast origin the forecast was made for. Thus, a one-month-ahead forecast would have lead time equals 1, a two-month-ahead forecast would have lead time equals 2, etc.

The bottom portion of the report displays cumulative statistics for different *lead times*.

Notice that the forecasts for lead time equals 1 are all shaded in blue, the forecasts for lead time equals 3 are all shaded in maroon and the forecasts for lead time equals 6 are all shaded in yellow. Click on the View Tracking Report and Graph icon ( ) to restore the Tracking Report Graph. The same color coding is used on the Tracking Report Graph. You can control the lead times to color code using the Tracking Report Settings dialog box.

Customizing the Tracking Report

Right click on the tracking report to invoke the context menu and select **Tracking Report Settings**.

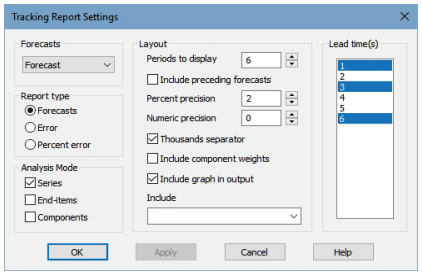

Notice that you can choose to view the statistical, baseline or adjusted forecasts. This allows you to determine if your adjustments are adding value.

You can also opt to view the forecasts themselves or the errors or the percent errors.

Select “Percent error” and click the Apply button. Notice that the display is updated accordingly. Click the “Include preceding forecasts” option and click Apply. This will display all archived forecasts that go into the cumulative statistics rather than just the subset in the triangular display.

The Analysis Mode setting controls what cumulative statistics to display for group-level data. The “Periods to display” controls the number of lead times to include in the report and the Lead time(s) box allows you to specify up to four lead times to color code.

Experiment with the report settings until you are comfortable with their operation. When you are finished, exit the program.

When reviewing tracking reports, you’ll often want to concentrate on items where the forecast accuracy has fallen outside of an acceptable range. Exception reports can be quite useful in this regard and is the subject of our next lesson.

This concludes the *Tracking Forecast Accuracy* lesson.

Lesson 10

Exception Reporting

Exception reports enable you to quickly find cases where your forecast error or some other performance metric has fallen outside of an acceptable range. Exception reporting reduces the need for manual review of your forecasts and allows you to focus on the items where human attention is most needed.

Forecast Pro provides a wide array of exception reports, allowing you to monitor the current history, forecasts and previously generated forecasts.

Working with Exception Reports

Start Forecast Pro, select **File>Open** and open the project, *Tutorial – Exception Reporting*. This project contains forecasts for the cakes and muffins data that we worked with in previous lessons.

Click the yellow View Exception Report icon ( ) to open the Exception Report window. Your display should now match the one below.

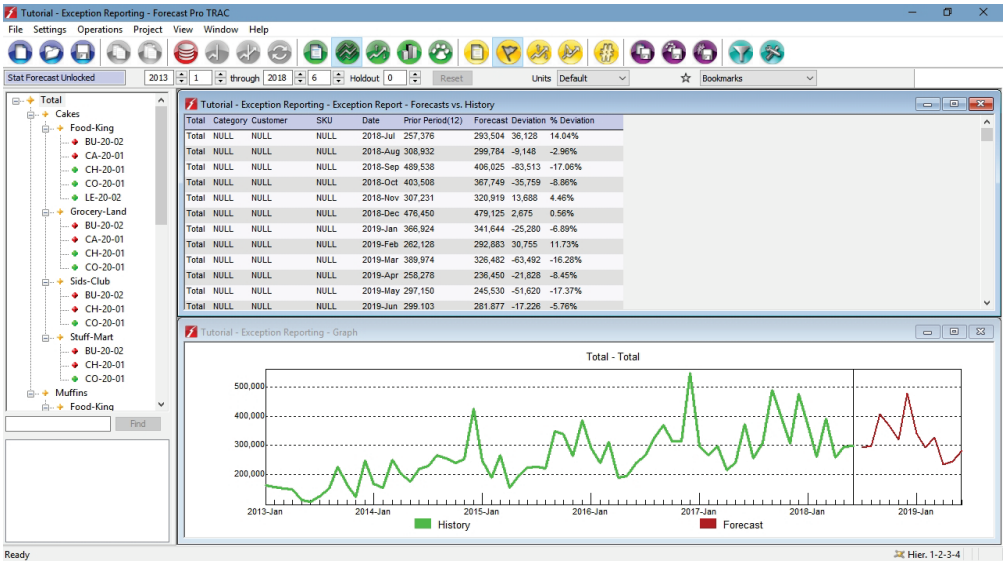

Notice that the window is labeled “Forecasts vs. History”. The default report compares the forecast for each period in the forecast horizon to the previous year’s historic value for the same period.

The first four columns of the report identify the item. The Date column identifies the forecast date. The next two columns (Prior Period(12) and Forecast) display the item’s historic value for 12 months prior to the date and the forecast value. The Deviation column shows the difference between the Forecast and the Prior Period(12) and the % Deviation column shows the difference as a percentage.

Right click on the exception report. Note that **All Items** is currently selected on the menu. We are also able to show only Hot List items or only the currently selected item on the Navigator. Select **Selected Item** and use the Navigator to change the selected item a few times to get a feel for how this works. Once you’re comfortable with the selected item mode of operation, right click on the exception report again to call up the context menu and select **All Items**.

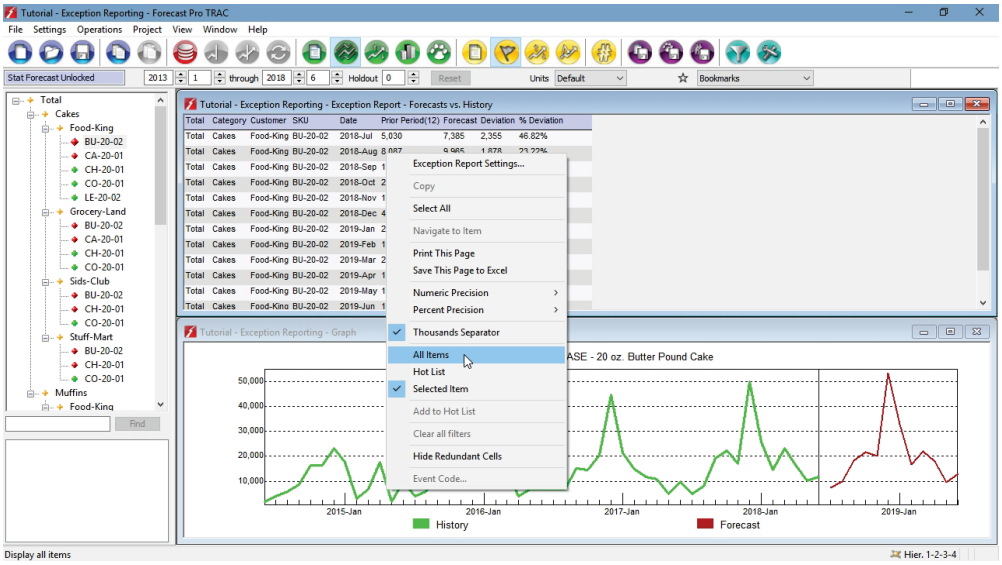

Let's sort the exception report by the percent deviations. Click the light blue Filter icon on the toolbar to activate the filtering and sorting mode. Click the drop down on the % Deviation column. Sort the display by checking both *Sort Largest to Smallest* and *Absolute Value* in the % Deviation filter drop down and clicking OK. Notice that the report is now sorted by the percent deviations.

Point to the first item on the exception report and double-click. Point to the second item on the report and double-click. Notice that when you double-click an item on the exception report, the Navigator jumps directly to that item. Thus, you can use the exception report as the equivalent of a Hot List to navigate through the listed items. All of the global report views (yellow-icon views) support this kind of navigation.

Let's now refine the report to only display items that have percent deviations greater than 30% or less than -30%. Click the drop down on the % Deviation column. Select *Numeric* on the drop down menu to bring up the Numeric Filter dialog box shown below.

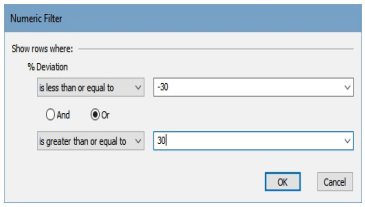

On the drop down below % Deviation (currently showing *equals*), select *is less than or equal to*. In the text box to the right enter "-30", select the *Or* radio button and on the bottom drop down select is *greater than or equal to*. In the text box to the right enter "30" and click Ok.

If you scroll to the bottom of the report you will see that the report shows NAs (Not Available) as well as items that have percent deviations below -30% or above 30%. Filter out NAs by unchecking "Not Available" in the % Deviation filter drop down and clicking OK.

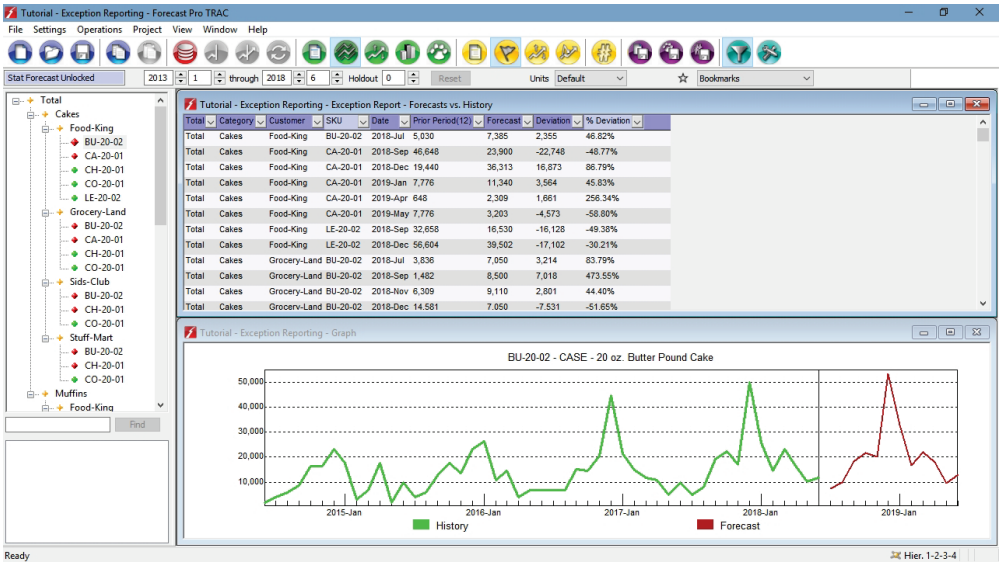

Now let's create a bookmark for this view. Click on the star to the left of the Bookmark drop down menu and select Add on the menu. The below dialog box should come up.

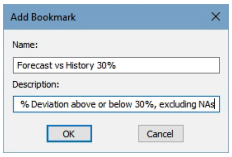

In the Name textbox, enter “Forecast vs. History 30%.” In the Description text box, enter "All items with % deviation above or below 30%, excluding NAs." Click OK to save the bookmark. We will use this bookmark to return to this display later in the lesson.

Exploring the Reporting Options

Right click on the exception report and select **Clear all filters** on the context menu.

Let’s take a closer look at the dialog box for the Forecast vs. History report. Right click on the exception report to invoke the context menu and select **Exception Report Settings**.

The *Comparison* section at the top of the dialog box is where you define the type of Forecast(s) you want to use (Baseline, Forecast, Statistical) and number of periods prior we want to compare to. By default, the report is showing the current forecast versus the same period last year for all periods in the forecast horizon. Uncheck the “Current forecast vs. same period last year” checkbox. Doing so allows you to customize the comparison periods.

The *Forecast periods to consider* section allows you to set the number of forecast periods to monitor and to specify if you want to look at each individual period or at the cumulative total across the specified forecast periods. The default setting (cumulative is not checked) monitors all forecast periods individually.

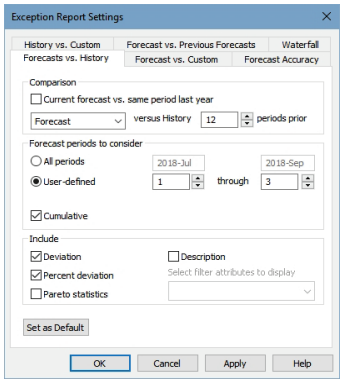

Let's compare the cumulative forecast for the next 3 months to the cumulative history for the corresponding 3 months last year. Select *User defined* and change the defined span to read 1through 3. Check the *Cumulative* checkbox.

The bottom section, *Include,* allows you to specify which statistics you want to show in the report. By default, the report shows Deviation and % Deviation. We will keep the defaults. Click the OK button.

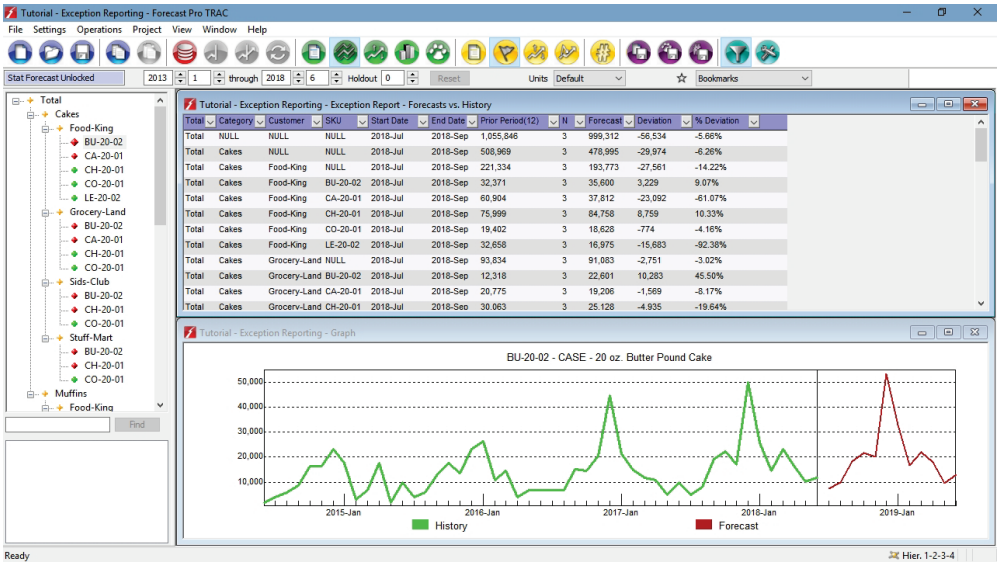

Notice that the report has updated to reflect the new settings. Columns are added to show the start date and the end date for the cumulative period. The newly added N column to the right of Prior Period(12) shows the number of prior periods that are included in the historical cumulative value. If historical

values are not available for all periods in the cumulative period, N will be the number of periods with available data.

Touring the Six Reports

Forecast Pro includes six different exception reports. Let's take a quick look at each of these using a set of bookmarks in our project.

Forecast vs. History

Open the bookmark drop down and select *R1:Forecast vs. History*. Double-click on the first entry on the report to jump to it. Right click on the exception report to invoke the context menu and select **Exception Report Settings**. Your screen should look like the one below.

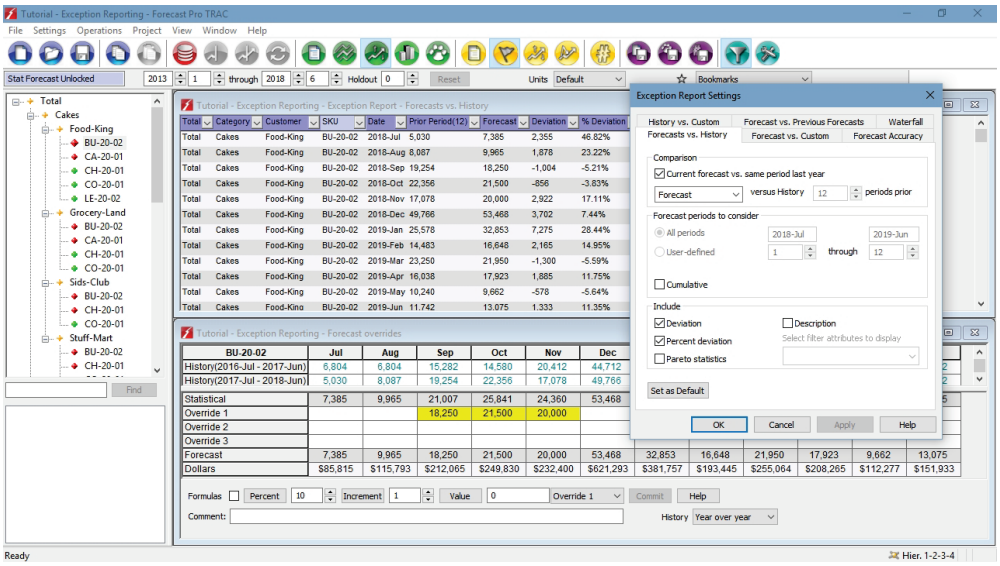

The Forecast vs. History report compares forecast values to prior historical values. We explored this report in detail earlier in the lesson.

The current report compares the forecast to the historic value 12 months prior. Notice that the values listed in the report for Forecast and Prior Period(12) match their counterparts in the override grid. Click OK to exit the Exception Report Settings dialog box.

Forecast vs. Custom

Open the bookmark drop down and select *R2:Forecast vs. Custom*. Double-click on the first entry on the report to jump to it. Right click on the exception report to invoke the context menu and select **Exception Report Settings**. Your screen should look like the one below.

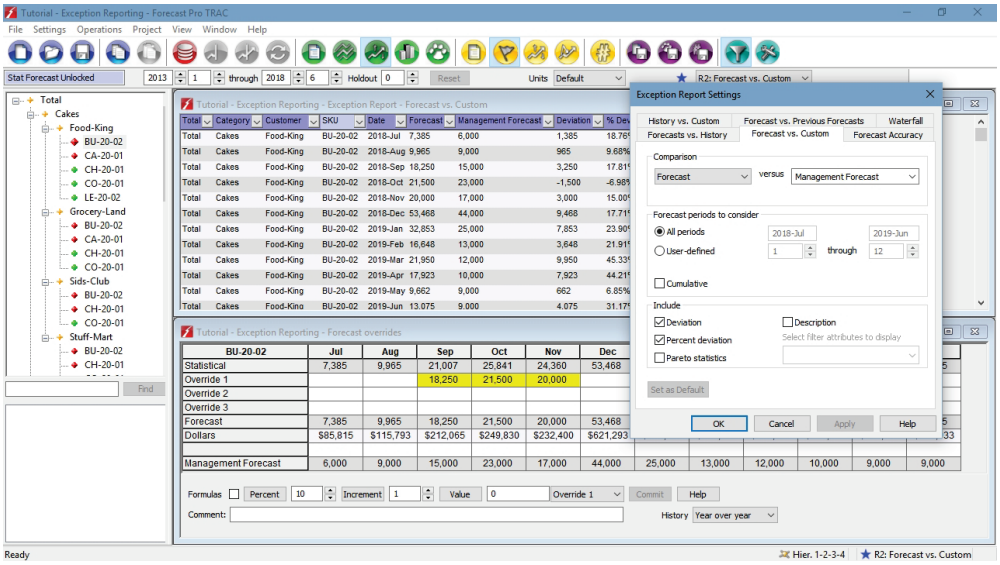

The **Forecast vs. Custom** report compares the current forecast to any row(s) available in the override grid. The current report compares the current forecast to the Management Forecast (an external row). Notice that the values listed in the report for Forecast and Management Forecast match their counterparts in the override grid.

Experiment with the report settings until you are comfortable with their operation. Click OK to exit the Exception Report Settings dialog box when you have finished.

Forecast Accuracy

Open the bookmark drop down and select *R3:Forecast Accuracy*. Double-click on the first entry on the report to jump to it. Right click on the exception report to invoke the context menu and select **Exception Report Settings**. Your screen should look like the one below.

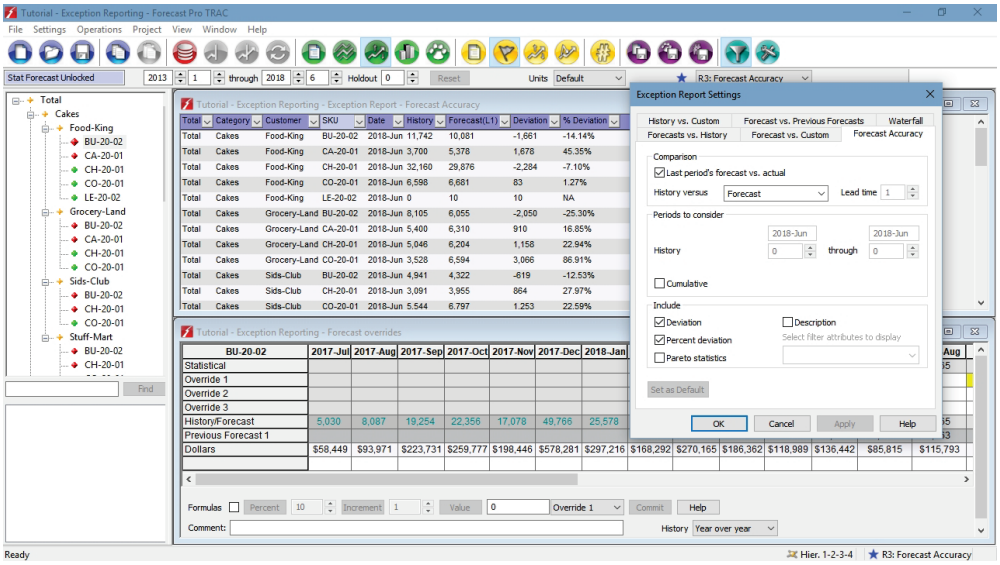

The **Forecast Accuracy** report looks at previously generated forecasts vs. what actually happened. The current report compares the last forecast we made for June 2018 against what happened. Notice that the values listed in the report for Forecast(L1) and History match their counterparts in the override grid.

Experiment with the report settings until you are comfortable with their operation. Click OK to exit the Exception Report Settings dialog box when you have finished.

History vs. Custom

Open the bookmark drop down and select *R4:History vs. Custom*. Double-click on the first entry on the report to jump to it. Right click on the exception report to invoke the context menu and select **Exception Report Settings**. Your screen should look like the one below.

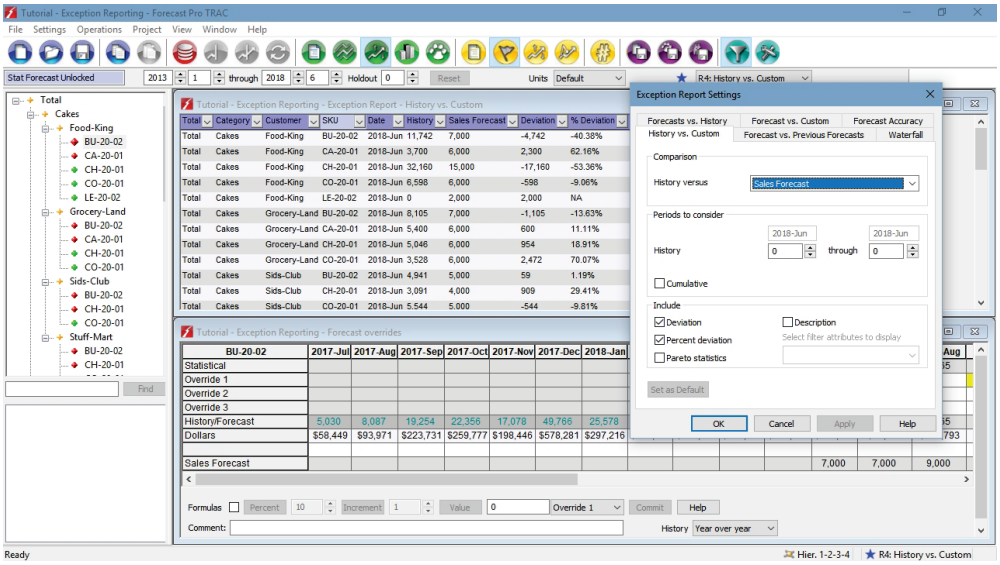

The **History vs. Custom** report compares historical values with any row(s) available in the override grid. The current report compares what happened in June 2018 vs. what the Sales Forecast was (an external variable). Notice that the values listed in the report for History and Sales Forecast match their counterparts in the override grid.

Experiment with the report settings until you are comfortable with their operation. Click OK to exit the Exception Report Settings dialog box when you have finished.

Forecast vs. Previous Forecast

Open the bookmark drop down and select *R5:Forecast vs. Previous Forecast*. Double-click on the first entry on the report to jump to it. Right click on the exception report to invoke the context menu and select **Exception Report Settings**. Your screen should look like the one below.

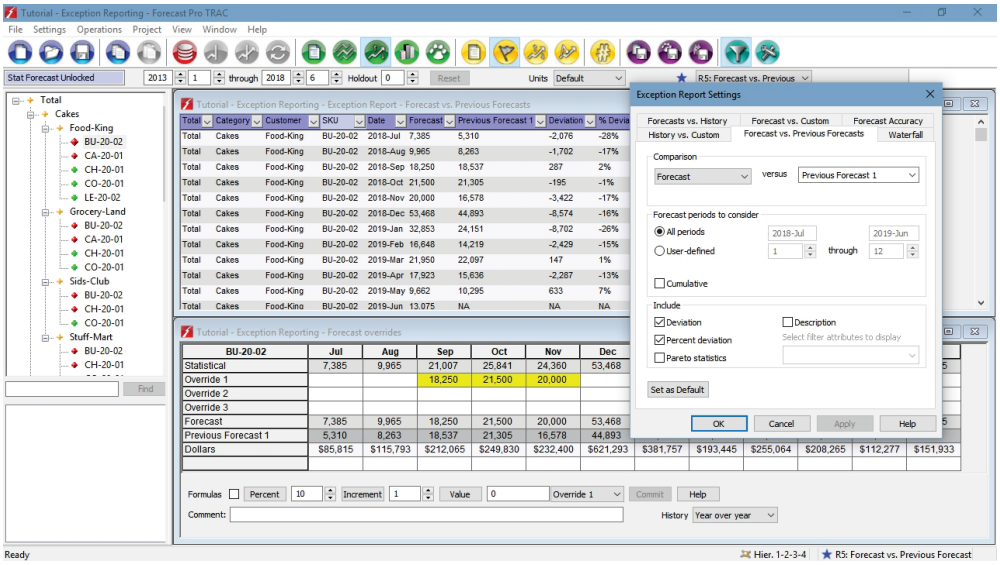

The **Forecasts vs. Previous Forecasts** report compares the current forecast to previously generated forecasts. The current report compares the current forecast against the forecast we made last month. Notice that the values listed in the report for Forecast and Previous Forecast 1 match their counterparts in the override grid.

Experiment with the report settings until you are comfortable with their operation. Click OK to exit the Exception Report Settings dialog box when you have finished.

Waterfall

Open the bookmark drop down and select *R6:Waterfall*. Double-click on the first entry on the report to jump to it. Right click on the exception report to invoke the context menu and select **Exception Report Settings**. Your screen should look like the one below.

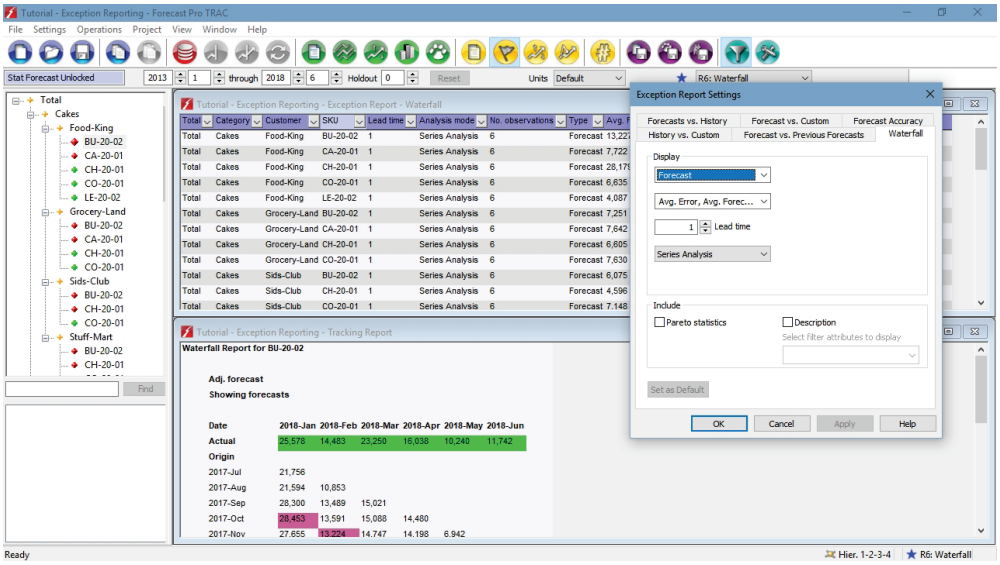

The **Waterfall** report provides a complete summary of forecast accuracy across lead times. It doesn't compare items and calculate deviations like the other exception reports. It simply allows you to list statistics from the waterfall reports in a global report that can be sorted and filtered to aid in your review process.

Experiment with the report settings until you are comfortable with their operation. Click OK to exit the Exception Report Settings dialog box when you have finished and then exit the program without saving changes to the *Tutorial – Exception Reporting* project.

This concludes the *Exception Reporting* lesson.

Lesson 11

Pareto Analysis

Pareto analysis is used to assign ABC classifications to your forecast items. Typically, important high volume items are categorized as type “A”, medium volume items are categorized as type “B” and slow moving items are categorized as type “C”. Thus, the code indicates the relative importance of an item and many organizations adopt different procedures for creating, reviewing and monitoring their forecasts based upon the codes.

Forecast Pro can automatically assign items into ABC categories based upon their volume. Alternatively, you can create a custom filter to define the classifications. In this lesson, we will explore both approaches and also illustrate how sorting and filtering reports on the Pareto output can be extremely useful.

Defining Pareto Classifications Within Forecast Pro

Start Forecast Pro, select **File>Open** and open the project, *Tutorial – Pareto*. This project contains forecasts for the cakes and muffins data that we worked with in previous lessons.

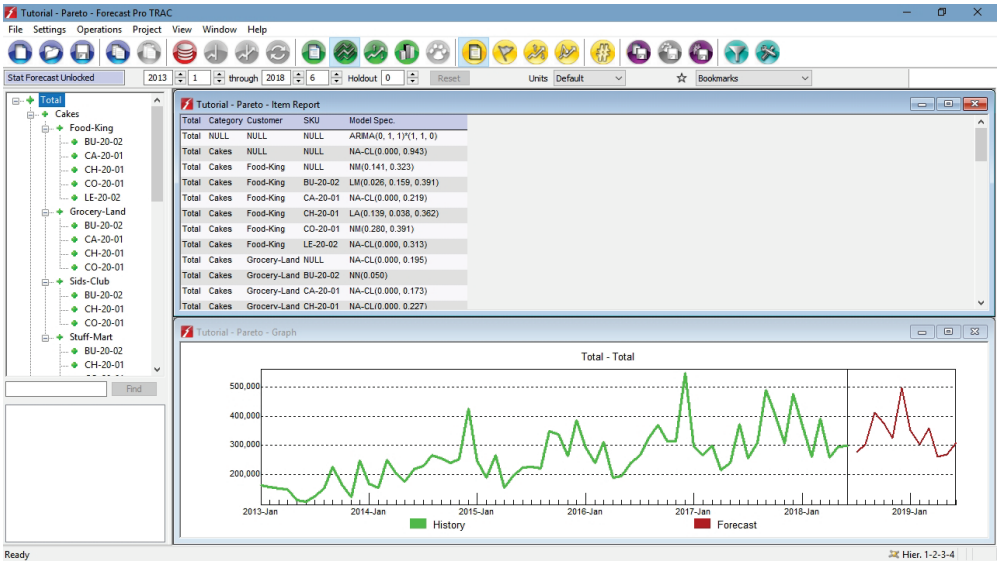

Click the light blue Options icon (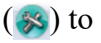) to invoke the Options dialog box and then select the Pareto Analysis tab.

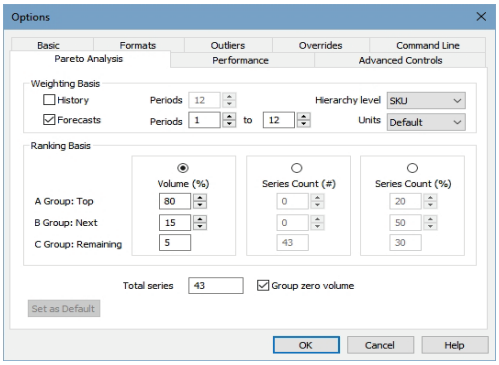

The Pareto Analysis tab allows you to specify how the ABC codes should be assigned.

The Weighting Basis section allows you to indicate whether the ranking should be performed using the history, forecasts or a combination of the two and allows you to set the span for your selection. The codes are assigned to a single level of the hierarchy. The Hierarchy level drop down allows you to pick the level to use.

The Ranking Basis section allows you to set the thresholds for the classifications.

The Volume (%) option allows you to set the codes based upon the cumulative volume of each group. For example, if A is set to 80%, B is set to 15% and C is set to 5%, Forecast Pro will first rank all of the items from highest volume to lowest volume and then sequentially place the top ranked items into the A group until their cumulative volume equals or exceeds 80% of the total volume. It will then sequentially place items into the B group until the cumulative volume of the A and B items equals or exceeds 95% (80% + 15%). The remaining items (5% or less of the cumulative volume) are assigned type C.

The Series Count (#) option allows you to assign specific numbers of items to each group. For example, if you had 1,000 items and you set the A group to 100, the B group to 200 and the C group to 700—the 100 highest volume items would be assigned to group A, the next highest 200 items would be assigned to group B and the remaining 700 would be assigned to group C.

The Series Count (%) assigns the specified percentage of the total number of items you have to each group. For example if A is set to 20%, B is set to 50% and C is set to 30%, Forecast Pro will first rank all of the items from highest volume to lowest volume and then sequentially place the top ranked items into the A group until the number of A group items equals 20% of the total number of items. It will then sequentially place items into the B group until the number of B group items equals 50% of the total number of items. The remaining items (30% of the total number of series) are assigned type C.

If selected, the Group zero volume option will place any items with zero volume into group “D”.

Make sure that your settings match the ones above and click OK.

Sorting and Filtering the Pareto Output

Right click on the Item Report view to display its context menu and select **Item Report Design**.

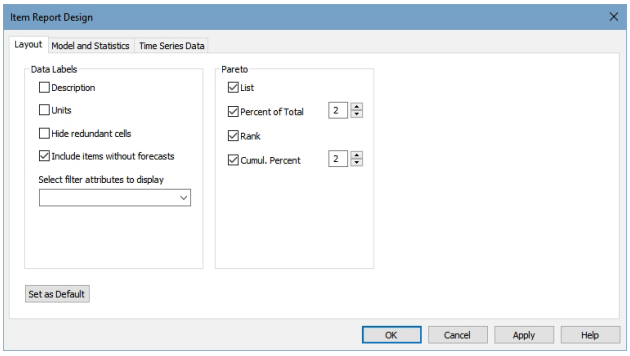

Click the check boxes for the four Pareto fields (List, Percent of Total, Rank, Cumul. Percent) and click OK.

Notice that the Item report now includes the four fields associated with the Pareto analysis—before we discuss these fields let's sort the report to make it easier to read.

Click the light blue Filter icon to activate the filtering and sorting mode. Click on the filter drop down for Ranking, select *Sort Smallest to Largest* and click OK. Your screen should now look like the one below.

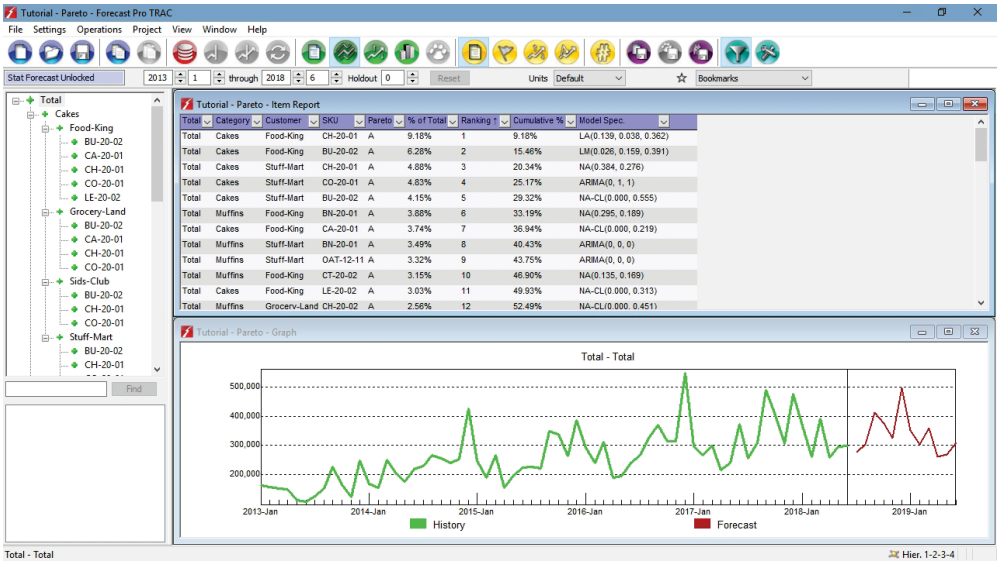

The Pareto field displays the assigned code. The % of Total field displays the percentage of total volume that the item represents. The Ranking field is self-

explanatory. The Cumulative % field displays the percentage of total volume for the current item plus all higher ranked items.

Scroll through the display and examine the output. Notice that the Cumulative % for the item ranked 25 is the first one to exceed 80% and that therefore the item ranked 26 is classified as a B.

Click on the filter drop down for the Pareto column, set the filter to only display type A and click OK. Examine the filtered report. Double-click on the top ranked item on the report and then double-click on the item ranked second. Notice that the Navigator immediately jumps to the selected item. The ability to sort, filter and navigate by double-clicking on items can be very convenient when reviewing your forecasts.

If you shuffle the hierarchy the Pareto information will automatically update.

Experiment with the Pareto and filtering options until you are comfortable with their operation. When you are finished, exit the program without saving changes to the *Tutorial - Pareto* project.

Defining Pareto Classifications via a Custom Filter

Forecast Pro allows you to import custom filter fields into your project via the Definitions row on the Data Manager. This can be very useful when you wish to filter reports based on attributes that are not part of your forecasting hierarchy. In this exercise we will display custom Pareto codes that were defined as a custom filter.

Start Forecast Pro, select **File>Open** and open the project, *Tutorial - Pareto*.

Right click on the Item Report view to display its context menu and select **Item Report Design**.

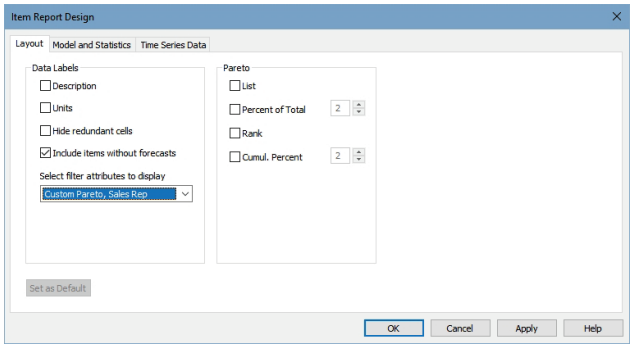

Open the *Select filter attributes to display* drop down. Notice that there are two custom filters available—Custom Pareto and Sales Rep. Click the checkboxes for both custom filters and click OK.

Notice that the Item report now includes the two filter fields.

Click the light blue Filter icon to activate the filtering and sorting mode. Click on the filter drop down for filter-Custom Pareto, select Sort A to Z and click OK. Your screen should now look like the one below.

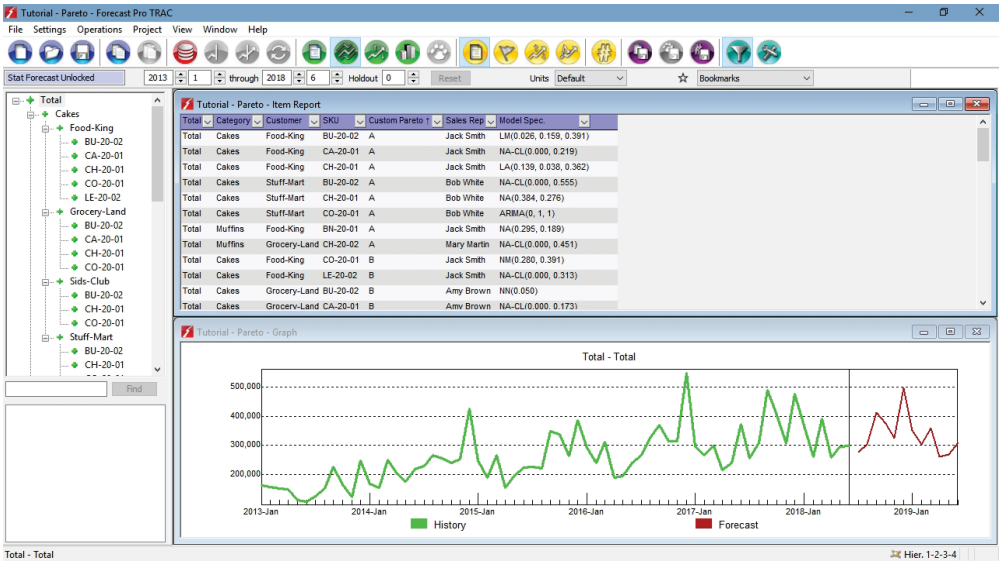

Scroll through the report. Notice that the custom filter defined codes for A, B, C, D and E rather than just ABC. Using a custom filter for your Pareto codes allows you to specify as many categories as you wish and assign membership in any fashion you desire.

Experiment with the filtering and sorting options until you are comfortable with their operation. When you are finished, exit the program without saving changes to the *Tutorial - Pareto* project.

This concludes the *Pareto Analysis* lesson.

Lesson 12

Operating From the Command Line

When the command line to execute the program includes a project file (FPZip or FPProj) and the \b parameter, Forecast Pro will read the project, create the forecasts, save all output files and then exit. This feature is particularly useful when you are integrating forecasting with other software systems and desire a "hands-off" approach.

If you are using a Collaborator license, this lesson requires a Forecaster license.

Command-Line Operation

In this lesson we will first use Forecast Pro interactively to create a project file. We will then exit the program and create forecasts for the project via the command line.

Start Forecast Pro and select **Settings>Options**. Set the forecast horizon to 12 and then click on the Command Line tab.

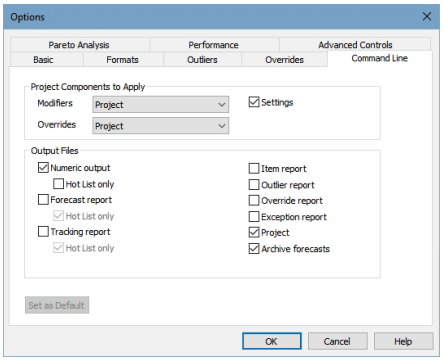

The *Project Components to Apply* section allows you to control how the forecasts are generated when the project is run from the command line.

The *Modifiers* drop down allows you to specify how modifiers are applied in the command-line mode. Click on the drop down, and you will see that there are four options: None, Project (default), Data Source and Project and Data Source. If None is selected, expert selection will be used for all items. If Project is selected, Forecast Pro will apply modifiers as specified in the Project's Navigator. If Data Source is selected, Forecast Pro will consider only modifiers specified in the Data Manager. If Project and Data Source is specified, Forecast Pro will apply modifiers specified in the Navigator for items that do not have a modifier in the Data Source and apply Data Source modifiers whenever specified.

The *Overrides* drop down has the same options and functionality as the *Modifiers* drop down, described above.

If *Settings* is selected, the command-line run will use all of the settings associated with the project (e.g., all settings in the Settings dialog box, numeric output formatting, etc.). If this option is not selected, the default settings (i.e., the settings that have been saved as the defaults for use with new projects) will be used when running from the command line.

The *Output Files* section allows you to specify the output files to generate when the project is run from the command line. The format and content of these files will match the current project's settings.

Make sure your settings match the ones above. Click the OK button to save your selections.

Click the red Data Manager icon ( ) to call up the Data Manager.

Click the Add drop down on the Historic Data row, select Excel and select *123 Bakery – Historic Data.xlsx*. Click the Exit button to exit the data manager. A warning will be displayed stating that "one or more data sources have not been applied", click OK to exit the Data Manager without reading the data.

Select **File>Save** to save the project. Name the project "Test". Exit Forecast Pro.

We will now run the project from the command line.

You can execute a command line using the Search box on the Windows Start Menu or using the Windows **Run** command. Execute the program including Test.FPZip as a parameter followed by " /b". If you installed Forecast Pro into a directory named "C:\Program Files\Forecast Pro v5" the command line should read (including the quotation marks):

"C:\Program Files\Forecast Pro v5\ForecastProTRAC.exe" Test.FPZip /b

After issuing the command, the program will read in the historic data, calculate the forecasts and create all of the requested output files. Examine your output directory to confirm the files were created.

For more detailed information regarding command-line operation consult the Reference section of this manual

3

TUTORIAL: Forecasting Operations

Lesson 13

Using Forecast Modifiers

In the very first lesson you prepared forecasts using expert selection—Forecast Pro's default model selection procedure. In this lesson you will use forecast modifiers to dictate the forecasting models to be used for specific items.

Specifying Forecast Modifiers for Individual Items

Start the program and click the red Data Manager icon ( ) to call up the Data Manager. Click the Add drop down on the Historic Data row, select Excel and select *ACME Hardware – Historic Data.xlsx* to add it to the Historic Data row.

Click the Read & Forecast button to read in the historic demand data and generate the statistical forecasts.

Click the Exit button to exit the Data Manager.

Fully expand the Navigator using the **Expand All** option on its context menu and open the Graph view by clicking its green icon ( ). Select *Total>USA>PG1-204-2*. Your display should now match the one shown below.

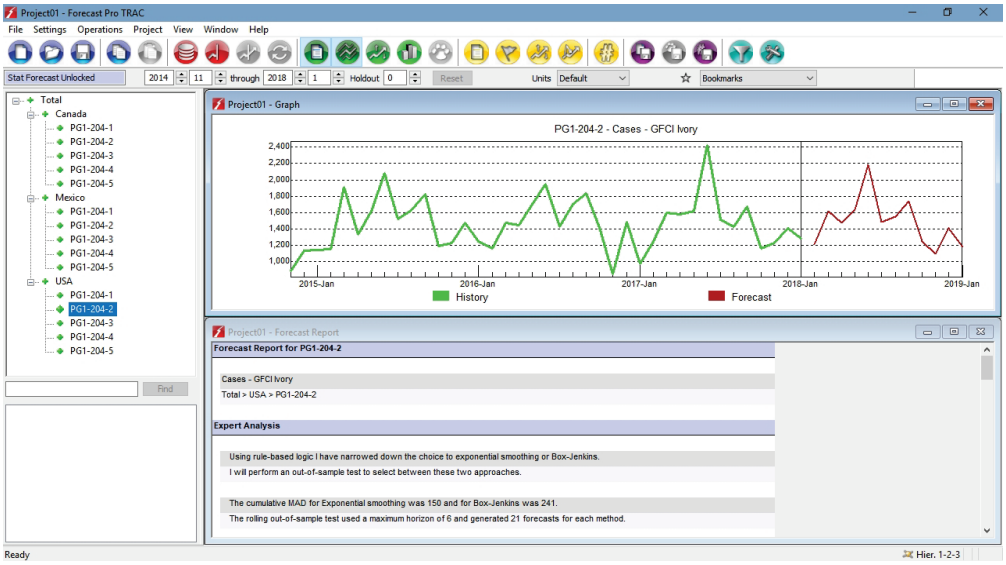

If you do not instruct Forecast Pro to use a specific forecasting technique it will use expert selection to select the appropriate technique for each item forecasted. If you examine *Total>USA>PG1-204-2*'s Forecast Report, you'll notice that it includes an Expert Analysis section, and that expert selection chose to forecast this item using a nontrended, seasonal exponential smoothing model.

Suppose we want to use a Winters exponential smoothing model for this item. Winters is a form of exponential smoothing which includes both a trend and a seasonal component.

To specify the model, right click on *Total>USA>PG1-204-2* to bring up the Navigator's context menu and then select **Model>Exponential Smoothing>Winters**.

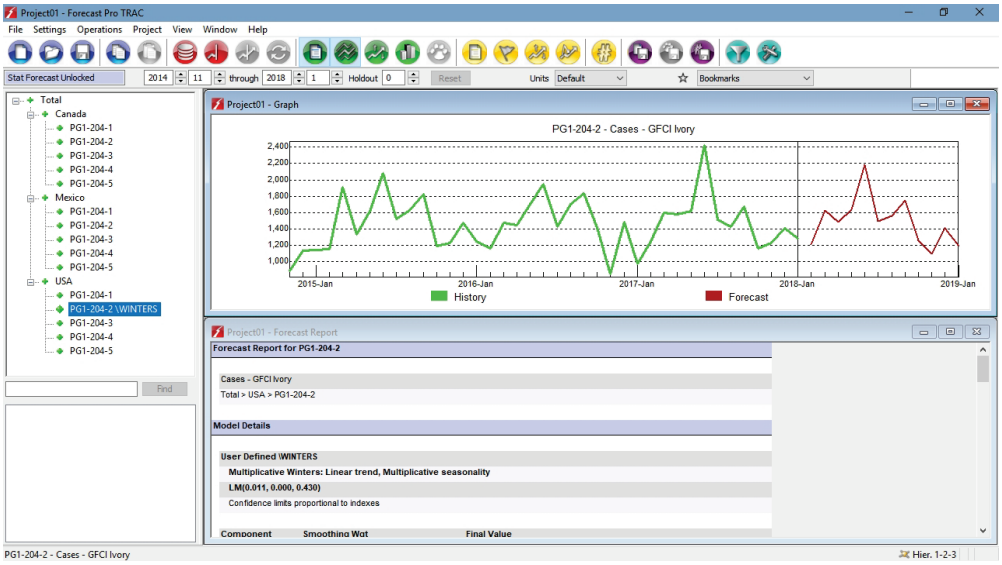

Notice that \WINTERS now appears next to *Total>USA>PG1-204-2* on the Navigator. \WINTERS is a forecast modifier. In this instance, it is used to indicate that you have opted to use a Winters exponential smoothing model for this item rather than expert selection.

Examine the Forecast Report to verify that the Winters model was used. You'll notice that the Expert Analysis section is no longer present and that the Model Details section indicates that a "User Defined \WINTERS" model was built.

Select *Total>USA>PG1-204-4* on the Navigator. Right click to call up the context menu and select **Model>Very Simple Models>Same as Last Year>Plus Percent**. Specify the percentage as "10" and click OK.

Notice that \SALYP=10 now appears next to *Total>USA>PG1-204-4* on the Navigator. This model will set the forecasts to equal the same values as the preceding year plus 10%.

Specifying Forecast Modifiers for Multiple Items

Select the group *Total>Canada* on the Navigator. Right click to call up the context menu and select **Model>Box-Jenkins>Auto**. Notice that \BJ now appears to the right of the group *Total>Canada*.

Suppose we wish to dictate that a Box-Jenkins model be used for every item in the group Canada. To accomplish this, select **Apply Modifier(s) To>Children** on the Navigator's context menu. Notice that the \BJ modifier now appears on all of *Total>Canada*'s "children" (i.e., all items one level down in the group *Total>Canada*).

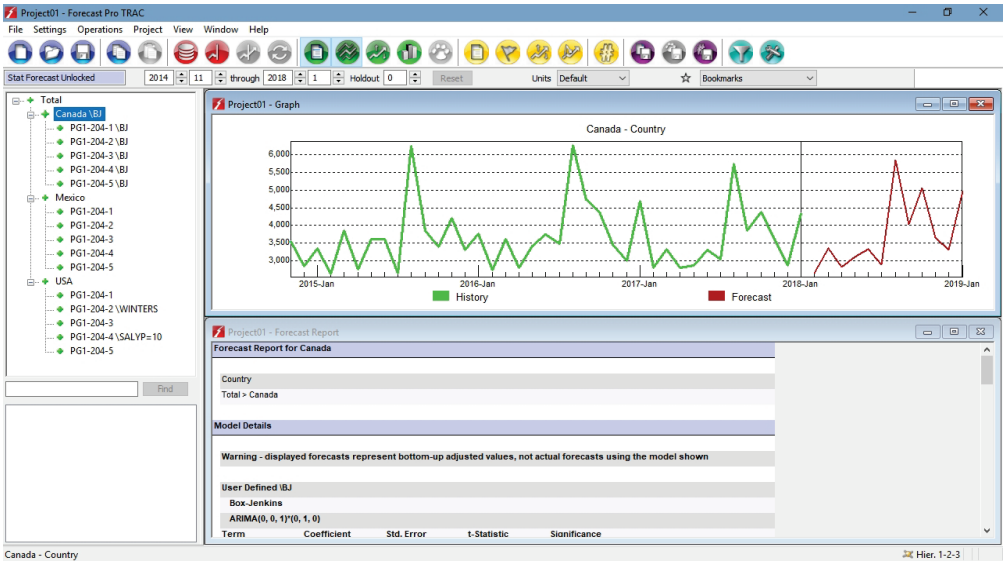

Specifying Custom Forecast Modifiers

In addition to the commonly used forecasting models found on the Model context menu, Forecast Pro supports a broad range of custom modeling modifiers. These modifiers accommodate a wide range of modeling options including confidence limit percentiles, safety stock lead times and power transformations.

Custom modifiers are entered using the **Custom Modifier** option on the Navigator's context menu.

Consult the *Using Forecast Modifiers and Model Specification Dialog Boxes* chapter of the Reference section or go to the help system in Forecast Pro to see a full list of the available modifiers.

Experiment with the various modeling options until you are comfortable with their operation. When you are finished, exit the program.

This concludes the *Using Forecast Modifiers* lesson.

Lesson 14

Building Event Models

Event adjustment models extend exponential smoothing by allowing you to adjust for events like sales promotions, strikes or simply for unexplained outliers. You can adjust for events of several different types. These could be promotions of different types or sizes, or different calendar effects like Easter and Independence Day weeks. Forecast Pro knows these occurrences simply as events of types 1, 2, etc.

Event adjustment models work almost the same as seasonal index models. In a seasonal index model, each month gets its own index, which is updated each time that month recurs. In an event adjustment model, each event type gets *its* own index, which is updated each time an event of that particular type recurs. The difference is that while January recurs every 12 months, an event of type 1 usually recurs irregularly.

Since Forecast Pro knows that January occurs every 12 months, you don't have to provide that information. To implement event adjustment, however, you must construct an event schedule which indicates where the events occurred historically, and if applicable, where they will occur during the forecast period. This is accomplished either interactively using the Event Manager or you can create the event schedule outside of Forecast Pro (for example in Excel) and import it using the data manager.

For more statistical details, please consult the statistical reference for exponential smoothing in the *Forecast Pro Statistical Reference Manual*.

In this lesson you will use an event models to capture the relationship between sales and promotions. You will also use an event model to capture seasonal patterns in weekly data.

Modeling a One-time Promotion

Start the program and click the red Data Manager icon ( ) to call up the Data Manager. Click the Add drop down on the Historic Data row, select Excel and select *Giveaway.xls* to add it to the Historic Data row.

Click the Read & Forecast button to read in the historic demand data and generate the statistical forecasts.

Click the Exit button to exit the Data Manager.

Open the Graph view by clicking its green icon ( ).Your display should now match the one shown below.

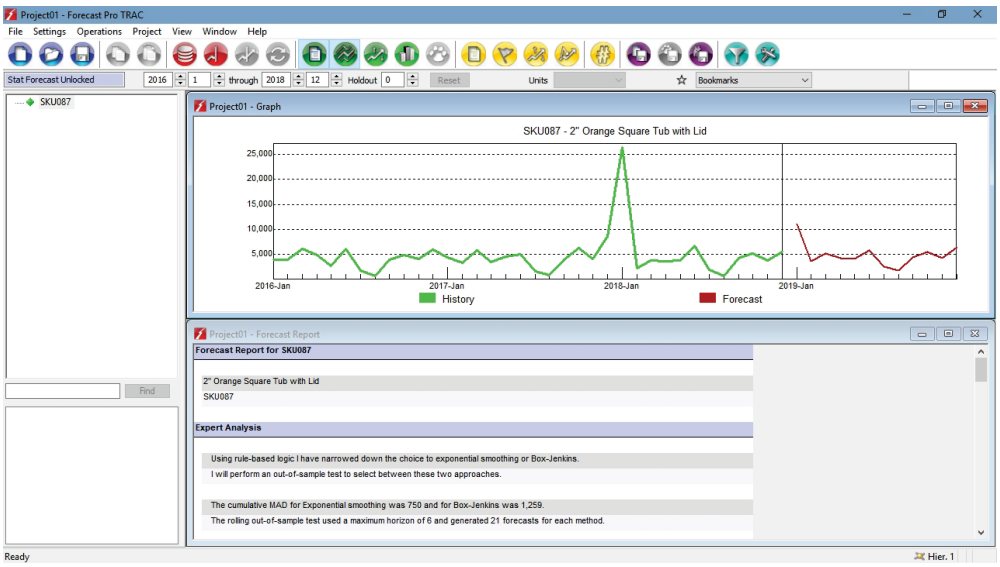

SKU087 represents monthly sales of a plastic storage container that is sold along with other storage products via in-home sales parties. The most notable feature in the data is the large peak in January of 2018.

One of the marketing strategies that this firm uses is to provide everyone who attends an in-home party with a free gift. In January of 2018, SKU087 was the free gift. Thus the volume associated with this period represents not only sales of the SKU but also the number provided free of charge.

By default, Forecast Pro modeled SKU087 using expert selection. Include the fitted values on the graph. (This is accomplished by toggling on **Fitted** on the graph's context menu).

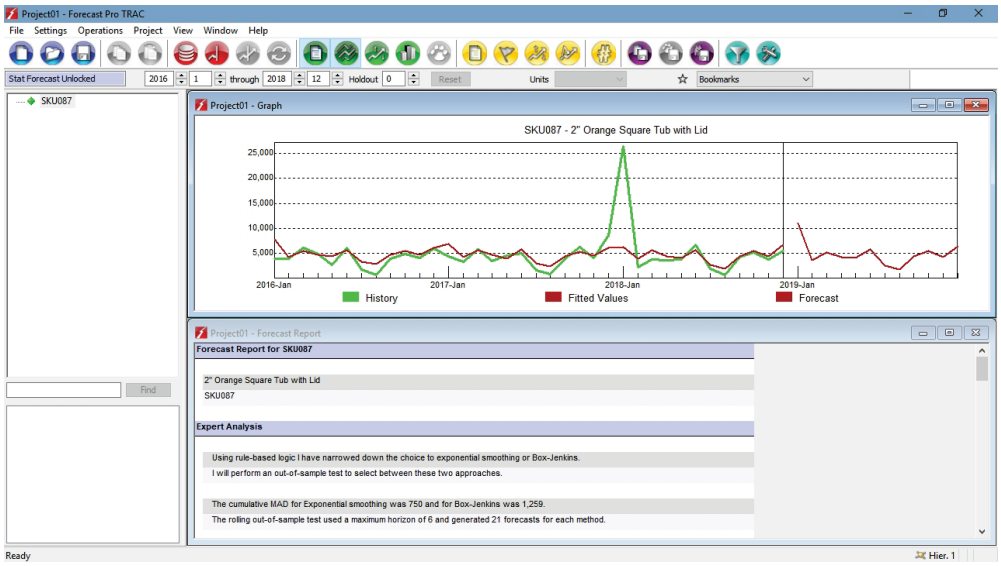

Notice that the model completely misses the January 2018 peak. This is because we are fitting this point based solely on the past history and a demand peak of this magnitude is unprecedented. Notice also that the forecast for January 2019 is substantially higher than the January peaks in 2016 and 2017. This is because the January 2018 peak has a substantial impact on the January 2018 forecast. Since we know that the January 2018 peak is due to a promotion, this is not an appropriate forecast.

We will now build an event model for SKU087. Right click on the graph to bring up the context menu and select **Manage Events**. This invokes the Manage Events dialog box, which is currently empty. The first step in building event models is to define the event types you'd like to use.

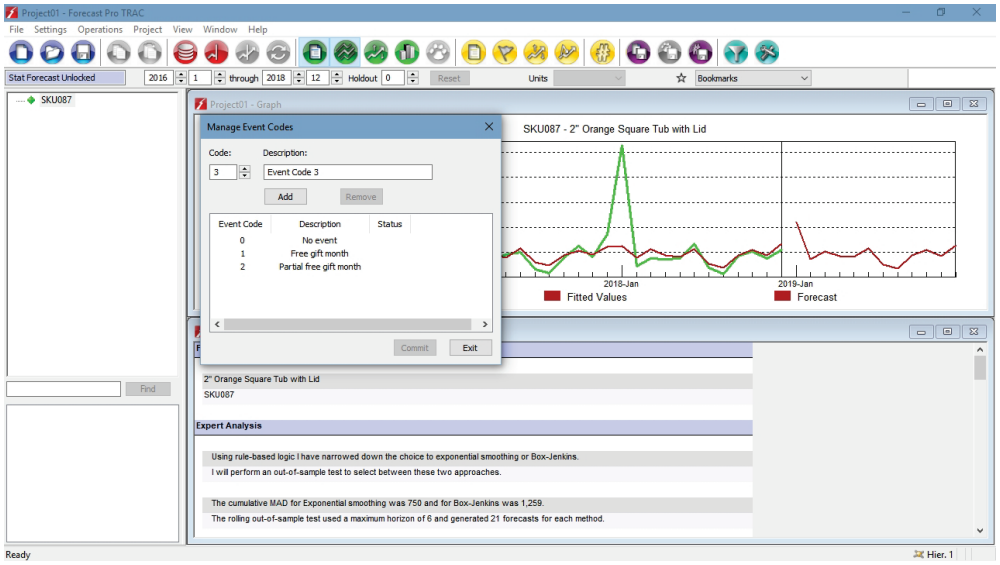

Click the Edit button to open the Manage Event Codes dialog box. The event code 0 is always used to indicate no event occurred. Therefore, it is predefined and *No Event* already appears in the bottom portion of the dialog box as a defined event type. Edit the description displayed for Event Code 1 to read *Free gift month* and click the Add button. Notice that the event type is added to the list with a status of "Pending" and that the dialog box is now ready for you to define Event Code 2. Edit the description displayed for Event Code 2 to read *Partial free gift month* and click the Add button. Your screen should now match the one shown above.

Click Commit to accept the new event codes and then Exit to return to the Manage Event dialog box. We are now ready to specify where the events occurred.

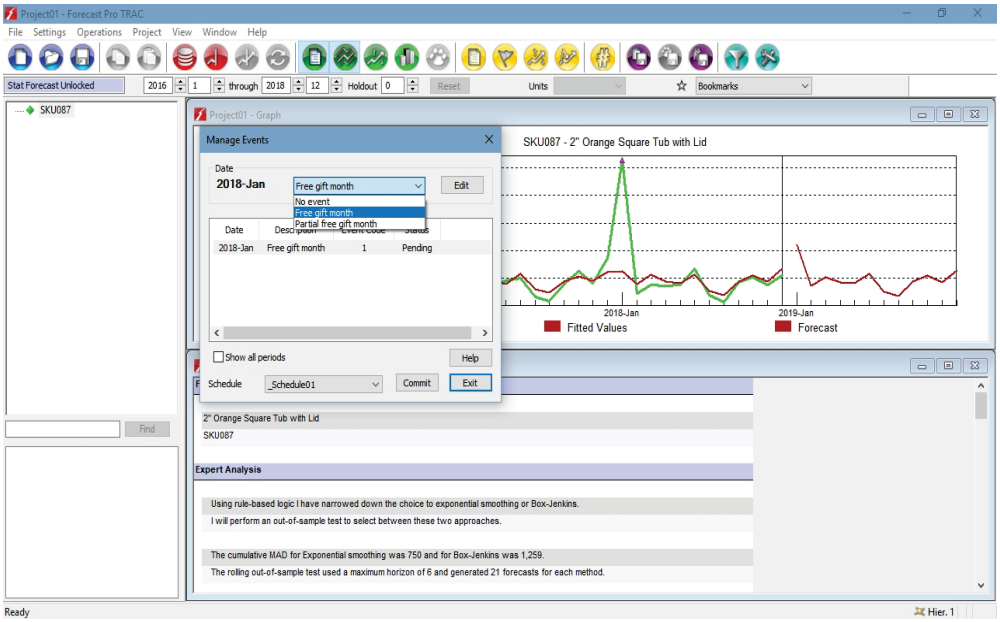

Click on January 2018 on the graph. Notice that the Manage Events dialog box displays the selected date and that a marker also appears on the graph to indicate that the period has been selected. Use the drop down control to assign the *Free gift month* event to the selection. Notice that the assignment now appears in the lower section with a status of "Pending." Click the Commit button to accept the assignment and build the event model.

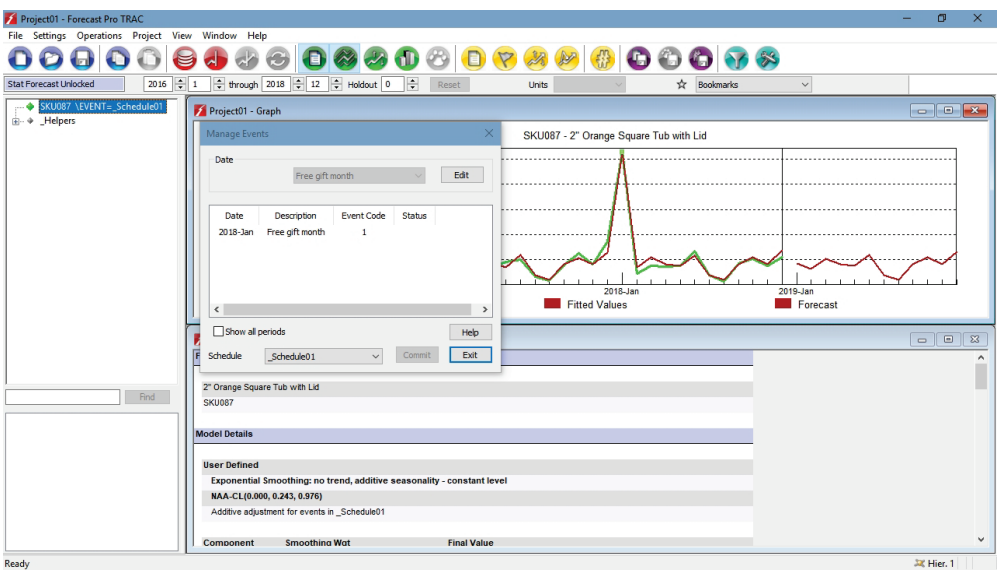

Notice the dramatic improvement in the fit and forecast. If you look closely at the graph, you might notice that the historic value for December 2017 is the second largest in the data set. This is due to the fact that SKU087 was the free gift for part of December 2017.

Click December 2017 on the graph and use the drop down to assign the period to the *Partial free gift month* event code. Click the Commit button to update the model.

This was a very simple example, where we only had two promoted periods during the history. Because the periods were promoted differently (one was a partially promoted month and the other was a fully promoted month) we used two event codes to capture the responses. In our next example we will build an event model for data with a recurring promotion.

Examine the model output and experiment with the event manager until you are comfortable with its operation. When you are finished, exit the program.

Modeling a Recurring Promotion

In this section we will be forecasting monthly sales of a nationally advertised brand of mouthwash.

Start the program and click the red Data Manager icon ( ) to call up the Data Manager. Click the Add drop down on the Historic Data row, select Excel and select *Minty Fresh Mouthwash – Historic Data.xls* to add it to the Historic Data row.

Click the Add drop down on the Helpers row, select Excel and select *Minty Fresh Mouthwash - Helpers.xls* to add it to the Helpers row. This file contains an event schedule indicating the months where the mouthwash was price promoted.

Click the Read & Forecast button to read in the historic demand data and generate the statistical forecasts.

Click the Refresh button on the Helpers row to read in the event schedule.

Click the Exit button to exit the Data Manager.

Open the Graph view by clicking its green icon ( ). Right click on the Navigator to invoke its context menu and select **Expand all** to fully expand the Navigator. Select *MintyFresh24* on the Navigator. Your display should now match the one shown below.

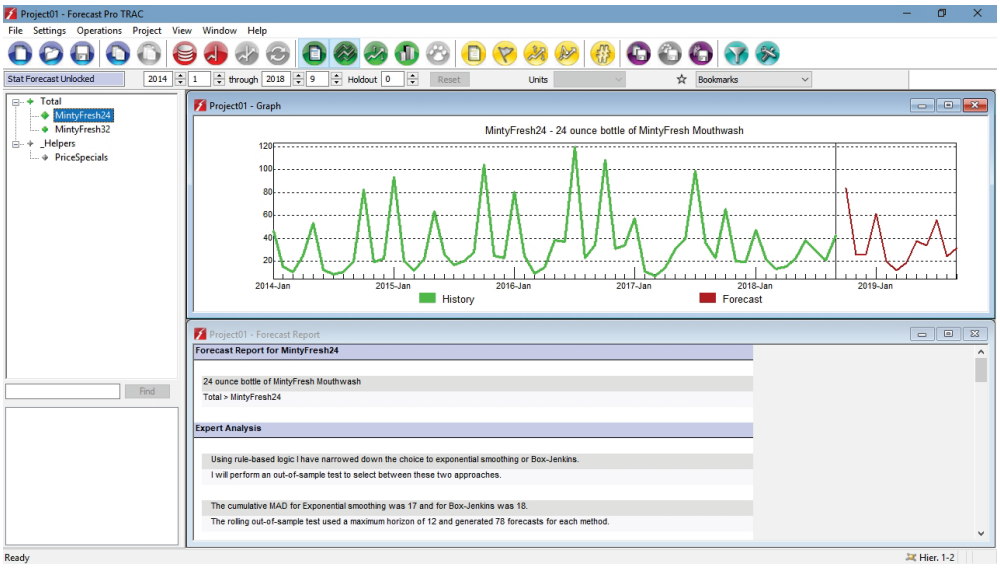

The prominent peaks in the historic demand (green line) on the graph are not due to seasonal patterns. Instead, they are the result of price promotions. Select *MintyFresh32* on the Navigator. This series represents sales of the same product in a different size package. *PriceSpecials* is an event variable—we will discuss this variable shortly.

Forecast Pro forecasted the data using expert selection. View the Forecast Report for MintyFresh24. Notice that a nontrended additive seasonality exponential smoothing is selected and that the adjusted R-square is 0.51. Exponential smoothing recognized the effects of the promotions as a form of irregular seasonality. The forecasts from such a model tend to repeat the promotional pattern of the last year of historic data. If, as is usually the case, your future promotions are patterned differently, the forecasts may be very bad.

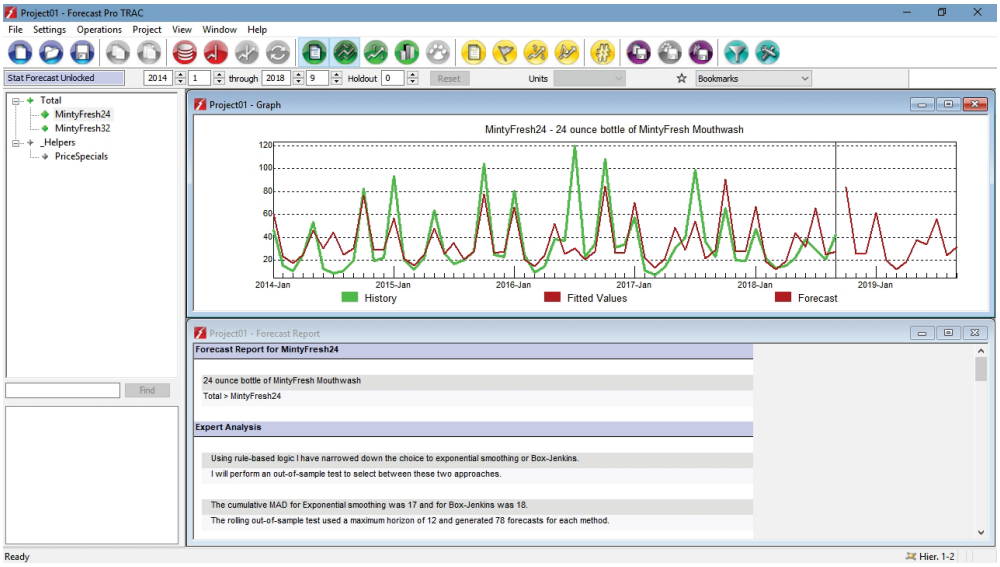

Include the fitted values on the graph. (This is accomplished by toggling on **Fitted** on the graph's context menu). Notice that the model predicted several peaks that did not occur as well as missing several that did. The model is trying to capture the peaks as part of the seasonal pattern. However, the promotions that caused these peaks did not always fall in the same months and thus cannot be modeled using a seasonal model.

Right click on the graph window to call up the context menu and select **Manage Events** to open the event manager. Use the Schedule drop down located at the bottom of dialog box to select the *PriceSpecials* schedule.

The *PriceSpecials* schedule uses Event Code 1 to indicate periods where the product was on promotion. Notice that the dates listed correspond to the strong peaks in the data and include two points during the forecast period where a promotion is scheduled.

Before we generate the forecasts, let's add a description for Event Code 1. Click the Edit button to open the Manage Event Codes dialog box. The Event Code 0 is always used to indicate no event occurred. Therefore, it is predefined and *No Event* already appears as the description. Select Event Code 1 in the lower section of the dialog box and then edit the description displayed for Event Code 1 to read *Price Promotion* and click the Add button. Notice that the description is updated with a status of "Pending." Click the

Commit button to accept the new description and then the Exit button to return to the Manage Events dialog box.

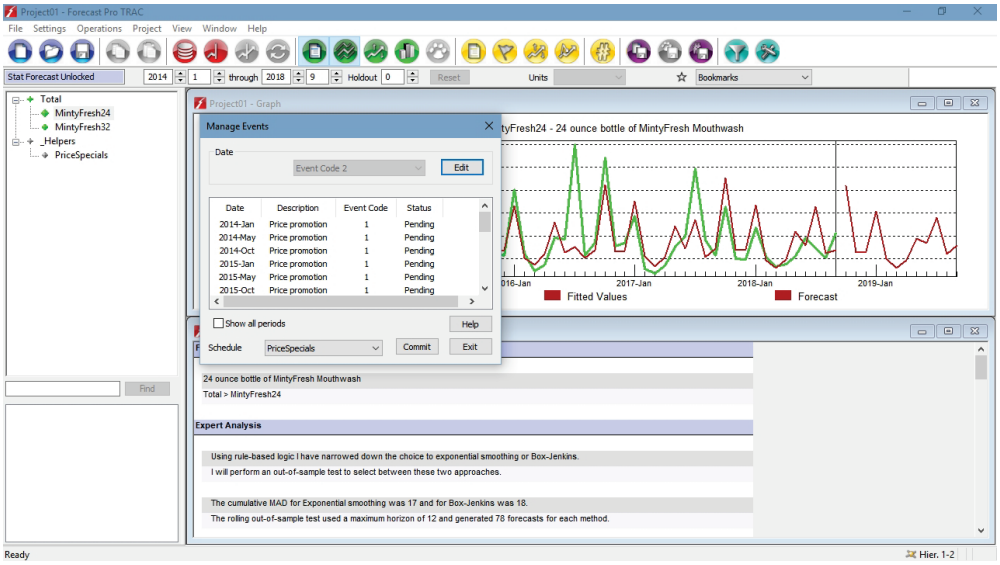

Click Commit to generate the event model forecasts.

Examine the Forecast Report for *MintyFresh24*. Notice that an exponential smoothing model was used and that the model includes an event smoothing weight. In addition, notice that the adjusted R-square is now 0.79.

The graph reveals a much better fit to the data and forecasts strong peaks that correspond to the months indicated as promoted.

Notice that on the Navigator MintyFresh24 now appears with a "\EVENT=PriceSpecials" modifier to indicate that an event model was specified. Select **Apply Modifier(s) to** on the Navigator's context menu and select **All Items.** This will specify that event models using the PriceSpecials schedule should be built for every item on the Navigator. In this example, this is appropriate since the same promotional schedule was used for both MintyFresh24 and MintyFresh32.

If the two products had been promoted differently we would have created two separate event schedules.

Exit Forecast Pro.

Modeling Weekly Data

Select **File>Open** and open the project, *Tutorial - Event Modeling*. In this exercise we will model weekly beer sales for a brand of beer in various package configurations.

Select C-6 on the Navigator. C-6 represents sales of beer packaged in cans and sold in 6-packs.

In order to get a better view of the data, right click the Graph view to invoke its context menu and select **Graph Settings**. Open the Components tab, remove the checkmark for "All", set "Periods" to 156 and click OK. Note that the display is now scrollable. *Periods* allows you to select how many data points will be shown on-screen when the *All* option is not selected.

Add the fitted values to the graph using the graph's context menu.

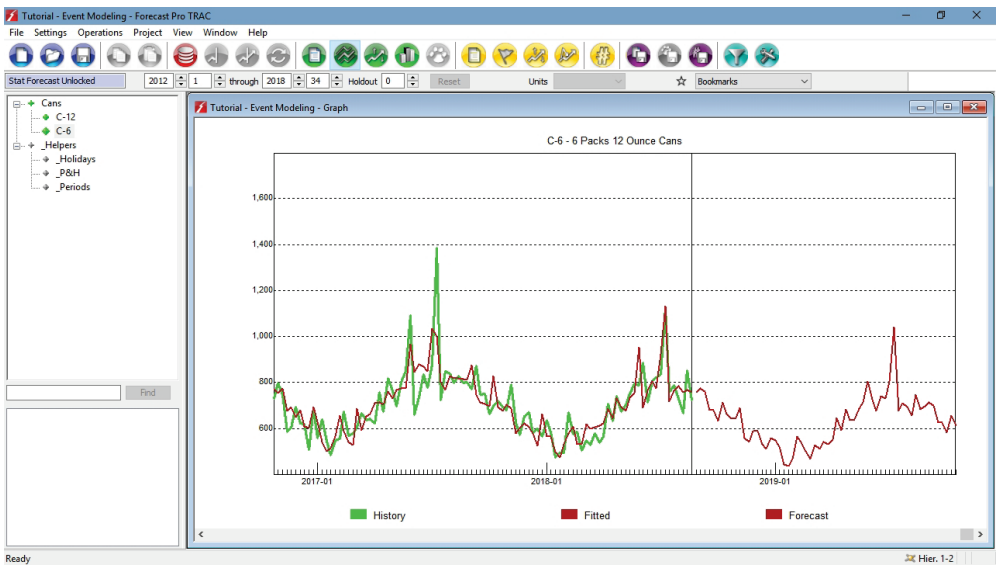

One noticeable feature of this data set is that sales for the weeks containing the three US summer holidays (Memorial Day, 4th of July, Labor Day) are higher than other weeks. If you look closely at the data you'll notice that these holidays sometimes change weeks, and when this occurs, the fit can be poor. For instance the 4th of July fell in week 27 before 2015, but fell in week 28 in

2015 and thereafter. The fitted values for week 27 in 2015 and 2016 exhibit strong peaks and “miss” the actual holiday peaks that fell in week 28.

Forecast Pro used a standard Winters model to forecast C-6. This model uses 52 seasonal indexes for the data and does not accommodate the holidays changing weeks from year to year.

Right click on the graph window to call up the context menu and select **Manage Events** to open the event manager. Use the Schedule drop down to select the _Holidays schedule.

_Holidays assigns event codes to each of the summer holidays. The event code “Memorial Day” is assigned to each week containing Memorial Day, the event code “4th of July” for the week containing the 4th of July and the event code “Labor Day” for the week containing Labor Day. Including an event index for each holiday allows us to model the holidays as they move around the calendar. (If we were to model this data as a monthly series, the holidays would not change periods and an event model would not be necessary.)

Click the Commit button to build the model.

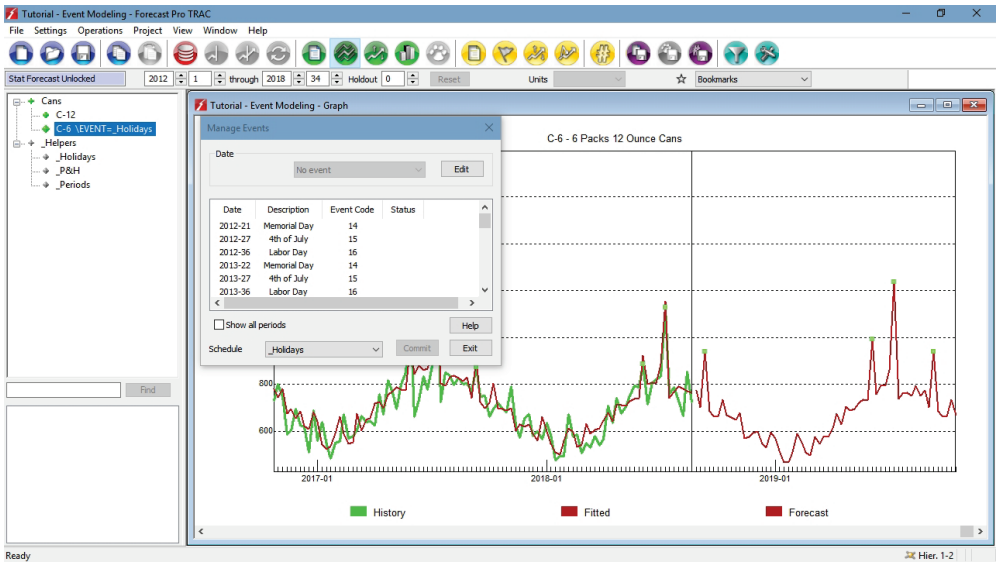

The model now includes 52 seasonal indexes to capture seasonality and three event indexes to capture the summer holidays. Examine the fit to the holiday

weeks and notice that the event model is not “missing” the holidays when they move around the calendar.

Although we would like the seasonal indexes to capture a smooth seasonal pattern in this case they do not. This is because of the high degree of random variation in the data and the complexity of using 52 indexes to model the seasonality.

Generally speaking, monthly models will more accurately capture seasonality. Nevertheless, many corporations *must* deal with weekly seasonal data because production and inventory control require it. Event adjustment models can help you with some of the problems that you will face.

Event models allow you to control how many indexes to include in the model. To illustrate, let's model C-6 using the _P&H event variable.

Use the Schedule drop down box to select _P&H. Notice that _P&H maps each week into one of 13 periods (i.e. weeks 1-4 are assigned to Period 1, weeks 5-8 are assigned to Period 2, etc.). The summer holiday weeks (Memorial Day, 4th of July and Labor Day) have their own unique event Codes.

Click the Commit button to build the model.

Turn on the Forecast Report view by clicking the Forecast Report icon ( ) and scroll down to view the event indexes. The resulting model uses 16 event indexes—13 to capture the seasonality and three for the summer holidays.

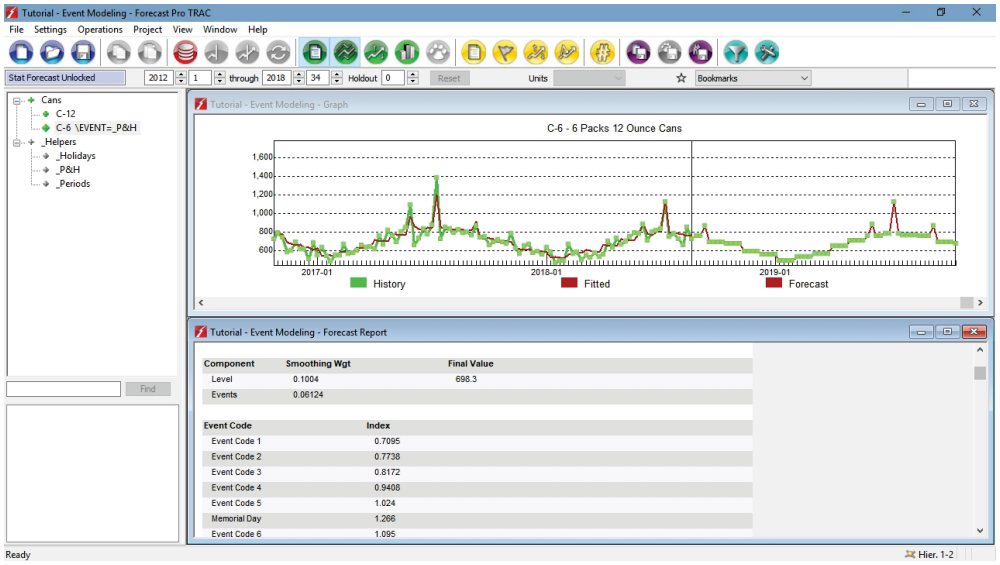

Notice that the forecasts exhibit much smoother seasonality than our previous model. That is because there are now fewer indexes and more historic observations for each index to be estimated.

Adjustment for Outliers

Historic data sets often exhibit the effects of one-time events that cause outliers. The event that causes the outlier may be known or unknown. Although exponential smoothing is a remarkably robust procedure, these outliers may decrease the quality of the forecasts and (especially) the confidence limits.

You can eliminate the effect of an outlier by coding it as a special event that occurs only once. If you have several outliers, each must be coded as a distinct event type. Forecast Pro will “explain” each outlier as the result of its associated event.

The impact of outliers on the forecasts and the confidence limits will be greatly reduced. Beware however: if outliers continue to occur in the forecast period, then the confidence limits are likely to be unrealistically narrow.

More Adjustments for Promotions

The event adjustment model provides a very flexible framework to treat promotional effects of many kinds. This section briefly notes some of them. Its purpose is to indicate some directions you may want to take with your own business data.

The examples that are given require assigning additional event types in your event variable. Each new event type provides additional ability for Forecast Pro to explain your historic data by making the event variable more complex. Keep in mind that if your event description is overly complex, the out-of-sample performance of your model may deteriorate. You must strike the right trade-off between goodness-of-fit to your historic data and model complexity. To do this will require experimentation and monitoring of actual model performance.

Example 1. A promotion in (say) September may have effects in August and October as well. Buyers may delay purchases in August and they may be overstocked in October. You can assign the pre- and post-promotional effects as event types of their own. These events will, of course, be associated with *decreases* in sales.

Example 2. Sometimes one SKU of a brand or product line is promoted but closely related SKUs are not. The result may be that the promoted SKU cannibalizes the sales of the other SKUs. You can treat this effect by assigning cannibalization events for these SKUs. But be cautious. Overuse of this technique will result in an overly complex model and possible deterioration of forecast performance.

This concludes the *Building Event Models* lesson.

Lesson 15

Detecting and Correcting Outliers

An outlier is a data point that falls outside of the expected range of the data (i.e., it is an unusually large or small data point). If you are forecasting a time series that contains an outlier there is a danger that the outlier could have a significant impact on the forecast.

One solution to this problem is to screen the historical data for outliers and replace them with more typical values prior to generating the forecasts. This process is referred to as outlier detection and correction.

Correcting for a severe outlier (or building an event model for the time series if the cause of the outlier is known) will often improve the forecast. However if the outlier is not truly severe, correcting for it may do more harm than good. When you correct an outlier, you are rewriting the history to be smoother than it actually was and this will change the forecasts and narrow the confidence limits. This may result in poor forecasts and unrealistic confidence limits when the correction was not necessary.

It is the author's opinion that outlier correction should be performed sparingly and that detected outliers should be individually reviewed by the forecaster to determine whether a correction is appropriate.

Forecast Pro incorporates an automated algorithm to detect and (optionally) correct outliers. In this lesson we will explore its operation.

How Outlier Detection and Correction Works

Start Forecast Pro, select **Settings>Options** and display the Outliers tab.

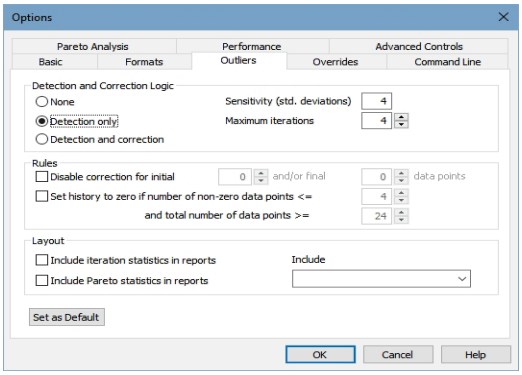

There are three modes for outlier detection and correction.

None turns outlier detection and correction off. This is the default and many Forecast Pro users do not use (nor need to use) outlier detection.

Detection only will detect outliers and display the suggested corrected values, however, the forecasts will be generated using the uncorrected history.

Detection and correction will detect outliers and will automatically use the corrected values when generating forecasts.

The detection and correction algorithm works as follows:

1. The specified forecasting model is fit to the time series, the residuals (fitted errors) are generated and their standard deviation is calculated.

2. If the size of the largest error exceeds the outlier threshold, the point is flagged as an outlier and the historic value for the period is replaced with the fitted value.

3. The procedure is then repeated using the corrected history until either no outliers are detected or the specified maximum number of iterations is reached.

In a multiple-level problem the detection is only performed on the end items (i.e., the nongroup level). If the correction option has been selected, after all end items are corrected, the group level totals are reaggregated to reflect the corrected values.

You can adjust the *Sensitivity* setting to make the outlier threshold more or less sensitive. The proper setting will depend on the stability of your data set.

Set the detection and correction mode to *Detection only*, make sure the other settings match the selections shown above and click OK.

Reviewing the Outliers

In this exercise we will perform the following steps:

1. Use the *Detection only* mode to identify the outliers.

2. Use the Outlier Report view and the Graph view to review the detected outliers and determine whether a correction is warranted.

3. Use the \OUTLIER=CORRECT modifier to correct the outliers that we feel should be corrected.

Click the red Data Manager icon ( ) to call up the Data Manager.

Click the Add drop down on the Historic Data row, select Excel and select *123 Bakery – Historic Data.xlsx* to add it to the Historic Data row.

Click the Read & Forecast button to read in the historic demand data and generate the statistical forecasts.

Click the Exit button to exit the Data Manager.

Fully expand the Navigator, click the yellow View Outlier Report icon ( ) to open the outlier report view and turn off any other open views.

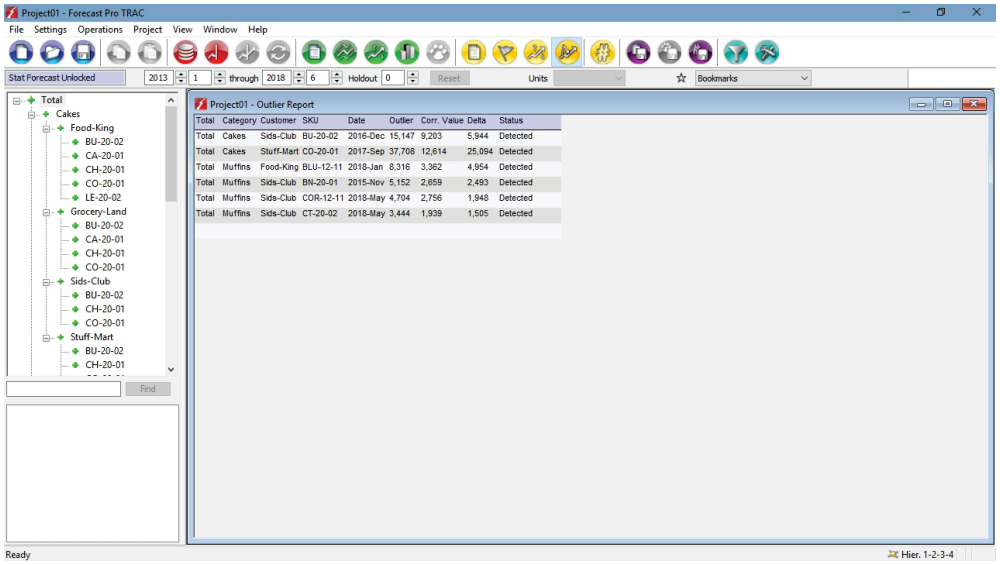

Notice that a total of 6 outliers have been detected, and their current Status is *Detected*.

Correcting the Outliers

Open the Graph view. Right click on the graph to bring up the context menu and select **Graph Settings**. Click on the Add or Remove button to bring up the Add or Remove dialog box and select *Corrected History*. Click on OK to close the Add or Remove dialog box. Make sure Corrected History is now included in the Layout tab of the Graph Settings dialog box and that the Display checkbox is selected. Click OK to close the Graph Settings dialog box.

Double-click on the first item listed on the outlier report. Your display should now match the one shown below.

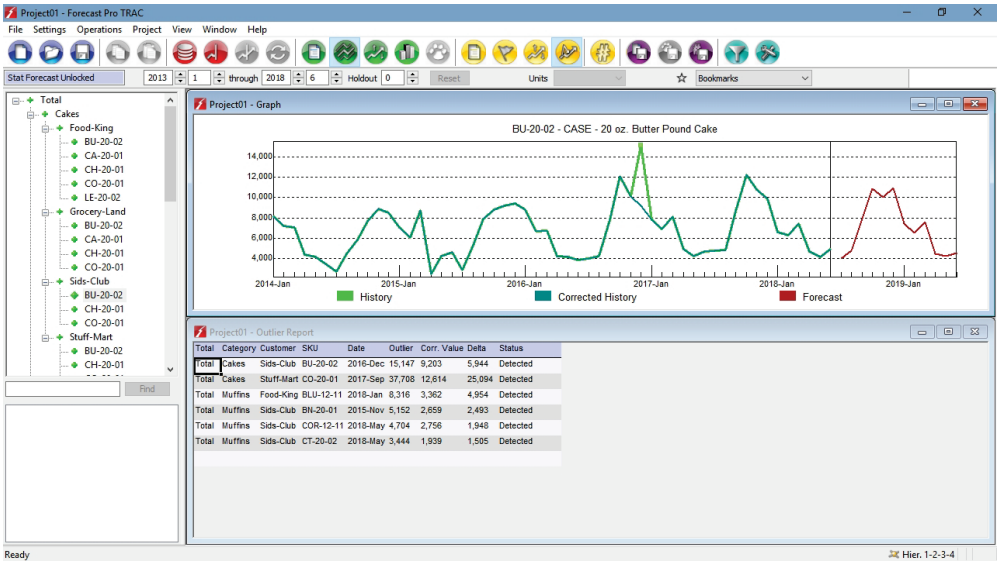

Notice that when you double-click an item on the outlier report, the Navigator jumps directly to that item. Thus, you can use the outlier report as the equivalent of a Hot List to navigate through the listed items. All of the global report views (yellow-icon views) support this kind of navigation.

Notice that December 2016 outlier is shown with a marker on the graph, and adding in Corrected History allows us to see the suggested correction. The data point is rather unusual and you can see that the forecast has a spike in December which is why it was flagged as an outlier.

Right click the item on the Navigator (not the outlier report) to bring up the Navigator's context menu. Select **Outliers>Correct**. Notice that the status in the outlier report has changed to Corrected and that correcting for this outlier has changed the shape of the forecasts.

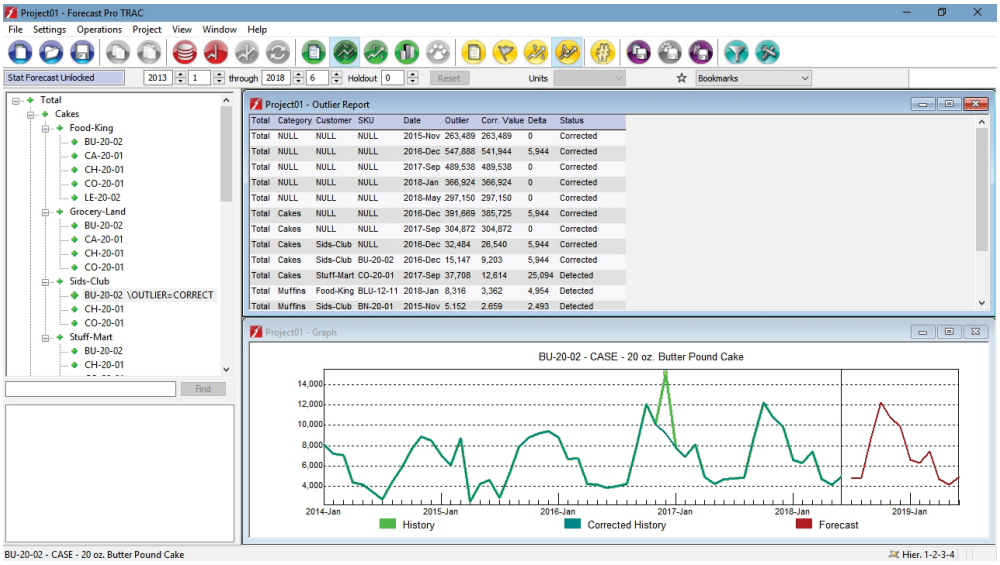

Double-click on the next item on the outlier report. Here too, we have a fairly significant outlier and will want to correct for it. Examine the remaining items on the Hot List and experiment with the Outlier options on the Navigator's context menu. When you are comfortable with their operation, exit the program.

This concludes the *Detecting and Correcting Outliers* lesson.

Lesson 16

Building Multiple-Level Models

In this lesson, you will learn how to define group variables and prepare multiple-level forecasts.

Defining Groups

Click the red Data Manager icon ( ) to call up the Data Manager.

Click the Add drop down on the Historic Data row, select Excel and select *Presto Camera Company.xls* to add it to the Historic Data row.

Click the Read & Forecast button to read in the historic demand data and generate the statistical forecasts.

Click the Exit button to exit the Data Manager.

Open the graph view and fully expand the Navigator. Your screen should now look like the one shown below.

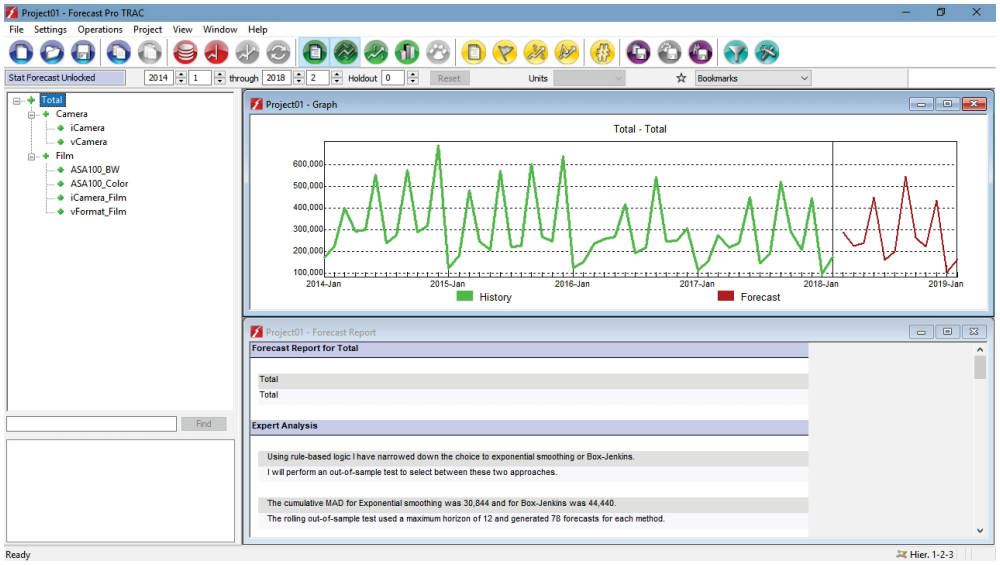

Notice that the Navigator tree structure matches the hierarchy shown below.

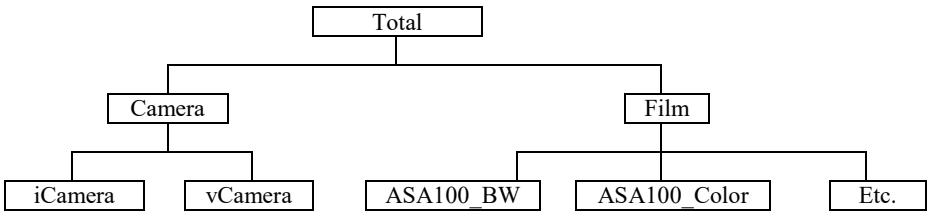

The tree contains three *group* variables (Total, Camera and Film). The tree also contains six end items (iCamera, vCamera, ASA100_BW, ASA100_Color, iCamera_Film and vFormat_Film). The demand histories for group variables are not in the database. They are created by Forecast Pro by aggregating the appropriate end items. Thus, the history for the group Camera is defined as the sum of iCamera + vCamera. The history for the group Film is defined as the sum of ASA100_BW + ASA100_Color + iCamera_Film + vFormat_Film. The history for the group Total is defined as the sum of the groups Camera + Film.

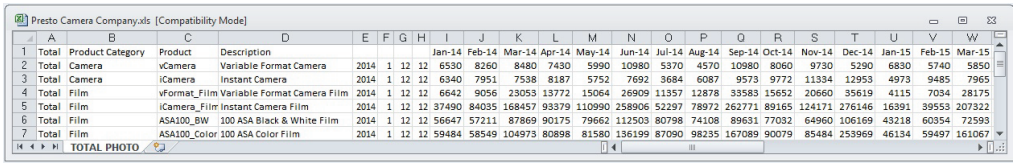

Presto Camera Company.xls [Compatibility Mode]

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W
1	Total	Product Category	Product	Description					Jan-14	Feb-14	Mar-14	Apr-14	May-14	Jun-14	Jul-14	Aug-14	Sep-14	Oct-14	Nov-14	Dec-14	Jan-15	Feb-15	Mar-15
2	Total	Camera	vCamera	Variable Format Camera	2014	1	12	12	6530	8260	8480	7430	5990	10980	5370	4570	10980	8060	9730	5290	6830	5740	5850
3	Total	Camera	iCamera	Instant Camera	2014	1	12	12	6340	7951	7538	8187	5752	7692	3684	6087	9573	9772	11334	12953	4973	9485	7965
4	Total	Film	vFormat_Film	Variable Format Camera Film	2014	1	12	12	6642	9056	23053	13772	15064	26909	11357	12878	33583	15652	20660	35619	4115	7034	28175
5	Total	Film	iCamera_Film	Instant Camera Film	2014	1	12	12	37490	84035	168457	93379	110990	258906	52297	78972	262771	89165	124171	276146	16391	39553	207322
6	Total	Film	ASA100_BW	100 ASA Black & White Film	2014	1	12	12	56647	57211	87869	90175	79662	112503	80798	74108	89631	77032	64960	106169	43218	60354	72593
7	Total	Film	ASA100_Color	100 ASA Color Film	2014	1	12	12	59484	58549	104973	80898	81580	136199	87090	98235	167089	90079	85484	253969	46134	59497	161067

TOTAL PHOTO

The spreadsheet *Presto Camera Company.xls* is shown above.

The groups Total, Camera and Film are defined using attribute fields (columns A and B).

If we had included additional attribute fields we would have created a larger hierarchy. Consult the *Setting Up Your Historic Data* section of this manual for complete details.

Bottom-up and Top-down Forecasting

When you forecast a multiple-level hierarchy Forecast Pro will automatically reconcile the forecasts. By default a bottom-up reconciliation will be performed (this is the case in our current example). That means that the group-level forecasts will be generated by aggregating their component forecasts. For example, the forecast for the group Camera would be made by summing the forecasts of vCamera and iCamera.

An alternative is to use top-down reconciliation. A top-down approach creates forecasts at a group level and then proportionally adjusts lower-level forecasts to sum to the group-level forecast. This can improve forecast accuracy in cases where the lower-level items form a homogenous group due to the higher volume (and usually lower variability) present at the group level.

Select *Total>Camera* on the Navigator and right click to display the Navigator's context menu. Select Top Down. Notice that the \TOPDOWN modifier now appears next to Camera.

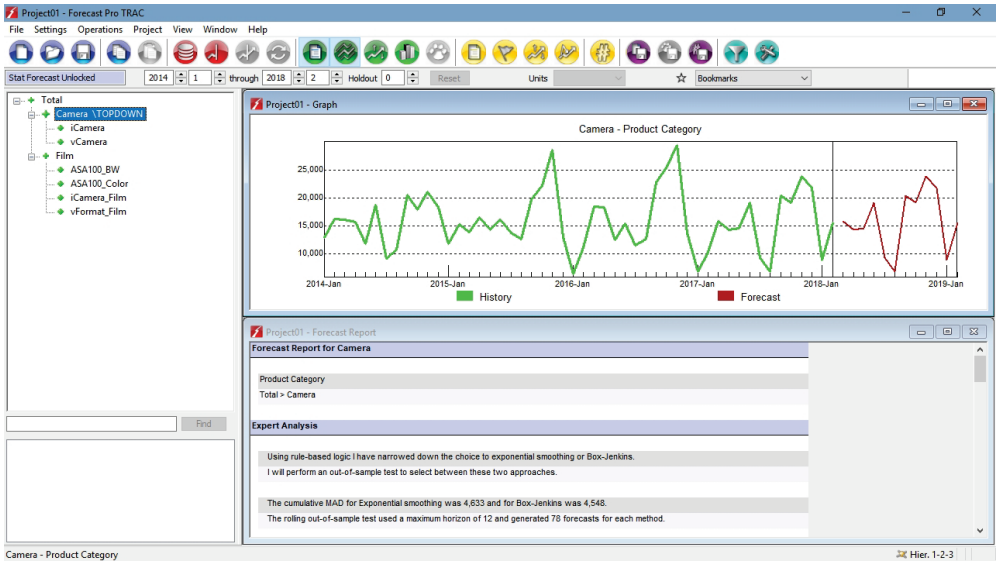

The \TOPDOWN group modifier instructs the program to adjust lower-level forecasts to sum up to the indicated level. Let's illustrate how this works by considering the various steps involved in forecasting our current example.

Step One: Prepare forecasts at all levels (i.e., prepare forecasts for Total, Camera, Film, iCamera, vCamera, ASA100_BW, ASA100_Color, iCamera_Film and vFormat_Film based on their past history).

Step Two: Since the group Camera is being forecasted top down, the forecasts of iCamera and vCamera are adjusted proportionally so that they sum to the Camera forecast.

Step Three: Since the group Film is being forecasted bottom up (the default) the Film forecast is replaced with the sum of the forecasts for ASA100_BW, ASA100_Color, iCamera_Film and vFormat_Film.

Step Four: Since the group Total is being forecasted bottom up (the default) the Total forecast is replaced with the sum of the forecasts for Camera + Film.

After you have examined the output, exit Forecast Pro.

There are two additional reconciliation modifiers—\CALL and \INDEXES. The \CALL modifier is discussed in the next section. \INDEXES instructs the program to estimate the seasonal indexes at the group level and use them for

all members in the group. This is particularly useful when some members of the group have short histories or small demand and seasonal factors are thus difficult to calculate.

Using a Custom Allocation

In this section we will explore using a custom allocation to construct the forecasts. A custom allocation approach allows you to dictate how a group-level forecast is allocated to its component series. This is appropriate when the breakdown between the two levels is known.

In this example, we will forecast demand for running shoes and use a custom allocation to determine how the statistically generated forecast for each specific shoe should be broken out to the different sizes.

Click the red Data Manager icon ( ) to call up the Data Manager.

Click the Add drop down on the Historic Data row, select Excel and select *Running Shoes – Historic Data.xlsx* to add it to the Historic Data row.

Click the Add drop down on the Definitions row, select Excel and select *Running Shoes – Custom Allocation.xlsx* to add it to the Definitions row.

Click the Read & Forecast button to read in the historic demand data and generate the statistical forecasts. Notice that since the on Read option is selected for the Definitions row, this action read in the definitions data as well.

Click the Exit button to exit the Data Manager.

Open the graph view and fully expand the Navigator. Your screen should now look like the one shown below.

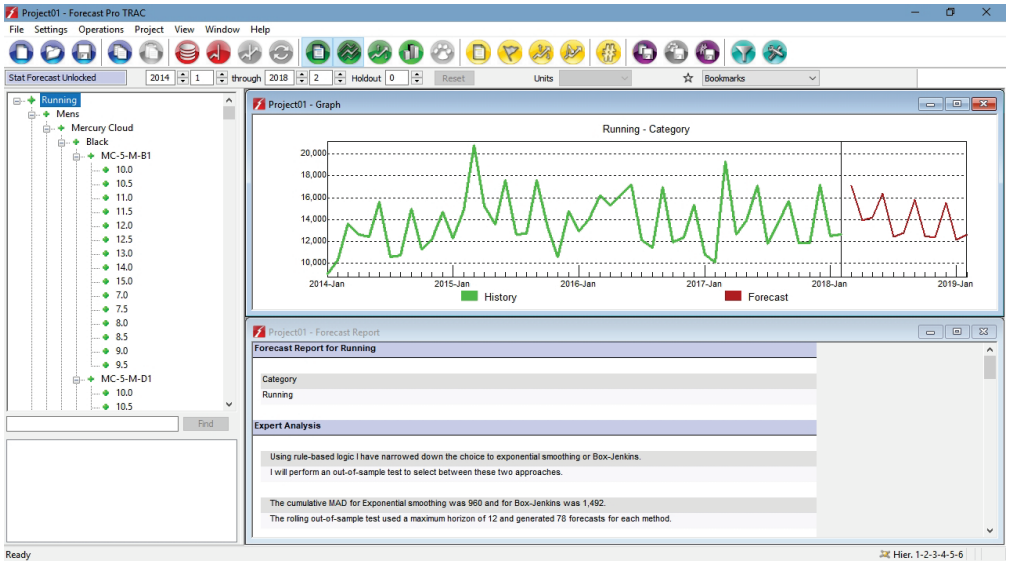

Notice that the data are organized into a 6-level hierarchy including category, gender, style, color, style ID and size.

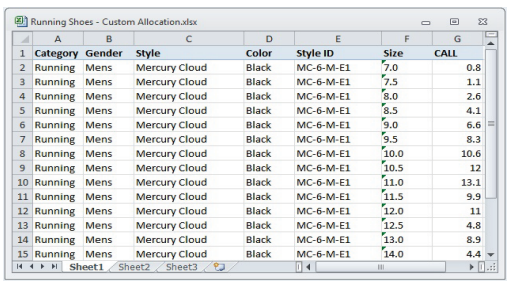

Running Shoes - Custom Allocation.xlsx

	A	B	C	D	E	F	G
1	Category	Gender	Style	Color	Style ID	Size	CALL
2	Running	Mens	Mercury Cloud	Black	MC-6-M-E1	7.0	0.8
3	Running	Mens	Mercury Cloud	Black	MC-6-M-E1	7.5	1.1
4	Running	Mens	Mercury Cloud	Black	MC-6-M-E1	8.0	2.6
5	Running	Mens	Mercury Cloud	Black	MC-6-M-E1	8.5	4.1
6	Running	Mens	Mercury Cloud	Black	MC-6-M-E1	9.0	6.6
7	Running	Mens	Mercury Cloud	Black	MC-6-M-E1	9.5	8.3
8	Running	Mens	Mercury Cloud	Black	MC-6-M-E1	10.0	10.6
9	Running	Mens	Mercury Cloud	Black	MC-6-M-E1	10.5	12
10	Running	Mens	Mercury Cloud	Black	MC-6-M-E1	11.0	13.1
11	Running	Mens	Mercury Cloud	Black	MC-6-M-E1	11.5	9.9
12	Running	Mens	Mercury Cloud	Black	MC-6-M-E1	12.0	11
13	Running	Mens	Mercury Cloud	Black	MC-6-M-E1	12.5	4.8
14	Running	Mens	Mercury Cloud	Black	MC-6-M-E1	13.0	8.9
15	Running	Mens	Mercury Cloud	Black	MC-6-M-E1	14.0	4.4

The CALL column in the file *Running Shoes – Custom Allocation.xlsx* contains the size breakdowns to use for the allocations. If you examine the file you will notice that all of the Men's shoes use consistent size proportions to define the allocation between the Style ID and Size levels. The Women's shoes also use (different) consistent proportions.

Since the size of runners' feet is not changing rapidly in time, using a size chart or a historic frequency analysis chart to perform the allocation rather than explicitly forecasting each shoe-by-size combination based on its historic demand is likely to yield a more accurate forecast.

Right click *Running>Mens>Mercury Cloud>Black>MC-5-M-B1* on the Navigator to invoke the context menu. Select **Custom Allocation**. Notice that the \CALL modifier now appears on *Running>Mens>Mercury Cloud>Black> MC-5-M-B1*. This instructs Forecast Pro to use a custom allocation to break down the *Running>Mens>Mercury Cloud>Black> MC-5-M-B1* group forecast to the various sizes.

In our example we wish to use the custom allocation for all of the size breakdowns. Right click *Running>Mens>Mercury Cloud>Black> MC-5-M-B1* on the Navigator to invoke the context menu. Select **Apply Modifier(s) to>This Level** to place the \CALL modifier on all Style ID's. Your screen should now match the one below.

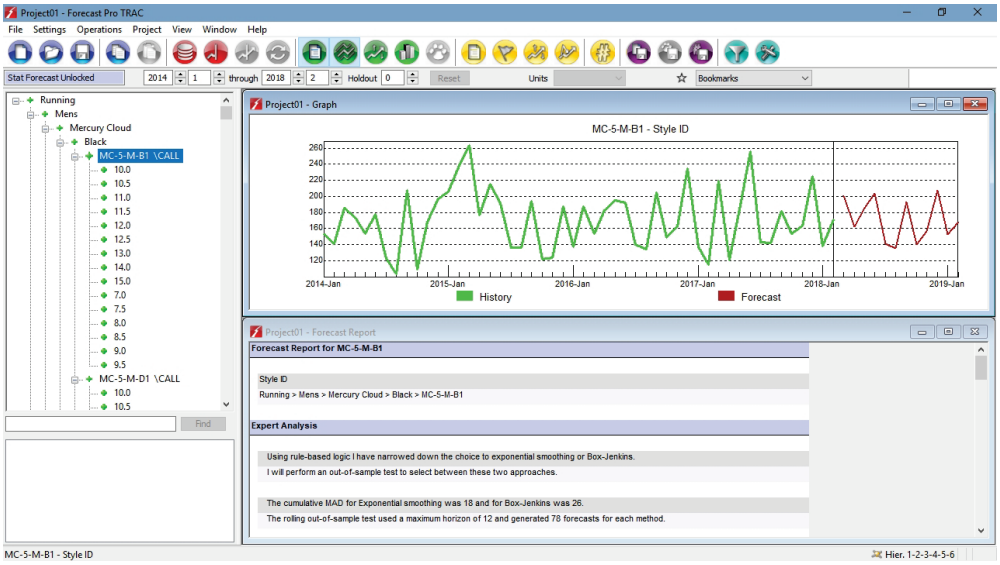

After you have examined the output, exit Forecast Pro.

Combining Multiple Level and Event Adjustment Models

In many cases, the SKUs or product lines in a multiple-level hierarchy are subject to promotional effects. You can use the event model methodology described in the *Building Event Models* tutorial to account for such effects.

Example 1. A product line, consisting of aggregated SKUs is promoted as a unit. You believe that all the SKUs in the product line will be affected similarly. You can deal with this by using the following approach.

```
LINE \EVENT=PROMO \TOPDOWN
        SKU1
        SKU2
        SKU3
```

This causes Forecast Pro to execute the following procedure.

Forecast the group LINE, taking promotions into account.

Forecast the SKUs *without* taking promotions into account.

Adjust the SKU forecasts so that they sum to the LINE forecasts. This propagates the group level promotional effects to the SKUs.

Example 2. Now suppose that only one SKU of the group is promoted. The following approach can account for this promotion.

```
LINE
        SKU1 \EVENT=PROMO
        SKU2
        SKU3
```

This causes Forecast Pro to forecast the SKUs individually. It then sums these forecasts to obtain forecasts for the group LINE.

What if you were to add the keyword \TOPDOWN to the group LINE in this example? The effect would be that the promotion of SKU1 affects it but does not affect overall sales at the group level. This might be the case if the promotion affected sales only by cannibalizing other SKUs in the group. This is a dubious effect. It is usually better to prepare forecasts bottom-up when there are distinct models for the lower-level units.

Short-Lived Products

A manufacturer often maintains a product line for a long period of time but frequently changes the SKUs that make up the line. A laser printer

manufacturer, for instance, often introduces new models and retires old ones. The result may be that the overall product line can be accurately forecasted, but the individual item histories are too short to support seasonal models. In these instances the top-down approach is particularly useful.

For instance assume that SKU1 and SKU2 have been phased out and replaced by SKU3 and SKU4. We also assume that at the end of the historic data, only SKU3 and SKU4 are alive. However, the histories for SKU3 and SKU4 are too short to generate seasonal forecasts. The manufacturer is interested in forecasting the group LINE, SKU3 and SKU4. The approach below takes care of the problem.

```
LINE \INDEXES
        SKU1
        SKU2
        SKU3
        SKU4
```

This causes Forecast Pro to follow the following procedure.

Forecast LINE, obtaining seasonal indexes.

Use the LINE seasonal indexes to deseasonalize the SKUs.

Forecast the resulting nonseasonal SKU-level data. These nonseasonal models require very little data.

Use the LINE seasonal indexes to reseasonalize the SKU-level forecasts.

By default, the LINE forecasts are then replaced by the summed SKU-level forecasts. If you do not want this to happen, you can add the keyword \TOPDOWN to the group LINE.

The approach presented in this example can also be used when the SKU-level histories are long lived. The result is that seasonality is accounted for at the LINE level. This is desirable when the SKUs are likely to have similar seasonal patterns, but the data are too irregular for accurate estimation of seasonal indexes at the SKU level.

This concludes the *Building Multiple-Level Models* lesson.

Lesson 17

Building Custom Component Models

In this lesson you will learn how to build custom component models. The model is based on the components found in an exponential smoothing model, specifically, level, trend, seasonality and events. In a standard smoothing approach, final values of each component are estimated from the data and used to assemble the forecasts.

In a custom component model you have the ability to either let Forecast Pro estimate the final value for a given component or to customize the values to be used.

Custom component models are useful in a variety of situations including (1) customizing the trend for longer-term forecasts, (2) customizing the seasonal pattern for short data sets and (3) defining the impact of future events that have not occurred historically.

Customizing the Trend

In this first example we will prepare a five-year forecast and experiment with customizing the trend.

Start Forecast Pro. The data we will be using is monthly and the first step will be to set the forecast horizon to 60 so that we will generate a 5-year forecast.

Click the light blue Options icon ( ) to open the Options dialog box shown below.

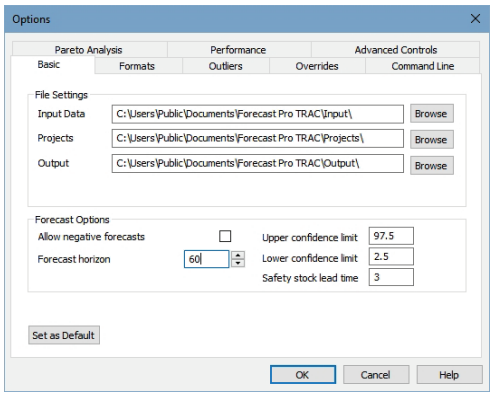

Set the Forecast horizon field to 60 and click OK. Click the red Data Manager icon ( ) to call up the Data Manager. Click the Add drop down on the Historic Data row, select Excel and select *Telecom installations.xls* to add it to the Historic Data row.

Click the Read & Forecast button to read in the historic demand data and generate the statistical forecasts.

Click the Exit button to exit the Data Manager.

Open the Graph view by clicking its green icon ( ). Your display should now match the one shown below.

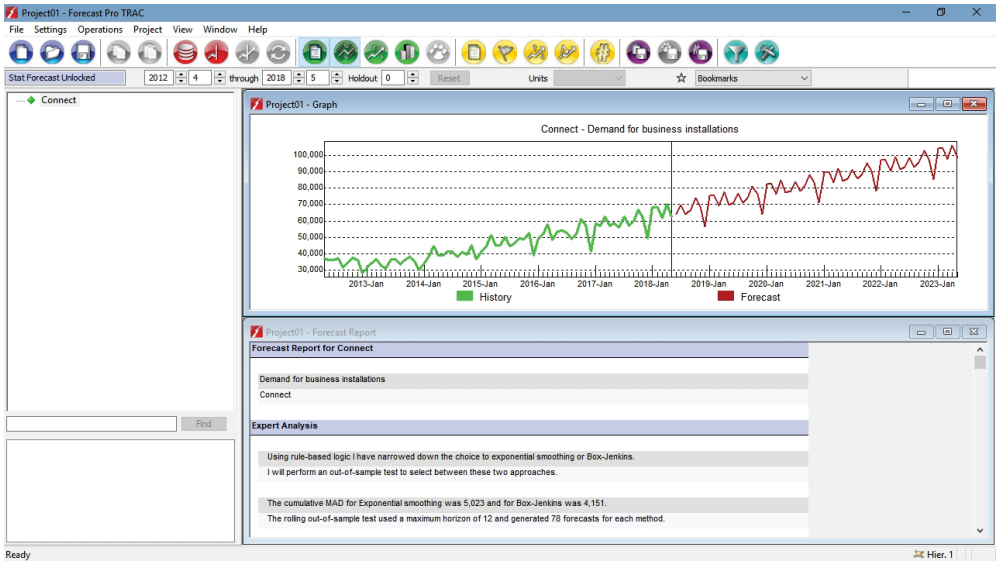

The data represent monthly demand for telephone installations. Notice the data and forecasts exhibit a strong trend and seasonal pattern. Notice also that expert selection chose to forecast this data set using a Box-Jenkins model.

Turn off the Forecast Report view by clicking its green icon (). Right click *Connect* on the Navigator to display the context menu and select **Model>Component** to invoke the Component Model Settings dialog box. Click the Apply button to build the initial model. Move the Component Model Settings dialog box so that you can clearly see the graph. Your screen should now match the one below.

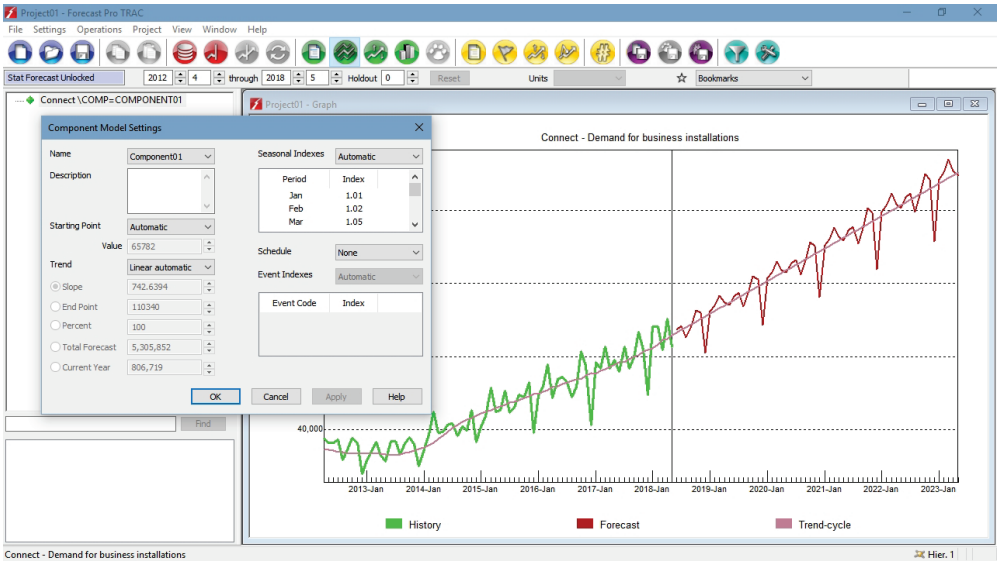

The initial model automatically calculates each component using an exponential smoothing approach. Notice that the Starting Point is currently set to Automatic and has a value of 65782. The Trend is set to linear automatic and has a slope of 742 units per month. The Seasonal Indexes are set to Automatic and are displayed in the dialog box. We have not defined an event schedule, so Event Indexes are not being used.

The dark green line running through the history and forecasts on the graph is the trend component. As noted above the automatic approach yields a linear (straight line) trend that increases 742 units a month for the next 5 years.

This is a very aggressive long-term forecast. Let's say that our knowledge of the market leads us to believe that this kind of growth is unlikely to be maintained and we therefore wish to modify the underlying trend component of the forecast.

Open the Trend drop down. Notice there are three types of trends supported (None, Linear, Bent) and that the linear and bent types have both an automatic and custom option. Choose Linear custom.

Linear custom will use a linear trend but allow you to control the slope. You can use the Slope option to set the slope directly or use the End Point, Total Forecast or Current Year options to have Forecast Pro calculate the new slope.

Set the Total Forecast field to 4,500,000 and click the Apply button.

Notice the forecast now reflects the new slope (302 units per month) and the values for the End Point, Total Forecast or Current Year options have been updated accordingly.

Now let's explore using a bent trend. If you bend a trend downward, the monthly slope increment will decrease each forecast period resulting in a "bent" trend. If you bend a trend upward the monthly slope increment will increase each forecast period resulting in an exponential growth pattern.

There are two Bent trend options—automatic and custom. Bent automatic allows you to bend the automatically generated slope (which in our example was 742 units per month). Bent custom allows you to specify a custom slope and then bend it.

Use the Trend drop down to select Bent automatic. Set Total Forecast to 4,500,000 and click the Apply button. Your screen should now match the one below.

Experiment with the Component Model Settings box including changing the Starting Point and Seasonal Indexes. Be sure to click the Apply button whenever you want to see the results from your current specification. When you are comfortable with the dialog box's operation, exit Forecast Pro.

Customizing Events

Select **File>Open** and open the project, *Tutorial - Customizing Events*. In this exercise we will model weekly beer sales for a brand of beer in various package configurations.

Right click *Cans>C-6* on the Navigator to display the context menu and select **Model>Component** to invoke the Component Model Settings dialog box. Click the Apply button to build the initial model. Your screen should now match the one below.

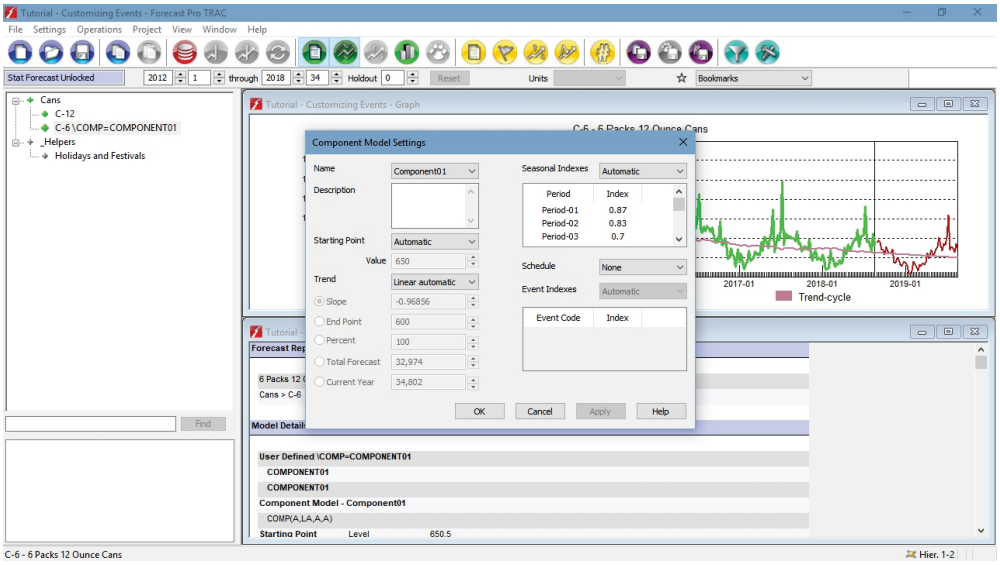

Notice that the initial component model uses automatic starting point, trend and seasonal indexes. Notice also that since there is not an event schedule specified there are no event codes listed.

In the event model example in the *Building Event Models* tutorial we illustrated that the forecasts for this data could be improved using an event model to capture the timing of the summer holidays. In this lesson we will use event indexes in the component model to not only capture the impact of the summer holidays, but also to capture the expected impact of *City Fest* a new summer festival that is scheduled for 2019.

Open the Schedule drop down menu, select *_Holidays and Festivals* and click Apply.

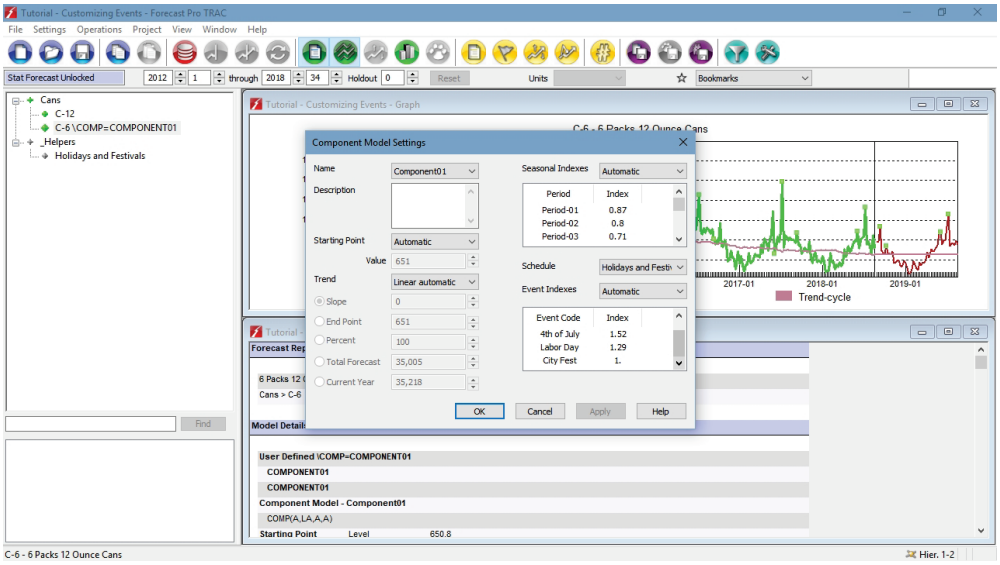

When you specify an event schedule the component model will rebuild all automatically generated components to reflect the events. Thus in our example, in addition to adding the event indexes to the model, the starting point, trend and seasonal indexes were also recalculated.

The Event Indexes mode is currently set to Automatic. Therefore the program has estimated the event indexes from the historic data. Because City Fest has never occurred historically, its index is set to 1.

Let's say that our market research leads us to believe that the City Fest event will raise beer sales by 20%. Therefore we would like to set the City Fest Index to 1.2 rather than 1.0. To accomplish this we will use a custom event approach rather than an automatic one.

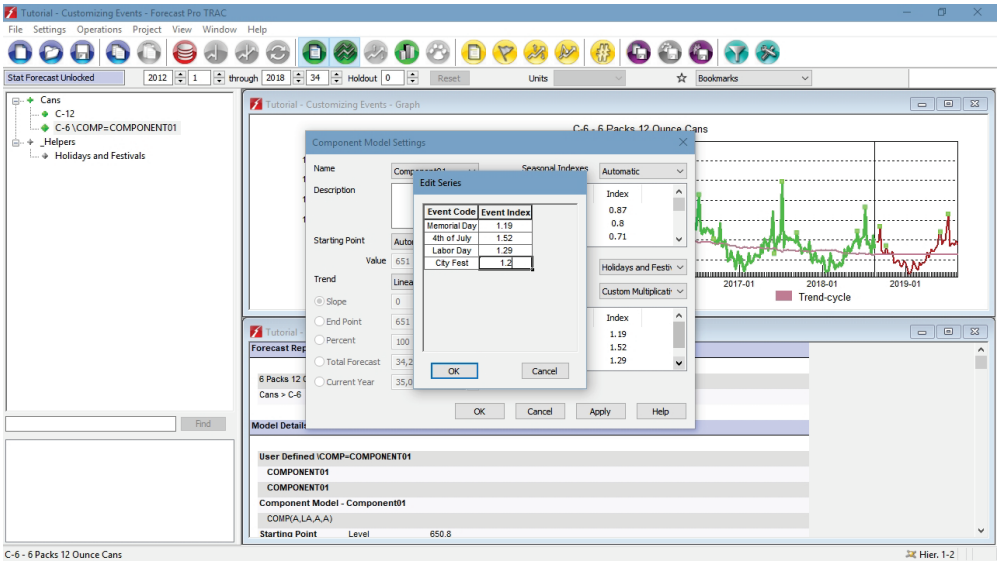

Open the Event Indexes drop down and select *Custom Multiplicative.* A dialog box appears allowing you to input custom indexes. Change the City Fest index from 1.00 to 1.20 and leave the other indexes at their current values. Click OK to exit the Edit Series dialog box and click Apply to accept the new model. Notice that the new City Fest index is now being used.

Experiment with the Component Model Settings box until you are comfortable with its operation and then exit Forecast Pro.

This concludes the *Building Custom Component Models* lesson.

Lesson 18

New Product Forecasting

In this lesson, you will explore different approaches to forecasting new products. The lesson begins by examining some of the options available in Forecast Pro for generating forecasts prior to the product being launched (i.e., when historic data are not yet available). It then illustrates how to apply two forecasting methods designed specifically for forecasting new products—forecast by analogy and the Bass model.

Working with No Data

Start the program, click the red Data Manager icon, specify *New Product-Historic Data.xlsx* on the Historic Data Row, click the Read & Forecast button and then click Exit to exit the Data Manager. Open the graph view. Your screen should now look like the one shown below.

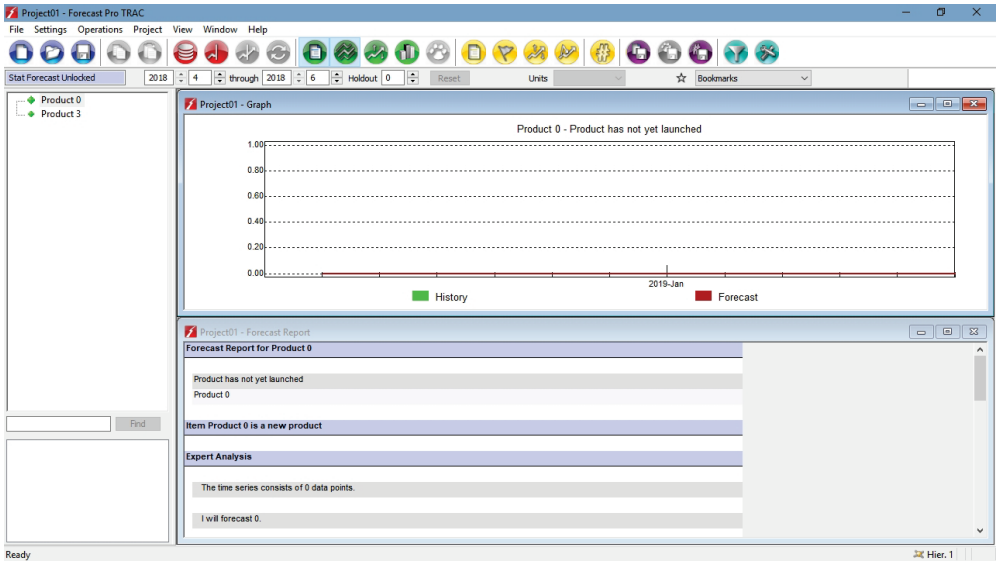

Notice that *Product-0* has no demand history. By default, Forecast Pro will generate a forecast of zero for a time series with no history. Notice also the confidence limits are set equal to the forecasts. The first forecast period is established by the historic data as a whole—in this example, July 2018.

One option is to use the override facility to enter the desired forecast in the form of an override.

Another option is to use the fixed forecast value model. This model allows you to set the statistical forecast to a specific value. Right click *Product-0* on the Navigator to invoke the context menu and then select **Model>Very Simple Models>Fixed Forecast Value**.

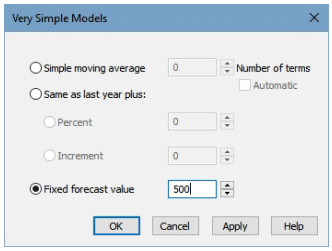

A dialog box will appear allowing you to specify the value. Set the value to 500 and click OK. Examine the Graph and Forecast Report views to verify that the forecast is now set to 500.

A related modeling option to Fixed Forecast Value is Fixed Value After. This option is used when you wish to set the statistical forecast to a specific value after a specific point in time. This can be useful when you are discontinuing a product and wish to set the forecast to zero part way through the forecast horizon or for a new product that won't come on-line until part way through the forecast horizon.

For our example, let's suppose that *Product-0* will not come on-line until October 2018. We'd therefore like the forecast to be zero until September 2018 and 500 units per month thereafter.

To accomplish this we must first remove the Fixed Forecast Value setting. Right click *Product-0* on the Navigator to invoke the context menu and then select **Remove Modifier(s)**. This will restore the default forecast of zero.

Right click *Product-0* on the Navigator to invoke the context menu and then select **Model>Very Simple Models>Fixed Value After**.

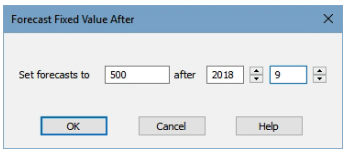

Enter 500 as the Set forecasts to value and set the after date to September 2018. Click OK and verify the forecast is now zero until September 2018 and 500 thereafter.

Exit Forecast Pro.

Forecast by Analogy

In this exercise we'll create forecasts for a new product using a technique known as forecast by analogy. This approach is sometimes also referred to as "looks like" analysis.

The concept is a very simple one. You are launching a new product and you expect the initial sales pattern to be similar to an analogous product's initial sales pattern or to a "launch profile" that you've created. To use this method, you must supply the launch profile or "analog series" in the form of a helper

variable. Helper variables are used in conjunction with event models, forecast by analogy models and weighting transformations.

If the product has not yet launched (i.e., there are no historic data available) then you must also supply an estimate of the initial sales over a specific period of time (the “launch total” over the “launch horizon”). Forecast Pro will then create the forecast by proportionally allocating the launch total over the launch horizon using the analog series to define the proportions.

If any historic demand for the new product is available, you may either supply the launch total and launch horizon or allow Forecast Pro to calculate the forecasts by fitting the existing data to the analog series.

In this exercise, we will look at two examples—one where no data are available for the new product and the other where the initial 3 months of sales are known.

Start the program, click the red Data Manager icon, specify *New Product-Historic Data.xlsx* on the Historic Data Row and then specify *New Product-Helpers.xlsx* on the Helpers row. Click the Read & Forecast button, click the Refresh button on the Helpers row and then click Exit to exit the Data Manager. Open the graph view and fully expand the Navigator. Your screen should now look like the one shown below.

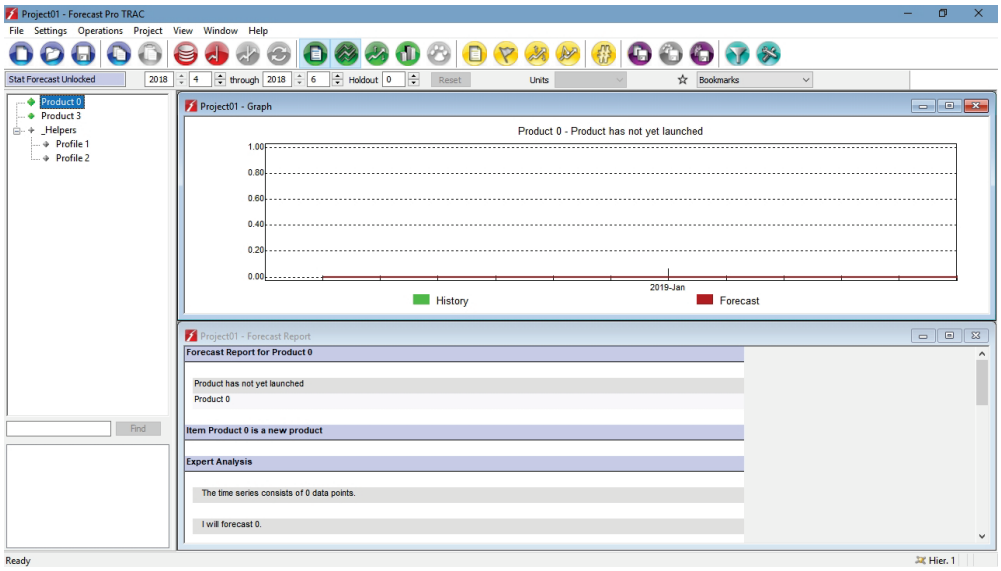

We have read in two launch profiles from *New Product- Helpers.xlsx* that we may use as analog variables, *Profile 1* and *Profile 2*. Each of these profiles contains the initial year and a half of monthly sales for a prior new product launch. Sales for our new product are likely to follow the same initial sales pattern as the product sales shown in *Profile 1*. Click on *Profile 1* to display the graph.

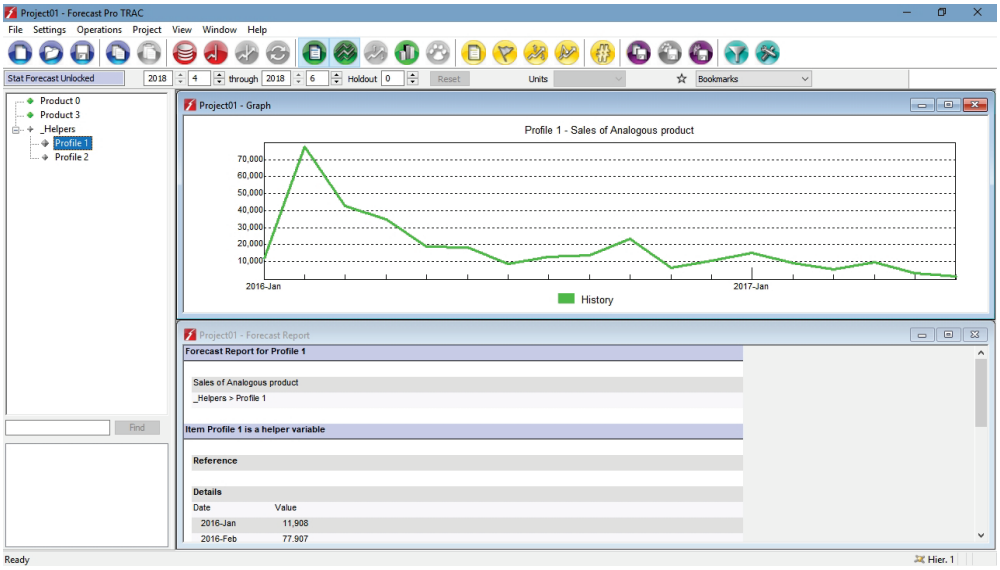

Right click *Product-0* on the Navigator to invoke the context menu and then select **Model>By Analogy**.

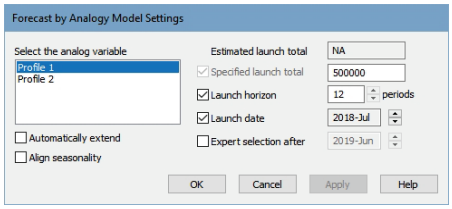

We must begin by selecting the analog variable to use. In our example there are two available. Select *Profile 1* by clicking it. Since *Product-0* does not contain any historic demand, a launch total must be specified. The default value for the Specified launch total is to calculate it using the values specified in the analog variable—in our example this yields 279,587. Let's say that our market research leads us to believe that this new product will sell 500,000 units in the first 12 months. To have our forecast reflect this assumption, edit

the Specified launch total to equal 500000. Check the boxes for Launch horizon and Launch date. Notice that the default Launch horizon is set to 12 (the length of the data plus the forecast horizon) and that the default Launch date is July 2018 (the first forecast point). Click Apply to build the model and OK to exit the dialog box.

Examine the output. Notice that the forecast now reflects the shape of the analog variable and Total forecast equals 500,000.

Now we will use the forecast by analogy method to prepare forecasts for the same product after we have 3 months' worth of history.

Right click *Product-3* on the Navigator to invoke the context menu and then select **Model>By Analogy**. Select *Profile 1* as the analog variable by clicking on it. Select Launch horizon and set the number of periods to 12. The dialog box should now look like the one below.

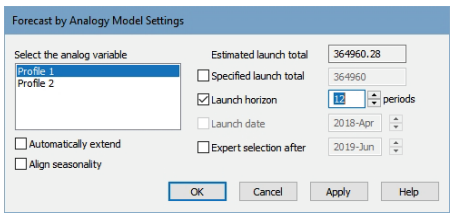

Now that we have some historic data, Forecast Pro is able to fit the Analog model and calculate the estimated launch total from the historic demand. The estimated launch total predicts that at our current rate, we will sell 364,960 units during the Launch horizon (the first 12 months in our example).

This is substantially less than the 500,000 units our market research predicted. Select Specified launch total and set it to 500000. Click Apply to build the model and OK to exit the dialog box. Right click the Graph window to invoke its context menu and toggle on Fitted to display the fitted values.

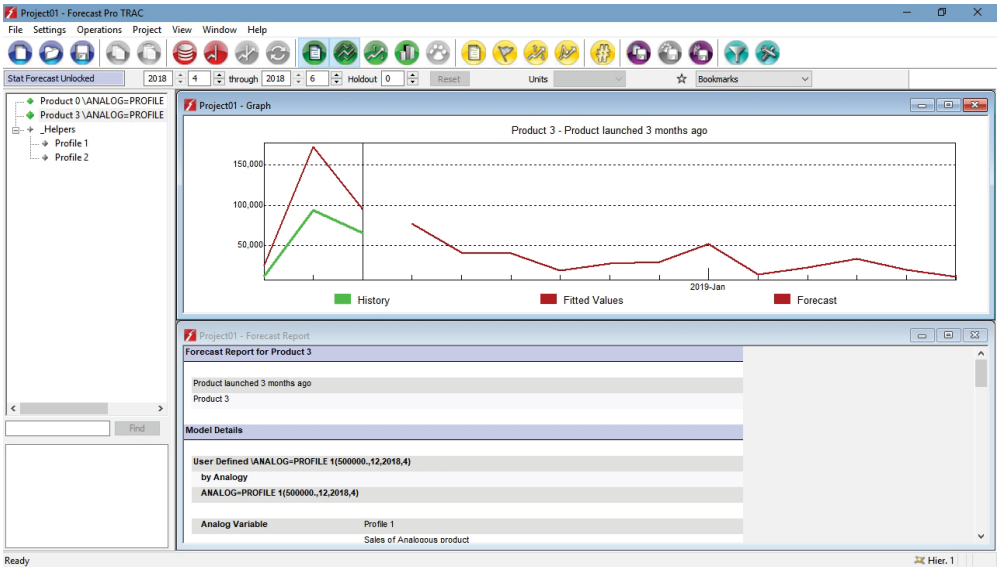

The forecast shows the sales now required to achieve the specified launch total. The fit shows the historic volume that would normally be associated with the current forecast. The green line shows the actual sales to date—which in this example is substantially lower than the fit.

Right click *Product-3* on the Navigator to invoke the context menu and then select **Model>By Analogy** to return to the Forecast by Analogy dialog box. Remove the check mark from the Specified launch total option and click OK to build the model using the estimated launch total and exit the dialog box.

This ability to compare your specified launch totals with the estimated launch totals is quite powerful and will often illustrate the need to revisit your assumptions as the actual demand comes in.

Exit Forecast Pro.

Using the Bass Model

In this exercise we'll create forecasts for a new product using the Bass diffusion model. The Bass model is most often used to forecast first time purchases of new-to-world products.

The model tries to capture the adoption rates of two types of users—innovators and imitators. Innovators are early adopters of new products and are driven by their desire to try new technology. Imitators are more wary of new technology—they tend to adopt only after receiving feedback from others.

The Bass model uses two coefficients to quantify the adoption rates. The *Coefficient of Innovation,* referred to in the literature as "p", controls the rate for the innovators. The *Coefficient of Imitation,* referred to in the literature as "q" controls the rate for the imitators.

If you have 5 or more historic data points, these coefficients can be fit to the data. To build a Bass model with fewer than 5 data points you must set the values for these coefficients along with the total number of potential adopters.

In this exercise, we will look at two examples—one where no data are available for the new product and the other where the initial 6 years of adoptions are known.

Start the program, click the red Data Manager icon, specify *Home Electronic Product.xls* on the Historic Data Row, click the Read & Forecast button and then click Exit to exit the Data Manager. Open the graph view. Your screen should now look like the one shown below.

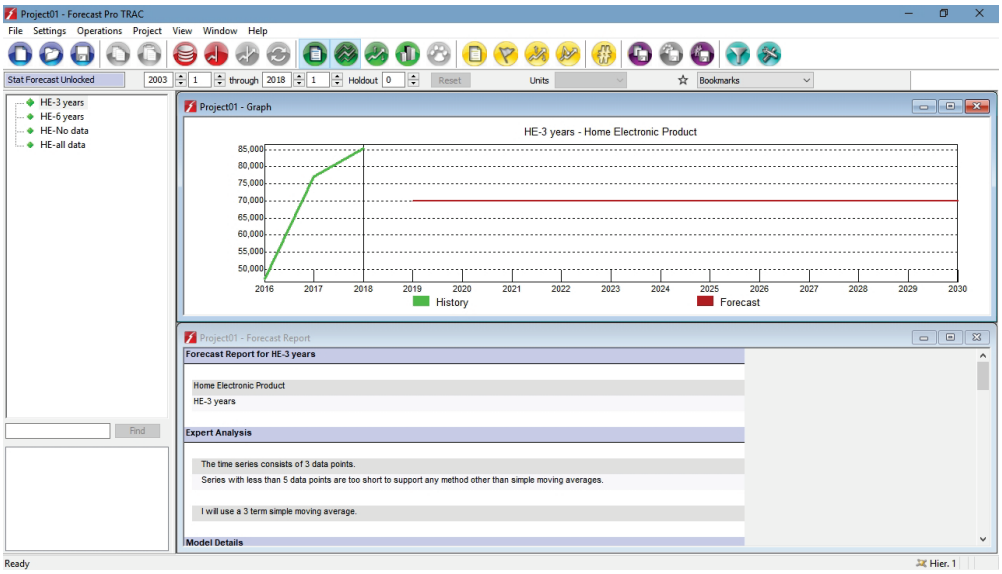

Right click *HE-No data* on the Navigator to invoke the context menu and then select **Model>Bass Diffusion**.

In this example we have no historic data so we will need to specify the Coefficient of Innovation, the Coefficient of Imitation and the Potential Market. The coefficients could be set using values from an analogous product's model. There is also a considerable body of literature on the Bass model including published coefficients for different types of technology. Consult the Forecast Pro Statistical Reference Manual for details.

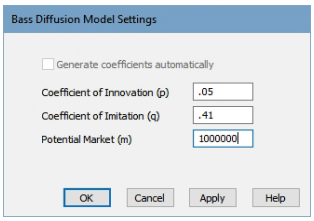

The example we are working with is a home electronic device. Our market research suggests the total market is 1 million units. A similar product had coefficients of innovation and imitation of 0.05 and 0.41 respectively. Edit the settings to match the ones above and click OK to build the model.

Click on the green Forecast Report icon ( ) to close the Forecast Report. Click the Advanced Diagnostics icon ( ) to open the Advanced Diagnostics view. Right click the Advanced Diagnostics view to invoke its context menu and select **Cumulative**. Your display should now match the one below.

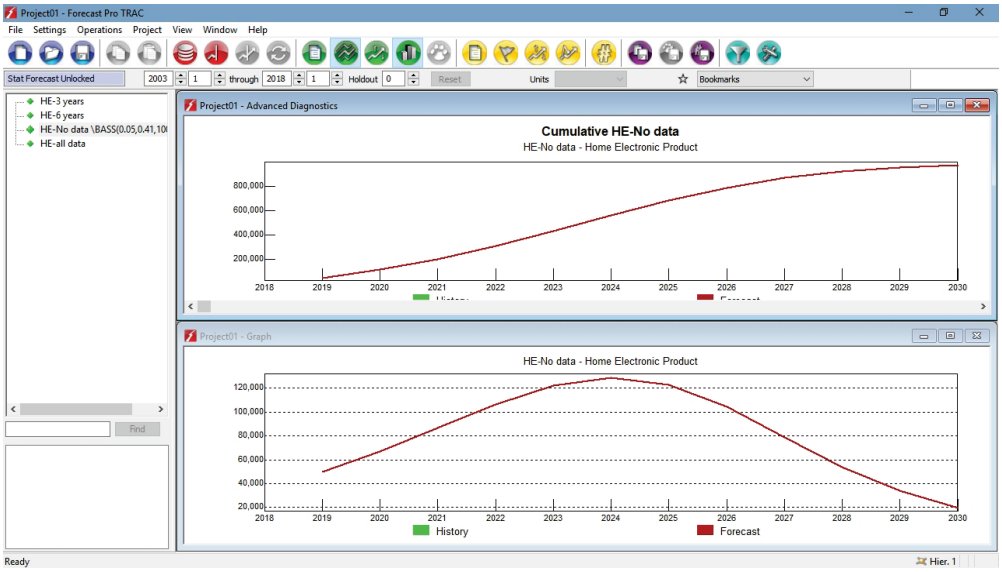

The bottom graph shows the forecasts growing for the first five years and then starting to decline. The top graph shows a graph of the cumulative basis (therefore each period displays total adoptions to date). Notice the cumulative forecasts resemble an elongated “S”. This characteristic shape is why the Bass model is often referred to as an S-curve model.

Right click *HE-6-years* on the Navigator to invoke the context menu and then select **Model>Bass Diffusion**.

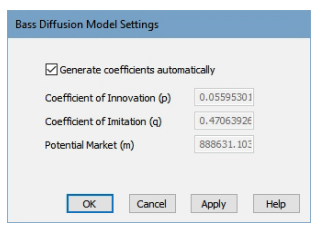

Notice that now that we have more than 5 data points, Forecast Pro can automatically fit the coefficients and estimate the potential market from the data. The estimated coefficients are pretty similar to the ones we used prior to having data, however, the estimate of the potential market is a good bit lower. Click OK to build the model.

Forecasting sales for a new-to-world product prior to launch is extremely difficult. Regardless of the forecasting method used, you should be constantly

revisiting your model assumptions as data becomes available and adjusting the model when appropriate.

This concludes the *New Product Forecasting* lesson.

Lesson 19

Building Dynamic Regression Models

In this lesson you will learn how to build dynamic regression models using Forecast Pro's Automatic Dynamics option, which largely automates the model building process, and learn how to apply a regression model specification across multiple items.

Building a Model Using Automatic Dynamics

Start the program and click the red Data Manager icon ( ) to call up the Data Manager. Click the Add drop down on the Historic Data row, select Excel and select *Electricity – Historic Data.xlsx* to add it to the Historic Data row. Click the Add drop down on the Explanatory Variables row, select Excel and select *Electricity – Explanatory Variables.xlsx* to add it to the Explanatory Variables row

Click the Read & Forecast button to read in the historic demand data and generate the statistical forecasts.

Click the Exit button to exit the Data Manager.

Right click on the Navigator to invoke its' context menu and select **Expand All** to fully expand the Navigator. Open the Graph view by clicking its green

icon ( ). Turn off the Forecast Report view by clicking the Forecast Report icon ( ). Your display should now match the one shown below.

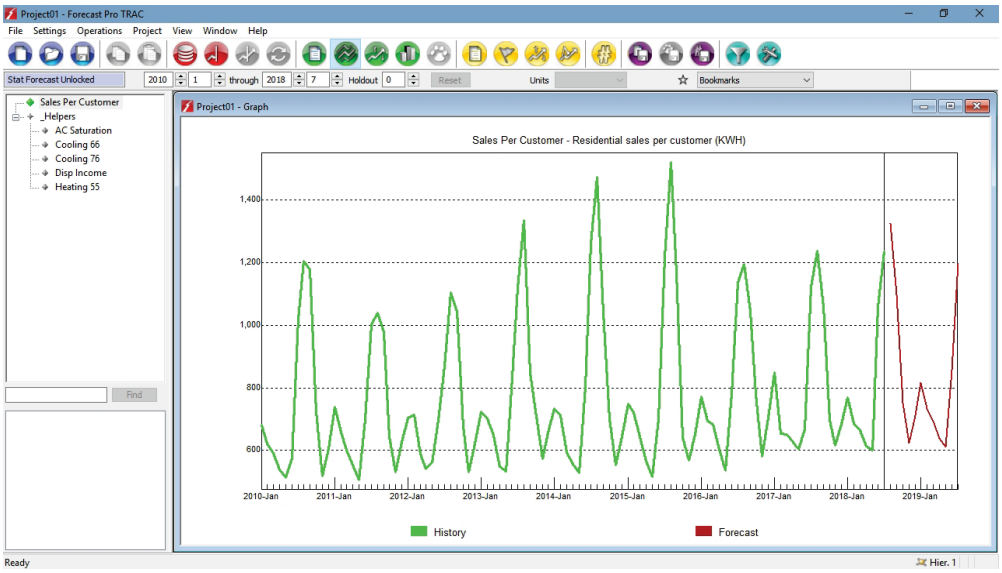

The data series for this lesson consist of:

Sales per Customer	Residential electricity sales (KWH) per customer.
AC Saturation	Air conditioning saturation rate.
Cooling 66	Cooling degree days at base temperature 66 degrees.
Cooling 76	Cooling degree days at base temperature 76 degrees.
Disp Income	Disposable income per household.
Heating 55	Heating degree days at base temperature 55 degrees.

Sales per Customer is the dependent variable (the one we will be forecasting). The remaining variables are explanatory variables (also called independent variables). Cooling 66 and Cooling 76 are measures of summer heat and Heating 55 is a measure of winter cold.

Our first step is to examine the dependent variable graphically. Notice the strong seasonal patterns in the data and the varying heights of the summer and winter peaks. Much of the variation is due to air-conditioning and heating, which is primarily driven by the temperature. We will try to capture this relationship by using the weather variables.

Right click on the graph to invoke the graph's context menu, select **Graph Settings** and select the Items tab. This tab displays the variables available on the Navigator and allows you to graph up to five of them together.

Select Sales per Customer and Cooling 66, uncheck the Include forecasts option (this turns off the statistical forecasts for the dependent variable—Sales per Customer in this case), and click OK.

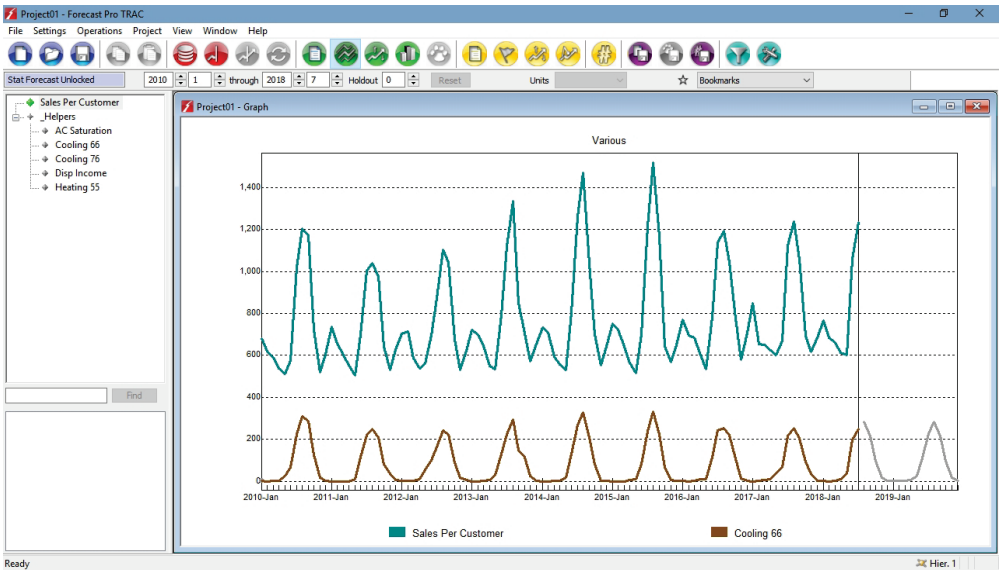

Notice that the peaks in Cooling 66 correspond to the summer peaks in Sales per Customer. This suggests that Cooling 66 is a powerful explanatory variable.

Notice also that forecasts are supplied for Cooling 66. To generate a forecast for your dependent variable your explanatory variables needs to be forecasted. Normally, you will want to supply appropriate forecasts for each explanatory. This will generate forecasts for your dependent variable using the forecast scenarios you've provided for the explanatory variables. If you do not supply forecast scenarios for all of your explanatory variables (or if your forecast scenarios do not cover the entire forecast horizon), you can use Forecast Pro's *Auto Extend* functionality to generate the required forecasts using expert selection. *Auto Extend* is a project-level option that is found in the Advanced Controls tab of the Options dialog box.

Try graphing some of the other variables to see how they relate to the Sales per Customer data. Keep in mind that you can use the *Relative scaling* option to view variables with very different scales.

When you are finished exploring the data, click on one of the explanatory variables on the Navigator to revert to the default graph mode and then click on Sales per Customer. Right click the graph to invoke its context menu and select Fitted to toggle on the fitted values. Your screen should now match the one below.

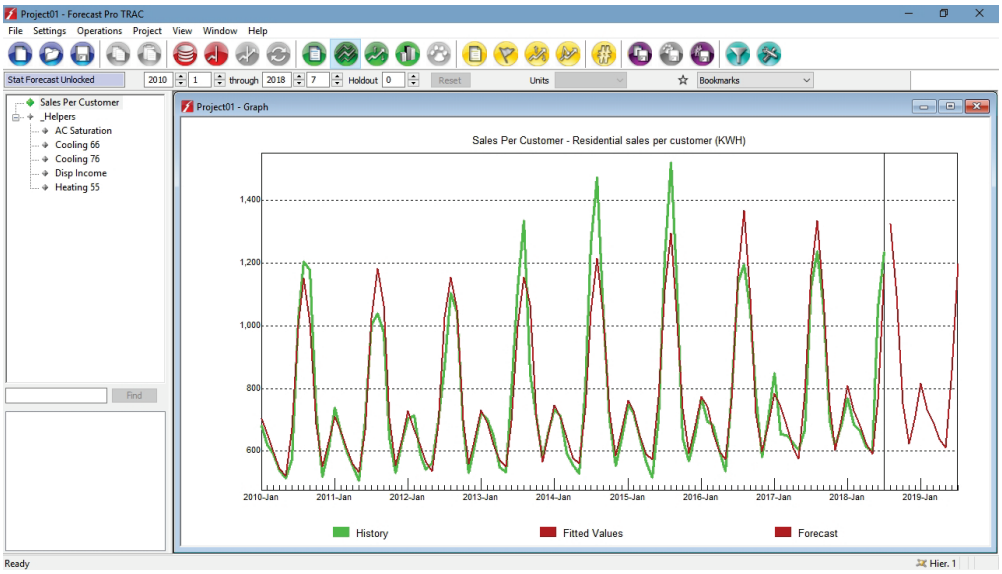

Open the Forecast Report by clicking the Forecast Report icon ( )

The current forecast was generated using a seasonal exponential smoothing model chosen via expert selection (expert selection does not consider dynamic regression). Notice that the adjusted R-square for this model is 0.89 and that although the model is capturing a seasonal pattern it is not capturing the amplitude of the seasonal peaks very well. These peaks are driven by weather conditions and we will be able to capture them much more effectively using dynamic regression.

Right click on Sales per Customer on the Navigator to invoke its context menu and select **Dynamic Regression>Manage**. This will invoke the Dynamic Regression dialog box shown below:

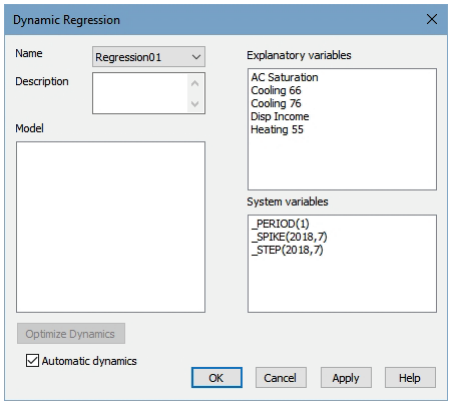

There are two types of variables that can be included in a dynamic regression models—explanatory variables and dynamic terms. Explanatory variables allow you to capture how the dependent variable changes in response to external variables. The Explanatory variables box lists all the variables read in with the Data Manager. The System variables box lists explanatory variables that can be defined in the program as well as dynamic terms. Dynamic terms capture how the dependent variable changes in time and can help forecast trends and seasonal patterns in a similar fashion to extrapolation methods such as Box-Jenkins and exponential smoothing.

Forecast Pro offers an Automatic dynamics mode whereby you specify the explanatory variables you wish to use and the program automatically determines the dynamic terms to include. When you use Automatic dynamics, the dynamic terms are not visible in the Dynamic Regression dialog box. This lesson focuses on building models with Automatic dynamics. Please see the reference section of this manual for details on building regression models without Automatic dynamics.

Building a regression model often involves a sequence of rounds. You posit an initial model, diagnose it and either accept it as final or make a change to improve it. Whenever you make a change, you generate a new model to either accept as final or make an additional change. The process continues until you are satisfied with the model.

There are many ways to build a regression model. Some people begin with a model that includes many terms and weed out the terms that aren't statistically significant. Others begin with a simple model and add variables one by one. We will take the former approach.

We'll begin by building a model which includes all of our explanatory variables. Double click AC Saturation in the Terms list box and notice that this adds the variable to the current model specification. Double click Cooling 66, Cooling 76, Disp Income, and Heating 55 to add them to the model specification.

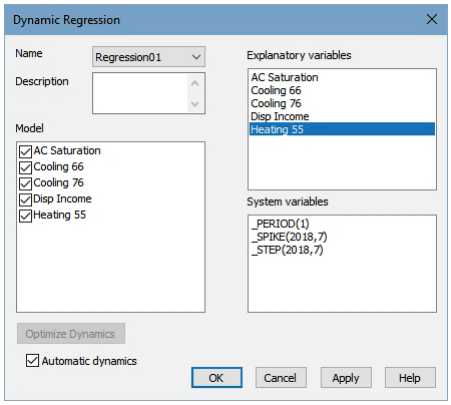

Your Dynamic Regression dialog box should match the one above. Notice that the Automatic dynamics option is selected. Click OK to build the model.

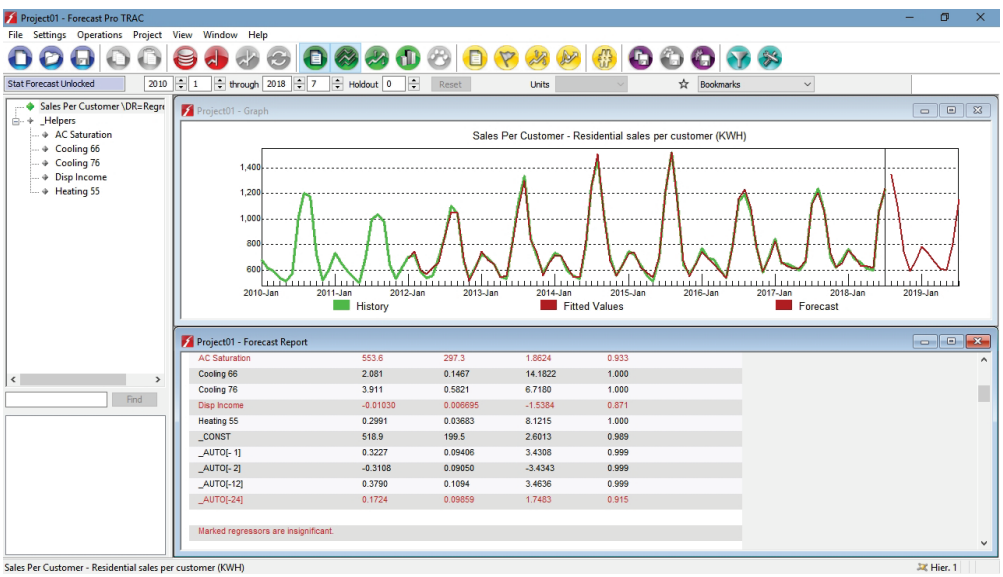

Take a look at the graph. Notice that the regression model fitted values capture the seasonal peaks much better than the seasonal exponential smoothing model. Close the Graph view by clicking its icon ( ).

Notice that the model includes the five explanatory variables that you specified as well as five dynamic terms. The dynamic terms consist of a constant (_CONST) and four autoregressive error terms (_AUTO). Models that include autoregressive error terms are referred to as Cochrane-Orcutt models. You can learn more about these models in the *Forecast Pro Statistical Reference Manual*.

Notice that two of our explanatory variables appear in red. This indicates that their T-statistics are insignificant and you should consider removing them from the model. A best practice when making a change to a regression model is to make only one change at a time. Let's begin by removing Disp Income which is the least significant of the flagged variables.

Right click on the Disp Income row on the forecast report to invoke the forecast report's context menu. Select **Remove 'Disp Income' from model**. Notice that this builds a new model which does not include Disp Income.

AC Saturation is still marked in Red, indicating that it is not significant. Remove AC Saturation from the model by right clicking on it and selecting **Remove 'AC Saturation' from model** on the context menu. Open the Graph view by clicking its green icon ( ). Your screen should now match the one below.

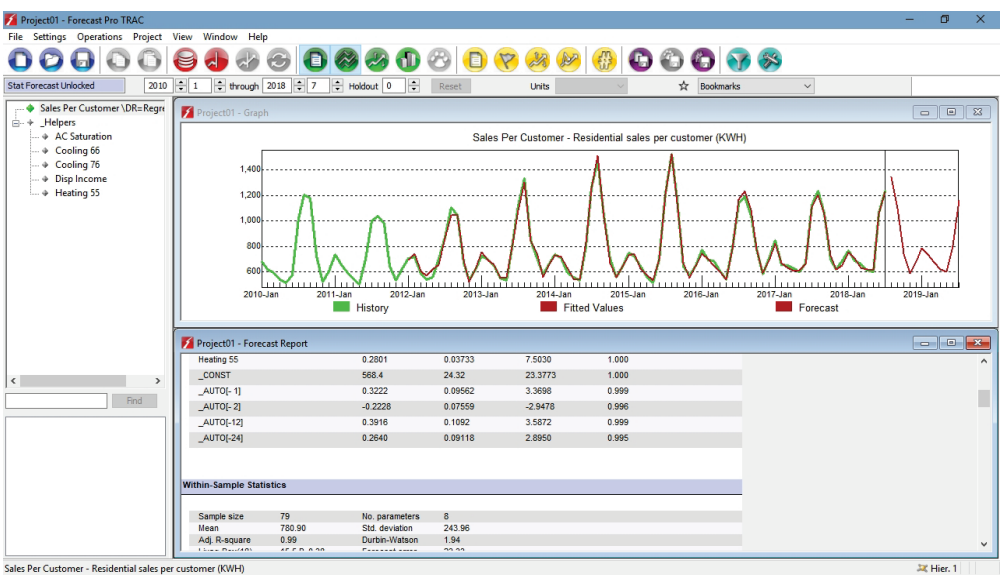

This dynamic regression model is substantially better than the exponential smoothing model selected by expert selection. All terms are significant, the adjusted R-square is 0.99 and the model captures the changing amplitudes of the seasonal peaks.

Exit Forecast Pro.

Appling a Model Specification to Multiple Items

When forecasting items that are similar, you may find that the same dynamic regression model specification works well across the different items. In this section, we will specify a price and promotion model for a specific market-level breakfast cereal Stock Keeping Unit (SKU) and then apply the same model specification to other breakfast cereal SKUs.

The data we will be using are weekly. Start the program and click the light blue Options icon ( ). On the basic tab, change the forecast horizon to 52 and click OK to close the Options dialog box. Click the red Data Manager icon ( ) to call up the Data Manager. Click the Add drop down on the Historic Data row, select Excel and select *Cereal - Historic Data.xlsx* to add it to the Historic Data row. Click the Add drop down on the Explanatory Variables row, select Excel and select *Cereal - Item Level Explantory Variables.xlsx* to add it to the Explanatory Variables row.

Click the Read & Forecast button to read in the historic demand data and click the Exit button to exit the Data Manager.

Right click on the Navigator to invoke its' context menu and select **Expand All** to fully expand the Navigator. Select *Cereal>C1 - Kids>Booth's>BCT-013*. Open the Graph view by clicking its green icon ( ). Turn off the Forecast Report view by clicking the Forecast Report icon ( ). Your display should now match the one shown below.

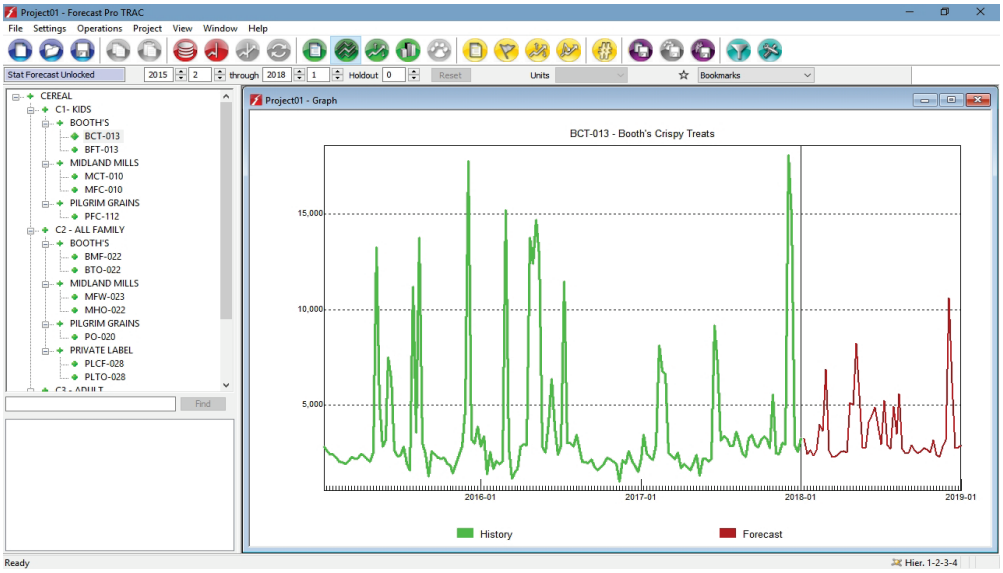

The data consists of weekly cereal sales for multiple sub-categories, brands and SKUs.

Forecast Pro allows you to specify two types of explanatory variables—global and item-specific. A global explanatory variable consists of a single time series (set of values) which can be included in a dynamic regression model for any item on the Navigator. An item-specific explanatory variable consists of a set of time series each of which is associated to a specific item on the Navigator.

To illustrate how this works, let's consider our current example where we will be forecasting different cereal SKUs and using Price as an explanatory variable. If the price was the same across all SKUs we could define Price as a global explanatory variable. However, if the price varied by SKU we would want to define price as an item-specific explanatory variable to allow us input different prices for different SKUs.

In the electricity example we used global explanatory variables and you may have noticed that they were displayed in the Navigator. Because the values for item-specific explanatory variables are tied to specific items they are not included on the Navigator.

In this example, we are using item-specific explanatory variables. Right click on the graph to bring up the graph context menu and select *Graph Settings*.

Click on the Add or Remove button to bring up the Add or Remove dialog box. On the filter dropdown, select Explanatory variables.

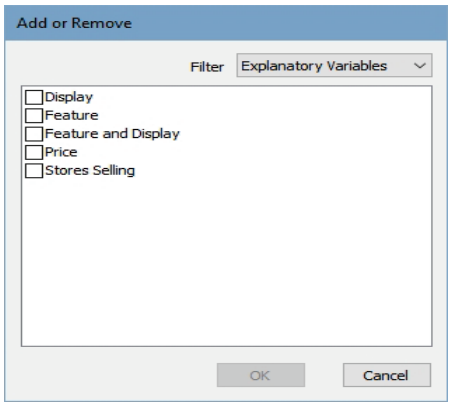

The Add or Remove dialog box now shows the five explanatory variables found in *Cereal - Item Level Explanatory Variables.xlsx.*

Display	Percent of stores with item on display (e.g. End Cap), not featured in weekly ad circular
Feature	Percent of stores with item featured in weekly ad circular
Feature and Display	Percent of stores with item on feature and display
Price	Average item price
Stores Selling	The number of Stores selling the item

Select Price and click OK to return to the Graph Settings dialog box. Select the Relative Scaling box for price in the Graph Settings dialog box and press OK.

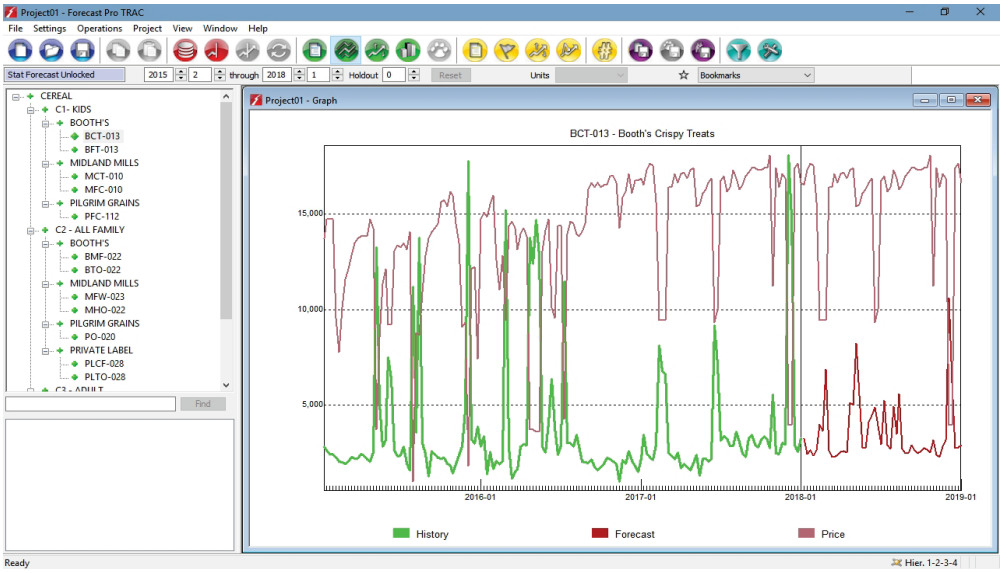

The graph now displays both weekly sales and price for Booth's Crispy Treats. Notice that many of the sales peaks coincide with a decline in price.

In this example, explanatory variables have been defined for all items. Select *Cereal* on the Navigator and notice that Price is shown for total Cereal as well. Look at some of the other items in the Navigator as well. Notice that sales increase when price declines for most items.

Use the Graph Settings dialog box to explore the other explanatory variables. While you can display multiple explanatory variables simultaneously, it is usually easiest to observe the relationship between a given explanatory variable and sales if you look at a single explanatory variable at a time. Be sure to use the Relative Scaling checkbox so that you can clearly see if changes in the displayed explanatory variable correlate with sales increases or declines.

When you are done exploring the data, re-open the Graph Settings dialog box, turn off the explanatory variables and turn on the Fitted Values. Click Apply. Your display should now match the one shown below.

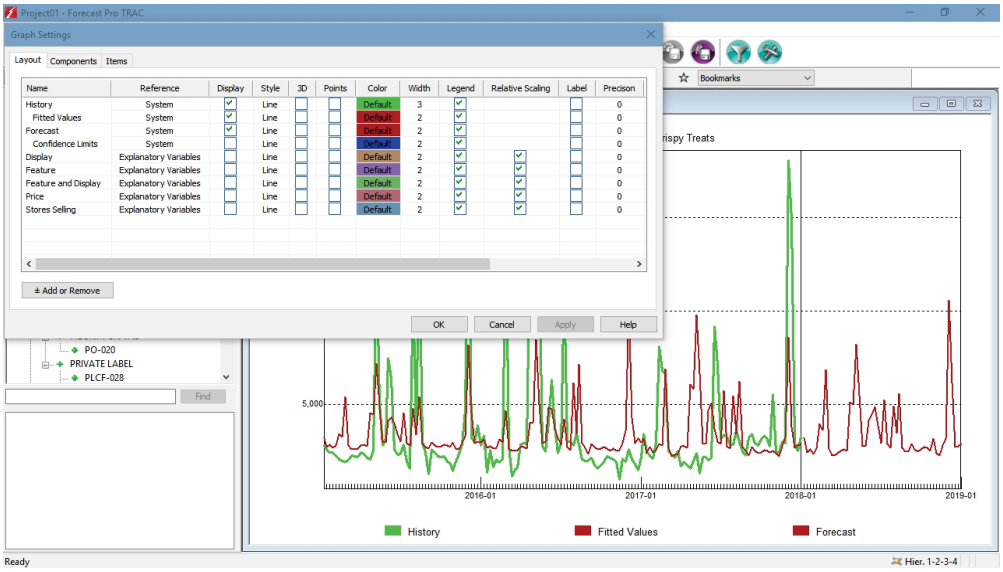

Click OK to close the Graph Setting dialog box. Open the Forecast Report by clicking the Forecast Report icon ( ). Scroll down and observe that the adjusted R-square for the expert selection model is 0.23.

We will now build a regression model for Booth's Crispy Treats. Price and Promotion models are often multiplicative. That is, changes in the explanatory variables drive percent changes in the dependent variable. We build a multiplicative model by taking a log transform of the dependent variable and, optionally, non-negative explanatory variables. To build a multiplicative model, right click on *Cereal>C1 - Kids>Booth's>BCT-013* in the Navigator and select Custom Modifier. In the Custom Modifier box, enter "Pow=Log". We will show how to take log transforms of the explanatory variables below.

Right click on *Cereal>C1 - Kids>Booth's>BCT-013* to bring up the Navigator's context menu again and select Dynamic Regression> Manage.

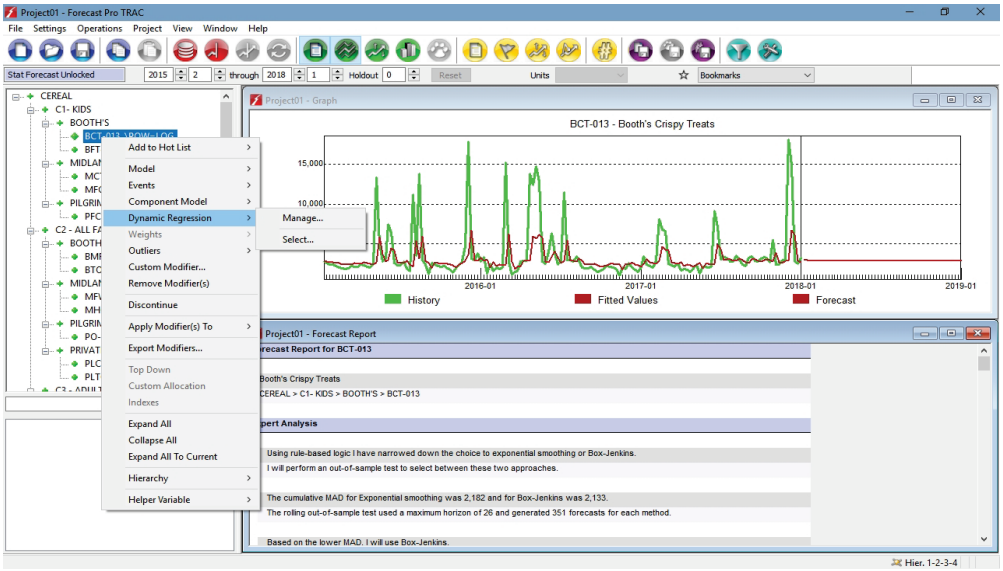

In the Dynamic Regression dialog box, double click on each of the explanatory variables to add them to the model. Click on the Name dropdown menu and select Save As. Enter “Price and Promotion” into the Save As dialog box and click OK. Right click on Price in the Model variable list box and select **Transform>Logarithmic**. Repeat this step for Stores Selling. The Dynamic Regression dialog box show look like below.

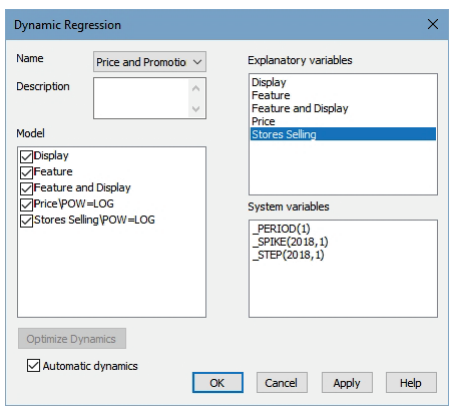

Click OK to apply the model and close the Dynamic Regression dialog box. Scroll down to the Within-Sample Statistics in the Forecast Report. Note that the adjusted R-squared has increased from 0.23 to 0.81, and the graph shows a much better fit. Scroll back up to show the regression coefficients.

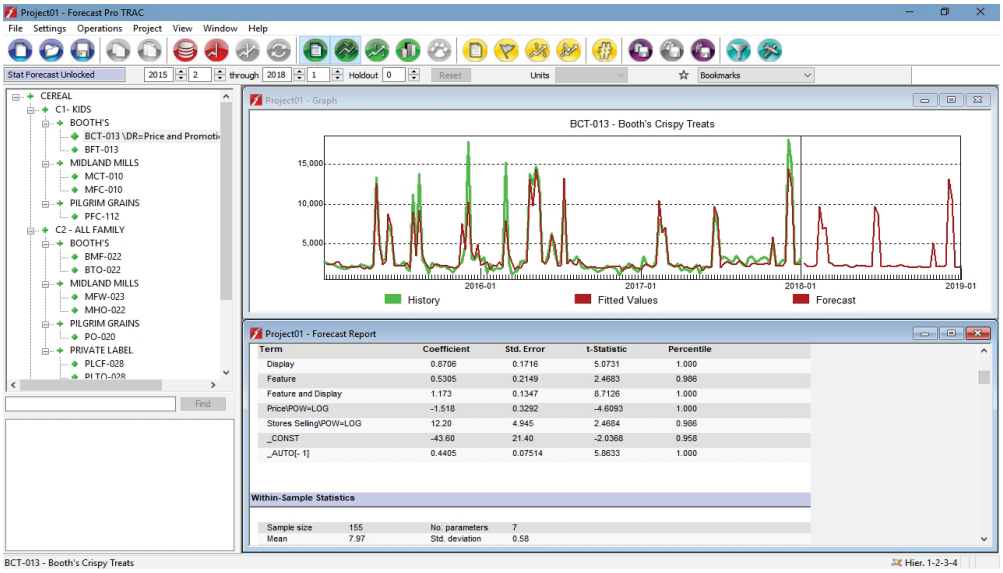

Note that all the explanatory variables are shown in black, meaning that they are significant, and the coefficients all have the expected sign, with features, displays and stores selling driving sales and price increases having a negative impact on sales.

Let's apply this regression model to all items in the project. Right click on *Cereal>C1 - Kids>Booth's>BCT-013* to bring up the Navigator's context menu and select **Apply Modifier(s) To>This Level**. We have now built price promotion models for all the end items.

Click on a few end-items to view the different price and promotion models. Note that while the explanatory variables are the same in each model, the dynamic terms are customized to the end item. Select *Cereal>C1 - Kids>Booth's>BFT-013* and compare to the model for *Cereal>C1 - Kids>Booth's>BCT-013*.

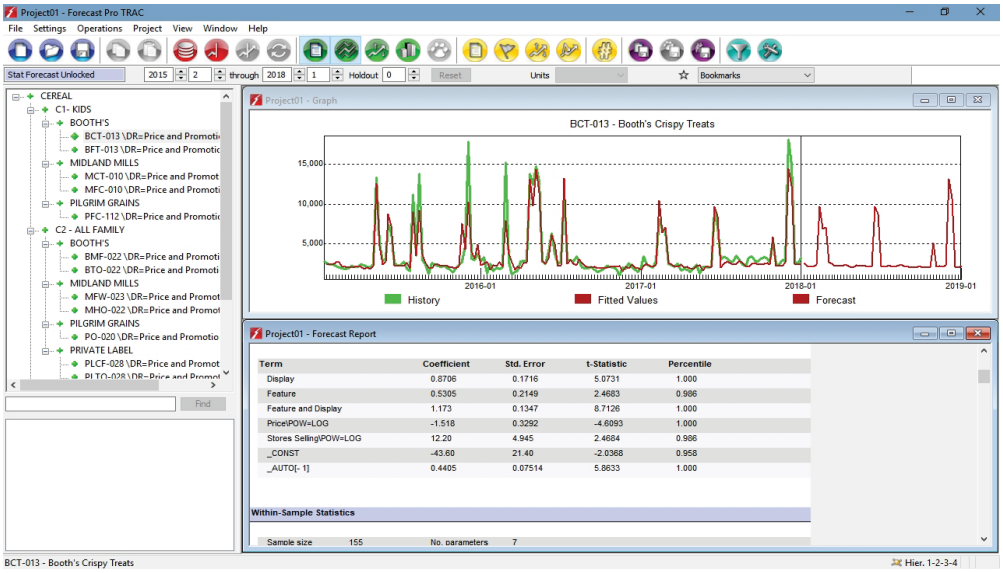

The model that we looked at above for *Cereal>C1 - Kids>Booth's>BFT-013* included AUTO[-1], while the model below for *Cereal>C1 - Kids>Booth's>BCT-013* includes AUTO[-1] and AUTO[-52].

Exit Forecast Pro.

This concludes the *Building Dynamic Regression Models* lesson.

Lesson 20

Using Weighting Transformations

Forecast Pro includes a weighting transformation which will divide your historical data by user-defined weights, forecast the resultant (deweighted) series and then multiply (reweight) the forecasts. This procedure can be useful in a variety of situations including adjusting for the number of working days in a month, defining a seasonal pattern and supplying a growth curve for a new product.

In this section we will use the weights procedure to adjust for the number of weeks in each month. The section will conclude with a discussion of how to apply weights in other situations.

Adjusting for 4 vs. 5 Week Periods

It is not uncommon for companies planning production or scheduling to divide the year into twelve periods each consisting of either four or five complete weeks. This is often referred to as a *4-4-5* calendar. In situations where the forecasts ultimately need to be broken down by week, this method may be preferable to using actual calendar months which include partial weeks.

If the number of weeks in each period is consistent from year to year (i.e., period 1 always has 4 weeks, period 2 always has 4 weeks, period 3 always has 5 weeks, etc.) then the variation due to the number of weeks in each period will be captured as part of the seasonality and no special action is required. In cases where the number of weeks in each period varies from year to year (i.e., some years period 1 has 4 weeks other years it has 5) a weighting transformation can be used to capture the variation due to how many weeks are in each period.

The procedure is very straightforward. You construct a weighting variable that classifies each period of the historical data and forecast period as containing either 4 or 5 weeks. Forecast Pro divides the historical data by the weights. (This converts sales per period into sales per week per period.) This deweighted series is forecasted. The forecasts are then reweighted by multiplying by the corresponding weights.

The data for this exercise are contained in two Excel files. *4-4-5 – Historic Data.xls* contains five time series that were collected using twelve 4 or 5 week periods per year. The number of weeks per period is not consistent from year to year. *4-4-5 - Helper.xls* is a "helper file" and contains the *helper variable* "Weeks".

Helper variables can be used as weighting variables, event schedules and analogy series. In this example the helper variable Weeks contains the number of weeks in each period. It covers the historical span as well as the forecast period.

Start the program, click the red Data Manager icon, specify *4-4-5 – Historic Data.xls* on the Historic Data Row and then specify *4-4-5 – Helpers.xls* on the Helpers row. Click the Read & Forecast button, click the Refresh button on the Helpers row and then click Exit to exit the Data Manager. Open the graph view, right click the Graph window to invoke its context menu and toggle on Fitted to display the fitted values. Your screen should now look like the one shown below.

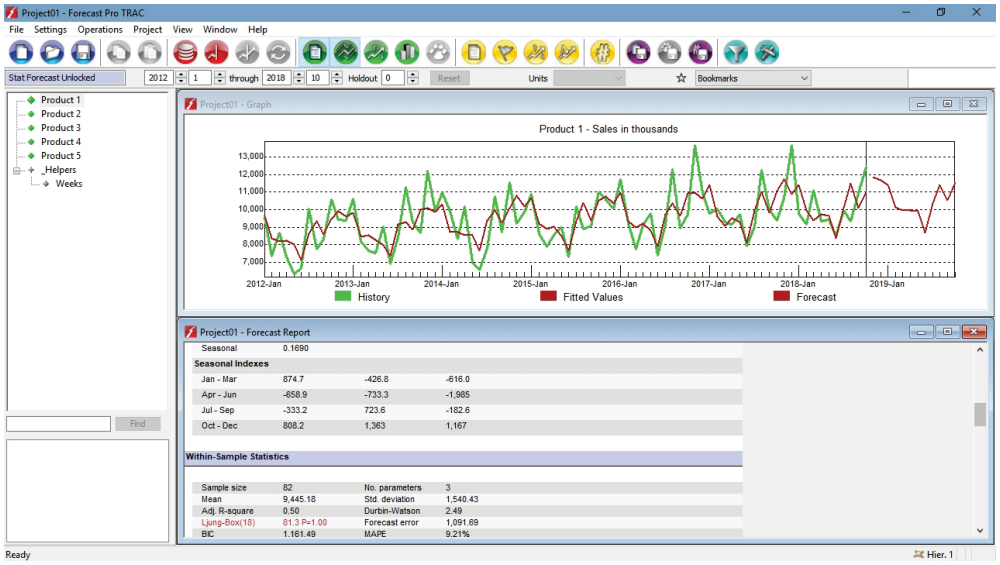

Notice that a seasonal model is selected and that the adjusted R-square is 0.50.

Right click on Product 1 to call up the Navigator's context menu. Select Weights and then select Weeks.

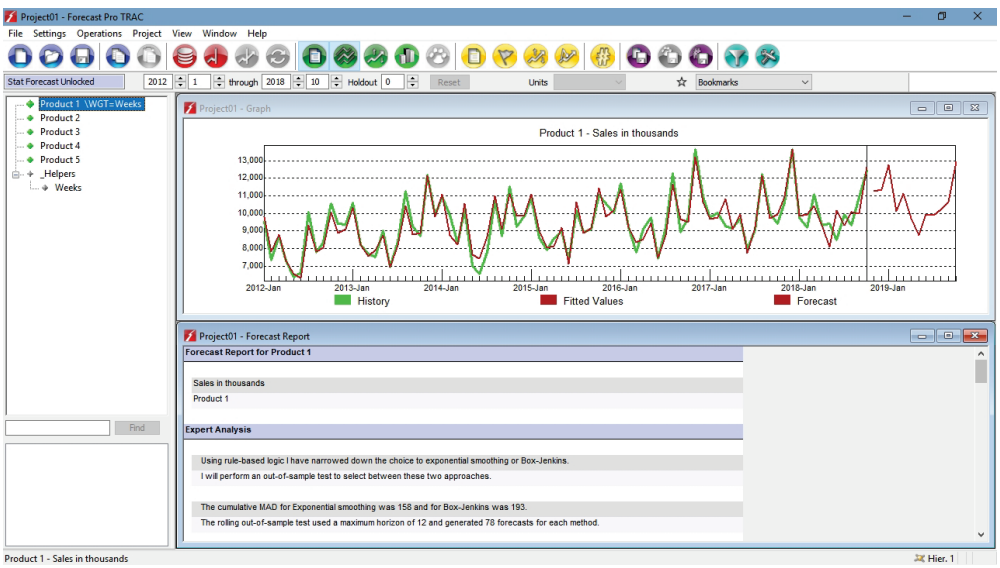

The second model's superior fit is easily discernible from a visual inspection of the graph. Notice that a seasonal model is selected and that the adjusted R-square is now 0.82.

Right click on Product 1 to call up the Navigator's context menu. Select **Apply Modifier(s) To>All Items**. Since the other four products all use the same calendar, applying the weighting transformation improves the model in all cases.

Other Uses for Weighting Transformations

The weighting transformation is useful in a wide variety of situations. This section briefly notes some of them. Its purpose is to indicate some directions you may want to take with your own business data.

Trading day corrections. Many businesses are sensitive to the number of working days per period. Consider a service provider who is closed on the weekends. The number of working days in January (and all other months) will vary from year to year depending on how many weekend days happen to fall in any given month. If the number of working days has an impact on sales, then it needs to be accounted for in the model. A simple solution would be to use a weighting transformation where the weights consist of the number of working days per month.

User-defined seasonality. At times you may wish to supply your own estimate of the seasonal pattern rather than trying to extract it directly from the data. This might be desirable if the data were short or very noisy. The weighting variable would consist of seasonal multipliers for the series.

Product phase outs and other forecast adjustments. There may be times where you wish to alter the statistical forecasts using a weighting variable rather than the forecast adjustment facility. For example, let's say that you plan on discontinuing a product and wish to use the statistical forecast until the product is discontinued. You could create a weighting variable that consists of all ones during the history and the forecast periods prior to the discontinuation date and equals zero for all periods thereafter. If the product would be phased out over a three month period rather than ending abruptly, then you could use weights like 0.75, 0.5 and 0.25 during the phase out period.

This concludes the *Using Weighting Transformations* lesson.

Lesson 21

Out-of-sample Testing

A good deal of the empirical knowledge about forecasting has come from comparisons of different methodologies. The M-Competition (Makridakis et al. [1982]) and M-3 Competition (Makridakis and Hibon [2000]) are the largest and most famous of these comparisons. Forecast Pro participated in the M-3 competition and outperformed all other software entrants and 16 out of 17 academic teams.

The rather simple comparison methodology for the original M-Competition was as follows.

The researchers assembled a collection of 1001 time series of yearly, quarterly and monthly data. The data were obtained from microeconomic, industry-level, macroeconomic and demographic sources. Twenty forecasting methods were tested for the entire sample of 1001 time series and three on a subset of only 111 time series.

A sample of time points (6 for annual series, 8 for quarterly, 18 for monthly) was held out from the end of each time series. Each forecast model was fitted to the remaining data and used to forecast the values of the holdout sample. The forecasts were then compared to the withheld data, and errors computed for each horizon, each time series and each forecast method. The errors were then summarized and analyzed in a variety of ways.

The most significant weakness in this methodology is that it uses only one forecast base for each time series, the last point in the fitting sample. One obtains only a "snapshot" of performance from one point in time. A forecast base just before or after a dramatic event in the data may completely change the results. Furthermore, you obtain only one forecast error for each horizon time from 1 to the end of the fit set. This procedure is referred to as a *static* evaluation.

Forecast Pro implements both a static and a *rolling-base* evaluation. The rolling-base procedure begins in the same way. However, after the forecasts have been made, the model is rolled forward by one period. Forecasts are then made from the new base to the end of the withheld data. This process is repeated until the withheld data sample is exhausted. If 6 data points have been withheld, then you obtain 6 1-step forecasts, 5 2-step forecasts, 4 3-step forecasts, etc.

The model coefficients are *not* reestimated as each additional data point is assimilated. The forecast model is based entirely upon the original fit set.

Defining a Holdout Sample

Forecast Pro is shipped with data files containing the 111 series from the original M-competition. *M-data Annual.xls* contains 20 annual series, *M-data Quarterly.xls* contains 23 quarterly series and *M-data Monthly.xls* contains 68 monthly series.

In this lesson we will compare the out-of-sample performance of two different forecasting techniques using the monthly data set. The main idea of this facility in Forecast Pro is to give you the opportunity to test options for time series from your own company. This lets you tune Forecast Pro to your data.

Start the program, click the red Data Manager icon, specify *M-data Monthly.xls* on the Historic Data Row, click the Read & Forecast button and then click Exit to exit the Data Manager. Now set the holdout sample on the dialog bar to 6 and click the red Forecast icon to create the forecasts.

Examining the Analytic Output

After processing is complete, select the group *Monthly* on the Navigator and view the *Out-of-Sample Rolling Evaluation for Items in Group* section of the Forecast Report view.

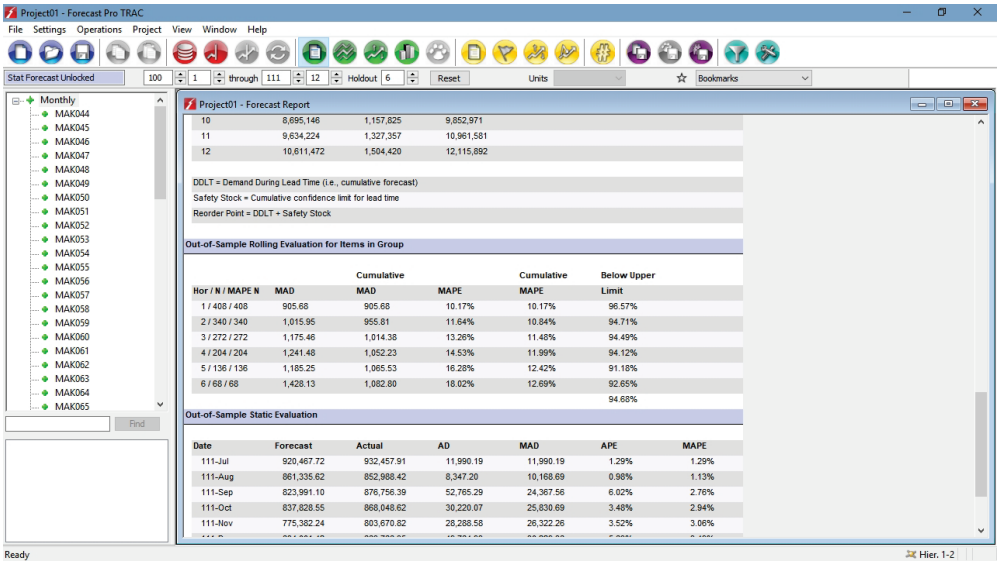

The tables display the following information. Means are always taken over all the variables forecasted for a particular display line.

Hor is the forecast horizon.

N is the number of cases used to compute all statistics for this horizon other than the MAPE (see below).

MAPE N is the number of cases used to compute the MAPEs for this horizon. This can be less than the total number of forecasts for the horizon when some of the actuals are zero, thus preventing computation of the MAPE.

MAD is the Mean Absolute Deviation for the horizon, over all series on the current display line.

Cumulative MAD is the cumulative MAD for all horizons up to and including the current horizon.

MAPE is the Mean Absolute Percentage Error for the horizon, over all series on the current display line.

Cumulative MAPE is the cumulative MAPE for all horizons up to and including the current horizon.

Below Upper Limit is the fraction of times that the actual was at or below the upper confidence limit as set in Settings Options. This allows you to calibrate the empirical upper limit to the theoretical upper limit. This is useful when you are using the upper limit to set stocking levels.

Advanced statistics such as Mad/Mean ratio, Symmetric MAPE and the Geometric Mean Relative Absolute Error can also be displayed if you customize the forecast report.

Figures listed in the row below the last forecast horizon refer to all horizons. Each group on the Navigator will produce such a report.

Comparing an Alternative Model

Right click *Monthly* on the Navigator to display the context menu. Select **Model>Very Simple Models>Simple Moving Average** and set the number of terms to 3 and click OK to build the model. Then right click *Monthly* a second time and select **Apply Modifier(s) to>All items**. We have now forecasted all items using a 3-month simple moving average.

Select the group *Monthly* on the Navigator and view the *Out-of-Sample Rolling Evaluation for Items in Group* section of the Forecast Report view.

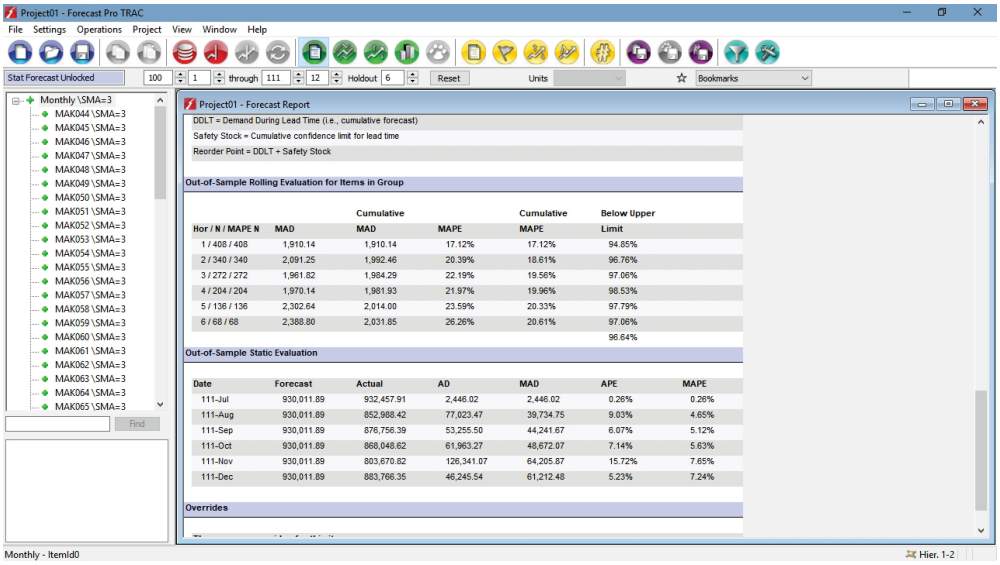

There are several ways to compare the results of this forecast and our expert selection forecast. The most common approach is to look at the MAPEs (Mean Absolute Percent Error) for the two. In our example, the MAPEs for expert selection are lower than the MAPEs for the moving averages for every time horizon—clear evidence that expert selection outperforms the moving averages.

This concludes the *Out-of-Sample Testing* lesson.

4

REFERENCE

Chapter 1

Setting Up Your Historic Data

Forecast Pro requires historic data to be input in a specific format. For many users, the data format is an initial stumbling block on the way to success in forecasting their own data. This usually happens when a user misunderstands some of the details of the data format. This chapter explains the data formats supported by Forecast Pro, advises you how to select a format and describes how to create the files. If you read the material carefully, and examine the sample data files that are shipped with Forecast Pro, you will be up and running quickly.

The first two sections describe data requirements and the different data file formats. After you have selected your file format, consult the appropriate section for instructions on how to create the files.

Overview

Forecast Pro works with historic time series, i.e., sequences of values of a variable at some time interval. Business series are usually expressed in terms of the calendar, and can be daily, weekly, monthly or quarterly. Forecast Pro can work with any type of time series but is oriented towards weekly, monthly and quarterly calendar data.

Collection and maintenance of reliable historic data is up to you and your IT department and can be implemented in many different ways, depending on information flow conditions in your organization.

Data length

Forecast Pro works by fitting a statistical model to your historic data and extrapolating it via the fitted model. Thus, your data must be long enough to provide reasonably stable estimates of the most important features of the data. Very short or very noisy historic records usually yield very simple models because the data are too short to support statistical estimates of important features like seasonality.

If the data are very short, say *four points or fewer*, Forecast Pro can pick up neither seasonality nor trend and reverts to the Simple Moving Average model.

For *more than four points but less than two years' worth of data*, Forecast Pro can fit and forecast trends but not seasonality. If your data are in fact nonseasonal, your forecasts are likely to be adequate. If your data are in fact seasonal, the forecasts are likely to be poor—Forecast Pro cannot extract or forecast the seasonality.

However, seasonal forecasts from short data sets are feasible using some of Forecast Pro's customized approaches. For example, when the short data sets are nested within aggregate product groups with longer histories, top-down forecasting can be used. If there are other longer series which exhibit the same seasonal pattern then the pattern can be estimated from the longer data set and applied to the short data sets using the INDEXES modifier, a custom component model or the forecast by analogy model.

Seasonality can be estimated and forecasted *from two to three years of data,* but this amount of data is marginal, especially when your data are noisy or intermittent. Patterns in the noise may be mistaken for seasonality, yielding inappropriate "seasonal" forecasts.

Robust capture of seasonality requires *three or more years* of data. *Four to seven years* is even better, since there is more information from which the program can separate seasonality and trend from the noise.

There is little additional payoff in accuracy *beyond about seven years* of data, and the cost in computer time can be substantial.

Missing values and zeros

Each of the statistical forecasting methods requires an unbroken stream of historic values. Unfortunately, missing values are common in business data. You therefore need a well-formulated approach to missing data. There are two parts to the approach. You must first decide how to encode missing values in the input data file. Then you must decide how these missing values are to be treated by Forecast Pro.

The encoding of missing values depends upon the data format. However, one problem is common to all of the formats—the distinction between missing values and zeroes. Forecast Pro ordinarily considers zeroes to be actual data values and treats them as such. If they were really intended as missing values, then the forecasts can be badly biased.

Forecast Pro treats a sequence of leading zeroes as missing values if you check "Ignore leading zeroes" on the Formats tab of the **Settings>Options** dialog box, and it treats trailing zeroes as missing if you check "Ignore trailing zeroes". If you can avoid using this strategy, do so. Its dependence upon these flag settings makes it a little risky.

When historic data is read into Forecast Pro, the program establishes a global ending date for the history which corresponds to the most recent observation available *for any item being read in*. If a given item does not have a historic observation for the global ending date, the item is flagged as "dead." Dead items are not forecasted but their demand history goes into the group totals. This can be an issue if your data source codes zeroes as missing (e.g., blank cells in Excel, no transaction in a table or query, etc.). If you check "Set trailing blanks to zero" on the Formats tab of the **Settings>Options** dialog box, Forecast Pro will replace missing trailing observations with zeroes. If you can avoid using this strategy, do so. Its dependence upon these flag settings makes it a little risky.

Zeroes in the middle of your data are *always* treated as numerical values. Therefore, embedded missing values *must be* explicitly coded as missing. Since the way you do this depends upon the data format, it will be addressed as we discuss each specific data format.

Forecast Pro interprets your input data and missing values as the data are read. Thus you cannot change treatment of missing data for a data set that has already been read except by rereading it from scratch. The interpretation of zeros and missing values depends upon the “Ignore leading zeroes” and “Ignore trailing zeroes” flags and upon the “Missing values” setting on the Formats tab of the **Settings>Options** dialog box. There are three possible settings for missing values—*Truncate*, *Impute* or *Zero*.

Truncate directs Forecast Pro to use only the most recent unbroken stream of data. All data up to and including the last missing value are discarded.

Impute directs Forecast Pro to discard leading and trailing missing values but to impute embedded missing values via linear interpolation.

Zero directs Forecast Pro to reset missing values to zeroes and to treat them as numerical values. This setting is appropriate when the data file is written from a database in which only the non-zero records are stored. In other cases it might lead to badly biased forecasts.

Header information

Once you have collected your time series data, you will need to define the *header* information. The required header information includes the following.

Variable name. The name by which the program refers to an individual item (often an SKU). An item name is limited to 256 of the following characters:

0-9 A-Z a-z ! # $ % & ? @ _ + - / * . ~ < >

However, you may also be limited by restrictions imposed by the software to which you export your data. SQL databases, for instance, sometimes require that the first character in a name is a letter and do not regard upper and lower cases as distinct.

If the variable name begins or ends with an underbar (_), Forecast Pro tags it as a potential helper variable and will not forecast it. Helper variables are used in conjunction with event models, forecast by analogy models and weighting transformations.

Variable description. A short description of the variable. Variable descriptions must not exceed 256 characters in length. Variable descriptions are displayed within Forecast Pro and can be included in Numeric Output Files and Formatted Forecast Report Files.

Starting year and period. The year and period corresponding to the first available data point. If the first available observation for a monthly series was April, 2011 then the starting year would be 2011 and the starting period would be 4. The starting year must be 100 or later. For data that are not calendar-oriented, enter 100 for the starting year and 1 for the starting period.

Periods per year. 12 for monthly data, 4 for quarterly data, etc. For data that is not calendar oriented, use 1. Note that most weekly corporate calendars involve occasional years of 53 weeks—these occurrences will cause the forecast date tags to be off by a week.

Periods per cycle. The number of periods per seasonal cycle, usually equal to the number of periods per year. An exception might be daily data, where periods per cycle could be set to 7 to capture weekly patterns and periods per year set to 365. Or, if weekend days are excluded, periods per cycle might be 5 and periods per year 260. For nonseasonal data, periods per cycle should be 1. This is a critical entry because Forecast Pro uses it for seasonal adjustments.

Selecting a Data Format

The program accepts three different data formats—Excel (XLS, XLSX & CSV), Text (MLT, CSV) and Open Database Connectivity (ODBC). Examples of each file type are shipped with the program.

Which data format you choose will depend on your application and database operations. Before making your choice you should be aware of the following advantages and disadvantages of each file type.

Excel (XLS, XLSX & CSV) files

Excel or CSV (Comma Separated Values) files are frequently used for managing Forecast Pro Data input files. You can make and forecast multiple spreadsheet files in one project if you wish.

PROS. If you are comfortable using a spreadsheet, then the spreadsheet format allows you to create, update and manipulate your data in a familiar environment. Processing time for xls files is faster than for ODBC. Creating forecast reports from Excel is straightforward.

CONS. The layout of your spreadsheet must follow the Forecast Pro rules. You may have to change your current spreadsheet layout to one you find slightly less convenient. Processing time for xlsx files can be slow (xls is faster and csv is faster still).

Text (MLT, CSV) files

Text files can also be used to hold your data inputs. You can make and forecast multiple text files in a single project if you wish.

PROS. Processing time is faster than for ODBC and spreadsheets (excluding CSV which is text-based). Most databases can output text files.

CONS. Editing, viewing and updating large MLT files can be clumsy. Not all editors accept very long files.

ODBC

Open database connectivity (ODBC) allows Forecast Pro to read and write data directly to databases for which an ODBC driver exists. Most popular databases support ODBC, including Access, Oracle and SQL server.

PROS. ODBC can provide direct access to data stored in the corporate database obviating the need for intermediate files.

CONS. Reading and writing directly to the corporate database can raise security issues. Data transfer can be slow for some configurations.

The next three sections supply the details you need to set up your data in one of the Forecast Pro formats. If you have chosen a format, you need read only the pertinent section.

Excel (XLS, XLSX, CSV) Formats

Forecast Pro can read Excel (XLS, XLSX) files and CSV (Comma Separated Values) files saved from Excel.

You must create your spreadsheet file from within your spreadsheet program. Forecast Pro checks data types as it reads the spreadsheet. If it sees text where it expects a number, or a floating point number where it expects an integer, an error message will be displayed when the data is read in. By default, a blank cell is interpreted as a missing value. This default can be adjusted for leading and trailing blanks on the Formats tab of the Options dialog box. Do not use zeros to represent missing values—they will be interpreted as numbers and will probably distort your forecasts badly.

Entire spreadsheet vs. named range

You may either devote an entire worksheet to the data (recommended) or specify a portion of a worksheet by giving it the range name BFSDATA.

In a multi-worksheet workbook the data are assumed to reside in the first (topmost) worksheet unless an alternate sheet has been specified by naming it BFSDATA.

The cell references in the examples below assume that the entire worksheet has been devoted to the Forecast Pro database. If you opt to specify a named range, cell A1 in the samples will correspond to the cell in the uppermost left corner of your named range.

Row format

If your data are already stored by rows, you will want to consider the row format first. You can probably alter your spreadsheet to the Forecast Pro row layout in just a few minutes. In row format each time series occupies a single row on the spreadsheet.

The data are assumed to reside in the topmost (first) spreadsheet in the workbook unless an alternate sheet has been specified by naming it BFSDATA.

Row 1 of the spreadsheet is devoted to column headings. The keyword "description" *must* be used as the column heading for the description column. The other cells in row 1 are ignored by Forecast Pro. You can use them for titles, calendar information, etc.

Here is an example of a row format spreadsheet. The file is installed with the program and is named *Sample Historic Data – Horizontal.xlsx*.

Sample Historic Data - Horizontal.xlsx

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P
1	Variable Name	Description	Starting Year	Starting Period	Periods Per Year	Periods Per Cycle	Nov-14	Dec-14	Jan-15	Feb-15	Mar-15	Apr-15	May-15	Jun-15	Jul-15	Aug-15
2	PG1-204-1	Cases - GFCI White	2014	11	12	12	727	677	871	446	641	749	346	612	410	439
3	PG1-204-2	Cases - GFCI Ivory	2014	11	12	12			1145	1159	1908	1332	1620	2081	1526	1627
4	PG1-204-3	Cases - GFCI Brown	2014	11	12	12	770	778	1397	2448	2801	2038	2909	3298	1130	1066
5	PG1-204-4	Cases - Dimmer Ivory	2014	11	12	12				763	886	720	677	1001	871	1850
6	PG1-204-5	Cases - Dimmer White	2014	11	12	12	691	382	497	504	583	583	490	576	468	684

Sheet1 Sheet2

The above example does not use attribute fields to define product groupings. This is often desirable and is discussed in the *Using Attribute Fields to Define a Hierarchy* section.

In our example, each data record consists of six header items in columns A through F, followed by the historic data in the remainder of the row, beginning in column G. Each item is placed in a separate cell.

Columns A through F contain the following six items required items.

A Variable name. Up to 256 characters.

B Variable description. Up to 256 characters.

C Starting year. Must be an integer.

D Starting period. Must be an integer.

E Periods per year. Must be an integer.

F Periods per seasonal cycle. Must be an integer.

See the Overview section at the beginning of this chapter for more details on these items.

The time series data begin in column G. The starting year (column C) and starting period (column D) refer to the year and period for column G, whether or not that cell actually contains data. Thus, in the example spreadsheet, the first value for PG1-204-2 is for January 2014 in column I, even though the

column C defines the starting year as 2013, and column D defines the starting period as 11.

Column format

If your data are already stored by columns, you will want to consider the column format first. You can probably alter your spreadsheet to the Forecast Pro row layout in just a few minutes. In column format each time series occupies a single column on the spreadsheet.

The data are assumed to reside in the topmost (first) spreadsheet in the workbook unless an alternate sheet has been specified by naming it BFSDATA.

Column A of the spreadsheet is devoted to row headings. The keyword "description" *must* be used as the row heading for the description row. The other cells in column A are ignored by Forecast Pro. You can use them for titles, calendar information, etc.

Here is an example of a column format spreadsheet. The file is installed with the program and is named *Sample Historic Data – Vertical.xlsx*.

Sample Historic Data - Vertical.xlsx

	A	B	C	D	E	F
1	Variable Name	PG1-204-1	PG1-204-2	PG1-204-3	PG1-204-4	PG1-204-5
2	Description	Cases - GFCI White	Cases - GFCI Ivory	Cases - GFCI Brown	Cases - Dimmer Ivory	Cases - Dimmer White
3	Starting Year	2014	2014	2014	2014	2014
4	Starting Period	11	11	11	11	11
5	Periods Per Year	12	12	12	12	12
6	Periods Per Cycle	12	12	12	12	12
7	Nov-14	727		770		691
8	Dec-14	677		778		382
9	Jan-15	871	1145	1397		497
10	Feb-15	446	1159	2448	763	504
11	Mar-15	641	1908	2801	886	583
12	Apr-15	749	1332	2038	720	583
13	May-15	346	1620	2909	677	490

Sheet1 Sheet2

The above example does not use attribute fields to define product groupings. This is often desirable and is discussed in the *Using Attribute Fields to Define a Hierarchy* section.

In our example, each data record consists of six header items in rows 1 through 6, followed by the historic data in remainder of the rows, beginning in row 7. Each item is placed in a separate cell.

Rows 1 through 6 contain the following six required items.

1 Variable name. Up to 256 characters.

2 Variable description. Up to 256 characters.

3 Starting year. Must be an integer.

4 Starting period. Must be an integer.

5 Periods per year. Must be an integer.

6 Periods per seasonal cycle. Must be an integer.

See the Overview section at the beginning of this chapter for more details on these items.

The time series data begin in row 7. The starting year (row 3) and starting period (row 4) refer to the year and period for row 7, whether or not that cell actually contains data. Thus in the example spreadsheet, the first value for PG1-204-2 is for January 2014 in row 9, even though the header defines the starting period as 11 and the starting year as 2013.

Transaction format

When demand history is stored in a database, it often consists of individual records for each transaction (e.g., each order or shipment) that includes the quantity and the date. If your data are currently available in Excel in this format, you will want to consider the transaction format. The advantage is that it may require less manipulation than other formats to bring the data into Forecast Pro. The disadvantage is that it is harder for a person to view and manipulate the data in Excel.

The data are assumed to reside in the topmost (first) spreadsheet in the workbook unless an alternate sheet has been specified by naming it BFSDATA.

Row 1 of the spreadsheet is devoted to column headings. The keyword "description" *must* be used as the column heading for the description column. The other cells in row 1 are ignored by Forecast Pro. You can use them for titles, calendar information, etc.

Here is an example of a transaction format spreadsheet. The file is installed with the program and is named S*ample Historic Data – Transactions.xlsx*.

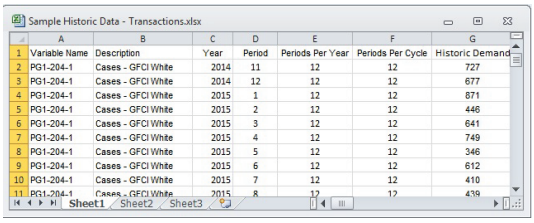

Sample Historic Data - Transactions.xlsx

	A	B	C	D	E	F	G
1	Variable Name	Description	Year	Period	Periods Per Year	Periods Per Cycle	Historic Demand
2	PG1-204-1	Cases - GFCI White	2014	11	12	12	727
3	PG1-204-1	Cases - GFCI White	2014	12	12	12	677
4	PG1-204-1	Cases - GFCI White	2015	1	12	12	871
5	PG1-204-1	Cases - GFCI White	2015	2	12	12	446
6	PG1-204-1	Cases - GFCI White	2015	3	12	12	641
7	PG1-204-1	Cases - GFCI White	2015	4	12	12	749
8	PG1-204-1	Cases - GFCI White	2015	5	12	12	346
9	PG1-204-1	Cases - GFCI White	2015	6	12	12	612
10	PG1-204-1	Cases - GFCI White	2015	7	12	12	410
11	PG1-204-1	Cases - GFCI White	2015	8	12	12	439

Sheet1 Sheet2 Sheet3

The above example does not use attribute fields to define product groupings. This is often desirable and is discussed in the *Using Attribute Fields to Define a Hierarchy* section.

In our example, each data record consists of six header items in columns A through F, followed by the historic data in column G. Each item is placed in a separate cell.

Columns A through F contain the following six items required items.

A Variable name. Up to 256 characters.

B Variable description. Up to 256 characters.

C Year. Must be an integer.

D Period. Must be an integer.

E Periods per year. Must be an integer.

F Periods per seasonal cycle. Must be an integer.

See the Overview section at the beginning of this chapter for more details on these items.

Formatting rules

Keep the following points in mind when you are creating your spreadsheet.

If you are not using attribute fields each variable name must be unique.

If you are using attribute fields each variable name within a given group must be unique.

If you use zeroes to pad the data prior to an item's availability, make sure that Ignore leading zeros is selected on the Formats tab of the Settings Options dialog box.

Text (MLT) Format

A multivariate text file contains information for the data series you wish to forecast.

This file is stored as an editable text file with the extension MLT. You must create this file outside of Forecast Pro, making sure that you format it appropriately.

The MLT file format is straightforward. The file begins with record headings followed by a semicolon. The record headings are required for the Variable Name and Description fields and are optional for all other records. You can use them for titles, calendar information, etc. if you wish.

After the record headings comes header information for the first variable, followed by decimal numbers representing the historical data and then a semicolon. This is followed by the same format for subsequent time series. Two semicolons in a row mark the end of the MLT file.

Since Forecast Pro recognizes spaces, tabs and returns as equivalent terminators, you have a great deal of flexibility in formatting your data. The most common format is to organize series by rows, as in the following example.

Here is an example of an MLT file. The file is installed with the program and is named S*ample Historic Data - Horizontal.mlt*.

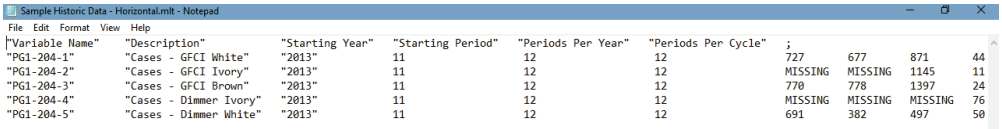


```
"Variable Name"    "Description"            "Starting Year"   "Starting Period"   "Periods Per Year"   "Periods Per Cycle"   ;
"PG1-204-1"        "Cases - GFCI White"     "2013"            11                  12                   12                    727       677       871       44
"PG1-204-2"        "Cases - GFCI Ivory"     "2013"            11                  12                   12                    MISSING   MISSING   1145      11
"PG1-204-3"        "Cases - GFCI Brown"     "2013"            11                  12                   12                    770       778       1397      24
"PG1-204-4"        "Cases - Dimmer Ivory"   "2013"            11                  12                   12                    MISSING   MISSING   MISSING   76
"PG1-204-5"        "Cases - Dimmer White"   "2013"            11                  12                   12                    691       382       497       50
```

The above example does not use attribute fields to define product groupings. This is often desirable and is discussed in the *Using Attribute Fields to Define a Hierarchy* section below.

The header consists of the first six items:

Variable name. Up to 256 characters enclosed in quotes.

Variable description. Up to 256 characters enclosed in quotes.

Starting year. Must be an integer.

Starting period. Must be an integer.

Periods per year. Must be an integer.

Periods per seasonal cycle. Must be an integer.

See the Overview section at the beginning of this chapter for more details on these items.

The remaining fields contain historic data. Each time series is terminated with a semicolon and the file is terminated with two semicolons.

Notice that we have elected to use the same starting period and starting year for each time series and then used the keyword "MISSING" in lieu of values for data points that are unavailable. Thus, in the example, the first value for PG1-204-2 is for January 2012, even though the header defines the starting period as 11 and the starting year as 2011. Aligning the data file this way makes it easier for somebody to view and modify the file. Alternatively, we could have omitted the "MISSING" keywords and altered the starting period and starting year to correspond to each individual time series' starting date. Forecast Pro will have no problem reading a file formatted in this fashion—but you might.

The one-time-series-per-line format illustrated above is easy to read, however, keep in mind you can use any combination of rows and columns.

Formatting rules

Keep the following points in mind when you are creating your .MLT files.

The record headers, variable name field and variable description field must always appear in quotes.

If you are not using attribute fields, each variable name must be unique. If you are using attribute fields, each variable name within a given group must be unique.

If you plan to use commas to separate items, or as place holders for large numbers, you will need to set a few options on the format tab of the Settings Options dialog box to insure that they are interpreted correctly.

Be careful not to use dollar signs, page breaks or other symbols that may confuse Forecast Pro.

Use the word MISSING instead of a number to indicate a missing value.

If you use zeroes to pad the data prior to an item's availability, make sure that Ignore leading zeros is selected on the Formats tab of the Settings Options dialog box.

Using ODBC

ODBC provides direct data communications between Forecast Pro and a wide variety of databases. This is accomplished through intermediary ODBC drivers that lie between Forecast Pro and your database. You must obtain the driver from the database manufacturer or a third party and install it according to the directions provided.

ODBC drivers are available for many database products including Access, Oracle and SQL Server.

Database structure

Forecast Pro reads data from structured tables or views defined in the database. Ordinarily, the tables created for the Forecast Pro interface are just a subset of the entire database.

Here is an example of a of a data table formatted for Forecast Pro. The file is installed with the program and is named *Sample Historic Data – ODBC.mdb*.

Historic_Data : Table

ItemId0	Description	Hist_Year	Hist_Period	PPY	PPC	Hist_Value
COR-12-11	CASE - 12 count Corn Muffins	2014	6	12	12	1944
COR-12-11	CASE - 12 count Corn Muffins	2014	7	12	12	3888
COR-12-11	CASE - 12 count Corn Muffins	2014	8	12	12	5832
COR-12-11	CASE - 12 count Corn Muffins	2014	9	12	12	8748
COR-12-11	CASE - 12 count Corn Muffins	2014	10	12	12	16524
COR-12-11	CASE - 12 count Corn Muffins	2014	11	12	12	16524
COR-12-11	CASE - 12 count Corn Muffins	2014	12	12	12	23328
COR-12-11	CASE - 12 count Corn Muffins	2015	1	12	12	17496
COR-12-11	CASE - 12 count Corn Muffins	2015	2	12	12	2916
COR-12-11	CASE - 12 count Corn Muffins	2015	3	12	12	6804
COR-12-11	CASE - 12 count Corn Muffins	2015	4	12	12	17496
COR-12-11	CASE - 12 count Corn Muffins	2015	5	12	12	1944

Record: 1 of 292

The above example does not use attribute fields to define product groupings. This is often desirable and is discussed in the *Using Attribute Fields to Define a Hierarchy* section.

In our example, each data record consists of six header items in columns 1 through 6, followed by the corresponding historic data point. It is important to note that each data record must contain the aggregated value for the given period.

The fields are defined as follows.

ItemId0. This is a text field containing the variable name. It can be up to 256 characters.

Description. This is a text field containing the variable description. It can be up to 256 characters.

Hist_Year. This is a number field with field size double containing the year.

Hist_Period. This is a number field with field size double containing the period.

Ppy. This is a number field with field size double containing the periods per year.

Ppc. This is a number field with field size double containing the periods per cycle.

Hist_Value. This is a number field with field size double containing the aggregated historic data value for the record's "date" as defined by the Hist_Year and Hist_Period.

See the Overview section at the beginning of this chapter for more details on these items.

The Forecast Pro project needs to connect to an ODBC database before you can access the tables and queries therein. The first time you select Add ODBC in the Data Manager for a given project you will be prompted to establish the file data source (select an ODBC driver) and then to select a database.

Once the project is connected to the database, the tables and queries are displayed, including those that have nothing to do with Forecast Pro. A table from the database is treated just like a file when you are using the Data Manager. Thus you should arrange data into the tables or views that make logical sense as entries in the Data Manager.

Formatting rules

ODBC drivers vary a great deal. If you experience any difficulties connecting to the database, the first thing to check is that you are using the latest ODBC driver available for your database.

With the exception of the last one (which only pertains to Oracle) all of the notes below have been verified using Microsoft Access. There is a chance that your ODBC driver might be more (or less) restrictive.

Missing values are indicated by whatever rules are used by the native database. Every database recognizes the distinction between zeroes and missing values. If you want, you can encode leading and trailing missing values as zeroes. In that case you must make sure that “Ignore leading zeros” and/or “Ignore trailing zeroes” is selected on the Format tab of the **Settings Options** dialog box.

Table names and variable names should begin with a letter and should not include blank spaces.

If you are not using attribute fields each variable name must be unique.

If you are using attribute fields each variable name within a given group must be unique.

Oracle Users: Some Oracle drivers will only work if the Hist_Value field is defined as FLOAT (not NUMBER). All other numeric fields can be NUMBER with the decimal places set to zero.

Using Attribute Fields to Define a Hierarchy

If you do not include any attribute fields, when you read the data into Forecast Pro, it will create a one-level hierarchy (i.e., there will be no group totals).

If your data file, table or query includes attribute fields prior to the Variable Name field, Forecast Pro will automatically build a multiple-level hierarchy. The first level of the hierarchy will correspond to the first attribute field, the second level of the forecasting hierarchy will correspond to the second attribute field, etc.

Let's illustrate this with an example. The spreadsheet *123 Bakery – Historic Data.xlsx* is a sample data file shipped with Forecast Pro. It is shown below.

123 Bakery - Historic Data.xlsx

	A	B	C	D	E	F	G	H	I	J	K	L
1	Total	Category	Customer	SKU	Description	Starting Year	Starting Period	Periods Per Year	Periods Per Cycle	Jan-13	Feb-13	Mar-13
2	Total	Cakes	Food-King	BU-20-02	CASE - 20 oz. Butter Pound Cake	2013	1	12	12			
3	Total	Cakes	Food-King	CA-20-01	CASE - 20 oz. Carrot Layer Cake	2013	1	12	12	20,829	72,576	11,016
4	Total	Cakes	Food-King	CH-20-01	CASE - 20 oz. Chocolate Layer Cake	2013	1	12	12			
5	Total	Cakes	Food-King	CO-20-01	CASE - 20 oz. Coconut Layer Cake	2013	1	12	12	6,763	7,661	9,561
6	Total	Cakes	Food-King	LE-20-02	CASE - 20 oz. Lemon Pound Cake	2013	1	12	12			
7	Total	Cakes	Grocery-Land	BU-20-02	CASE - 20 oz. Butter Pound Cake	2013	1	12	12			
8	Total	Cakes	Grocery-Land	CA-20-01	CASE - 20 oz. Carrot Layer Cake	2013	1	12	12	3,572	4,000	2,808
9	Total	Cakes	Grocery-Land	CH-20-01	CASE - 20 oz. Chocolate Layer Cake	2013	1	12	12	15,307	1,866	3,477
10	Total	Cakes	Grocery-Land	CO-20-01	CASE - 20 oz. Coconut Layer Cake	2013	1	12	12	25,404	450	906
11	Total	Cakes	Sids-Club	BU-20-02	CASE - 20 oz. Butter Pound Cake	2013	1	12	12			
12	Total	Cakes	Sids-Club	CH-20-01	CASE - 20 oz. Chocolate Layer Cake	2013	1	12	12	7,383	6,512	4,375

BFSData Sheet2

In our example, our first attribute field is Total. Notice that the value for the field Total is "Total" for all items. Since all items have the same value for the first attribute field, when we read this spreadsheet into Forecast Pro the first level of our forecasting hierarchy will provide a grand total for all items in our spreadsheet. Notice that the next attribute field is Category, which contains both Cakes and Muffins. Thus, when we read this spreadsheet into Forecast Pro the second level of our forecasting hierarchy will breakdown total sales into Cakes and Muffins. The next attribute field is Customer. Thus, the third level of our forecasting hierarchy will breakdown the Cakes and Muffins into Customers. The next field is the Variable Name (labeled "SKU") so the final level of our hierarchy will breakdown the Category-by-Customer into SKUs.

By default, the *Automatically sort* option on the Formats tab of the Options dialog box is selected, and Forecast Pro sorts the data alphabetically prior to reading it in. If *Automatically sort* is not selected, the order of appearance on

the spreadsheet dictates the structure of the hierarchy and you'll want to take care when constructing your spreadsheet.

ODBC attribute fields

The file *123 Bakery – ODBC.mdb* is a sample data file shipped with Forecast Pro. It is shown below.

Historic_Data : Table

ItemId0	ItemId1	ItemId2	ItemId3	Description	Hist_Year	Hist_Period	PPY	PPC	Hist_Value
Total	Muffins	Stuff-Mart	COR-12-11	CASE - 12 count Corn Muffins	2014	6	12	12	1944
Total	Muffins	Stuff-Mart	COR-12-11	CASE - 12 count Corn Muffins	2014	7	12	12	3888
Total	Muffins	Stuff-Mart	COR-12-11	CASE - 12 count Corn Muffins	2014	8	12	12	5832
Total	Muffins	Stuff-Mart	COR-12-11	CASE - 12 count Corn Muffins	2014	9	12	12	8748
Total	Muffins	Stuff-Mart	COR-12-11	CASE - 12 count Corn Muffins	2014	10	12	12	16524
Total	Muffins	Stuff-Mart	COR-12-11	CASE - 12 count Corn Muffins	2014	11	12	12	16524
Total	Muffins	Stuff-Mart	COR-12-11	CASE - 12 count Corn Muffins	2014	12	12	12	23328
Total	Muffins	Stuff-Mart	COR-12-11	CASE - 12 count Corn Muffins	2015	1	12	12	17496
Total	Muffins	Stuff-Mart	COR-12-11	CASE - 12 count Corn Muffins	2015	2	12	12	2916
Total	Muffins	Stuff-Mart	COR-12-11	CASE - 12 count Corn Muffins	2015	3	12	12	6804
Total	Muffins	Stuff-Mart	COR-12-11	CASE - 12 count Corn Muffins	2015	4	12	12	17496

Record: 1 of 2427

If you are using ODBC the first attribute field *must* be named ItemId0, the second attribute field must be named ItemId1, etc. All attribute fields must be text fields. Other than that, the operation is identical to the spreadsheet format described above.

Chapter 2

Setting Up Your Optional Data

Overview

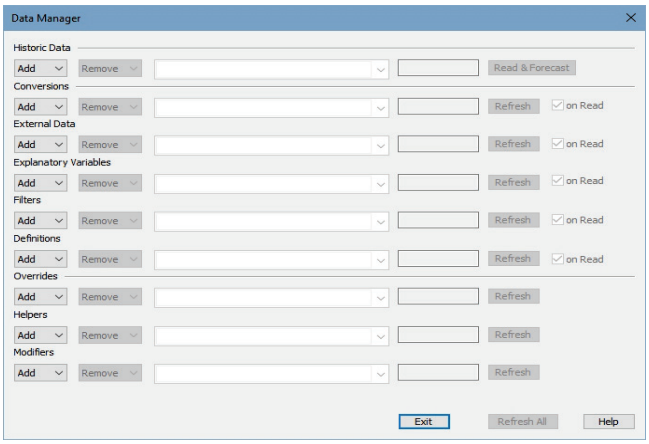

In addition to the historic data (which is required) Forecast Pro's Data Manager allows you to import other information to assist you in creating and working with the forecasts. Specifically, you can import:

Conversions Factors. Conversion factors allow you to display the history and forecasts in alternative units of measure (e.g., dollars, kilos, etc.).

External Data. Forecast Pro allows you to import external data rows (such as alternative forecasts, open orders, etc.) into the forecast override grid.

Explanatory Variables. Explanatory variables files contain variables that you wish to use in conjunction with dynamic regression models.

Filter Fields. Filter fields can be included in report views and filtered like any other fields. They can also be included in the numeric output file.

Definitions. Definition files contain values to use in conjunction with custom forecast allocations and item-level integer rounding quantities.

Overrides. In addition to entering forecast overrides interactively using the Forecast Overrides view, Forecast Pro allows you to import them.

Helper Variables. Helper variables are used in conjunction with event models, forecast by analogy models and weighting transformations.

Modifiers. In addition to specifying forecast modifiers interactively on the Navigator, Forecast Pro allows you to import them.

The remainder of this chapter will detail each of these data types.

Conversion Factors

The unit of measure used in the historic data file is the *default* unit of measure. Forecast Pro allows you to define conversion factors (i.e., multipliers) which can be used to display the history and forecasts in alternative units (e.g., dollars, kilos, etc.).

Examine the file *123 Bakery - Conversions.xlsx*.

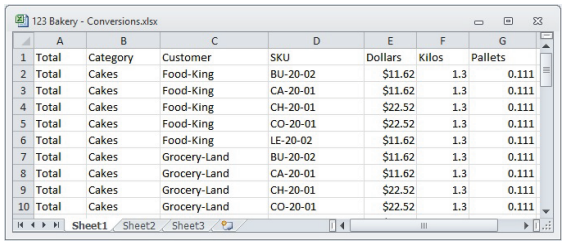

123 Bakery - Conversions.xlsx

	A	B	C	D	E	F	G
1	Total	Category	Customer	SKU	Dollars	Kilos	Pallets
2	Total	Cakes	Food-King	BU-20-02	$11.62	1.3	0.111
3	Total	Cakes	Food-King	CA-20-01	$11.62	1.3	0.111
4	Total	Cakes	Food-King	CH-20-01	$22.52	1.3	0.111
5	Total	Cakes	Food-King	CO-20-01	$22.52	1.3	0.111
6	Total	Cakes	Food-King	LE-20-02	$11.62	1.3	0.111
7	Total	Cakes	Grocery-Land	BU-20-02	$11.62	1.3	0.111
8	Total	Cakes	Grocery-Land	CA-20-01	$11.62	1.3	0.111
9	Total	Cakes	Grocery-Land	CH-20-01	$22.52	1.3	0.111
10	Total	Cakes	Grocery-Land	CO-20-01	$22.52	1.3	0.111

Sheet1 Sheet2 Sheet3

Notice that the attribute fields are present. In our example, columns E through G contain the conversion factors.

By default, the conversion definition must include a record for every end item in the historic data, or you will receive an error message when you attempt to read it. If you want to allow conversion factors with a value of 0 or missing conversion factors, you may do so by activating the *Allow zero valued conversion factors* option on the Formats tab of the Options dialog box.

If the *Allow zero valued conversion factors* option is selected, Forecast Pro will allow blanks or zeros or missing rows in a conversions file specified in the Data Manager. Forecast Pro will give warnings about the missing data, and changing project units using the Units drop down on the dialog bar is not enabled. However, the converted rows are available in the Overrides Report and Graph views. Blanks will be considered to be zeros, and items with either zero or blank conversion factors will consequently not be included in the totals for those conversion factors.

If you will be using a table or query to hold the conversion factors you'll want to examine the *Conversions* table in the file *123 Bakery - ODBC.mdb*.

Conversions : Table

ItemId0	ItemId1	ItemId2	ItemId3	Description	NumValue
Total	Cakes	Stuff-Mart	CA-20-01	Dollars	53.68
Total	Cakes	Stuff-Mart	BU-20-02	Dollars	38.68
Total	Cakes	Sids-Club	CA-20-01	Dollars	53.68
Total	Cakes	Sids-Club	BU-20-02	Dollars	38.68
Total	Cakes	Grocery-Land	CH-20-01	Dollars	55.88
Total	Cakes	Food-King	CO-20-01	Dollars	40.78
Total	Muffins	Stuff-Mart	COR-12-11	Kilos	1.3
Total	Muffins	Stuff-Mart	BRA-12-11	Kilos	1.3
Total	Muffins	Stuff-Mart	BLU-12-11	Kilos	1.3
Total	Muffins	Stuff-Mart	APP-12-11	Kilos	1.3
Total	Muffins	Sids-Club	COR-12-11	Kilos	1.3
Total	Muffins	Sids-Club	BRA-12-11	Kilos	1.3

Record: 1 of 129

Notice that the attribute fields are present. *Description* is a text field containing the conversion factor's name. *NumValue* is a number field with field size double containing the conversion factor.

Formatting rules

In a spreadsheet, a header row that includes the names of the conversion fields is required.

The conversion factors are assigned at the end-item level and must include a record for every end item in the historic data, or you will receive an error message when you attempt to read it.

External Data Files

At times it is useful to have visibility to external information while working with the forecasts. Forecast Pro allows you to import external data rows (such as alternative forecasts, open orders, etc.) into the forecast override grid. The external data can be imported from other Forecast Pro projects, from Excel spreadsheets (both row and transaction-style) and from databases using ODBC.

Importing external data rows requires that you set up the data in a specific format. The sample file shown below is installed into the Forecast Pro input directory and is named *123 Bakery - External Data - SKU-Level Forecast.xlsx.*

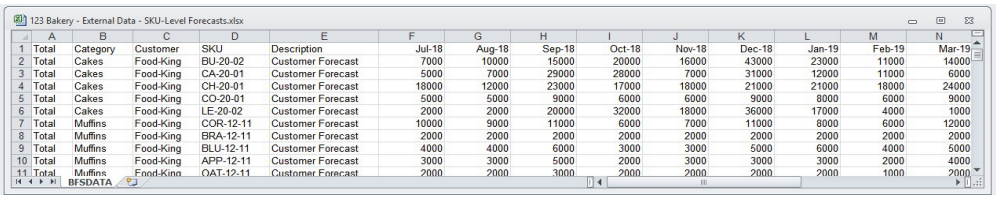

123 Bakery - External Data - SKU-Level Forecasts.xlsx

	A	B	C	D	E	F	G	H	I	J	K	L	M	N
1	Total	Category	Customer	SKU	Description	Jul-18	Aug-18	Sep-18	Oct-18	Nov-18	Dec-18	Jan-19	Feb-19	Mar-19
2	Total	Cakes	Food-King	BU-20-02	Customer Forecast	7000	10000	15000	20000	16000	43000	23000	11000	14000
3	Total	Cakes	Food-King	CA-20-01	Customer Forecast	5000	7000	29000	28000	7000	31000	12000	11000	6000
4	Total	Cakes	Food-King	CH-20-01	Customer Forecast	18000	12000	23000	17000	18000	21000	21000	18000	24000
5	Total	Cakes	Food-King	CO-20-01	Customer Forecast	5000	5000	9000	6000	6000	9000	8000	6000	9000
6	Total	Cakes	Food-King	LE-20-02	Customer Forecast	2000	2000	20000	32000	18000	36000	17000	4000	1000
7	Total	Muffins	Food-King	COR-12-11	Customer Forecast	10000	9000	11000	6000	7000	11000	8000	6000	12000
8	Total	Muffins	Food-King	BRA-12-11	Customer Forecast	2000	2000	2000	2000	2000	2000	2000	2000	2000
9	Total	Muffins	Food-King	BLU-12-11	Customer Forecast	4000	4000	6000	3000	3000	5000	6000	4000	5000
10	Total	Muffins	Food-King	APP-12-11	Customer Forecast	3000	3000	5000	2000	3000	3000	3000	2000	4000
11	Total	Muffins	Food-King	OAT-12-11	Customer Forecast	2000	2000	3000	2000	2000	2000	2000	1000	2000

BFSDATA

The first columns (A through D in our example) define the attributes and name of the item (i.e., the hierarchy information). The next column (E in our example) is labeled Description and contains the name of the external row. The remaining columns are labeled with dates and contain the data values.

If you plan to import your external data using a transaction style spreadsheet, examine the sample file *123 Bakery – External Data – SKU-Level Forecasts – Transactions.xlsx* located in the Forecast Pro input directory. Note that the file layout is the same as the ODBC file layout shown below.

If you plan to import your external data using ODBC, examine the *SKU-Level Forecasts* table in the sample Access database file *123-Bakery-ODBC.mdb* located in the Forecast Pro input directory.

SKU-level Forecasts : Table

ItemId0	ItemId1	ItemId2	ItemId3	Description	Hist_Year	Hist_Period	Hist_Value
Total	Muffins	Stuff-Mart	COR-12-11	Management Forecast	2018	7	6000
Total	Muffins	Stuff-Mart	COR-12-11	Management Forecast	2018	8	9000
Total	Muffins	Stuff-Mart	COR-12-11	Management Forecast	2018	9	15000
Total	Muffins	Stuff-Mart	COR-12-11	Management Forecast	2018	10	23000
Total	Muffins	Stuff-Mart	COR-12-11	Management Forecast	2018	11	17000
Total	Muffins	Stuff-Mart	COR-12-11	Management Forecast	2018	12	44000
Total	Muffins	Stuff-Mart	COR-12-11	Management Forecast	2019	1	25000
Total	Muffins	Stuff-Mart	COR-12-11	Management Forecast	2019	2	13000
Total	Muffins	Stuff-Mart	COR-12-11	Management Forecast	2019	3	12000
Total	Muffins	Stuff-Mart	COR-12-11	Management Forecast	2019	4	10000
Total	Muffins	Stuff-Mart	COR-12-11	Management Forecast	2019	5	9000
Total	Muffins	Stuff-Mart	COR-12-11	Management Forecast	2019	6	9000
Total	Muffins	Stuff-Mart	COR-12-11	Sales Forecast	2018	7	7000
Total	Muffins	Stuff-Mart	COR-12-11	Sales Forecast	2018	8	9000
Total	Muffins	Stuff-Mart	COR-12-11	Sales Forecast	2018	9	16000

Record: 1 of 1188

The format is very similar to the format for the historic demand. Notice that the attribute fields are present. *Hist_Year*, *Hist_Period* and *Hist_Value* are number fields identifying the year, period and values respectively.

Formatting rules

In a spreadsheet, a header row is required. The Description and date fields must also be present and labeled accordingly. The dates in the header row need to be either in an Excel format, or in a Year-Period format, e.g. 2015-July, or 2015-52.

The external data rows are imported at one level of the hierarchy only.

If an item has no entries for a given external data row, it does not need to be included in the spreadsheet, table or query.

Explanatory Variables

Explanatory variable files are used to import explanatory variables (i.e., independent variables) that you wish to include in a dynamic regression model.

Forecast Pro allows you to specify two types of explanatory variables—global and item-specific. A global explanatory variable consists of a single time series (set of values) which can be included in a dynamic regression model for any item on the Navigator. An item-specific explanatory variable consists of a set of time series each of which is associated to a specific item on the Navigator.

To illustrate the difference, consider a variable like holidays which will likely be the same for all items within a hierarchy and therefore should be defined as a global variable. Contrast holidays with a variable like price, which will likely take on different values for different items and therefore should be defined as an item-specific variable. In this context "items" is being used to include both end items and aggregate levels in the hierarchy.

Explanatory variables *must* include values for the historic period. Ideally, you would also include forecast values for your explanatory variable, but you are not required to do so. If an explanatory variable does not have values provided for the complete forecast period, you are able to use the "automatically extend" option, and Forecast Pro will use expert selection to forecast explanatory variables where necessary.

Global Explanatory Variables

Setting up global variables is straightforward. The first column is labeled Variable and contains the name of the explanatory variable. The remaining columns are labeled with dates and contain the date values. These date values should match what is used in Forecast Pro (e.g. the dates shown in the Numeric Output view).

The sample file shown below, "*Electricity - Explanatory Variables.xlsx*", is provided with the software.

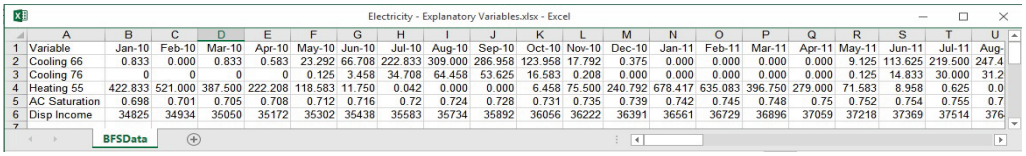

Electricity - Explanatory Variables.xlsx - Excel

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U
1	Variable	Jan-10	Feb-10	Mar-10	Apr-10	May-10	Jun-10	Jul-10	Aug-10	Sep-10	Oct-10	Nov-10	Dec-10	Jan-11	Feb-11	Mar-11	Apr-11	May-11	Jun-11	Jul-11	Aug-
2	Cooling 66	0.833	0.000	0.833	0.583	23.292	66.708	222.833	309.000	286.958	123.958	17.792	0.375	0.000	0.000	0.000	0.000	9.125	113.625	219.500	247.4
3	Cooling 76	0	0	0	0	0.125	3.458	34.708	64.458	53.625	16.583	0.208	0.000	0.000	0.000	0.000	0.000	0.125	14.833	30.000	31.2
4	Heating 55	422.833	521.000	387.500	222.208	118.583	11.750	0.042	0.000	0.000	6.458	75.500	240.792	678.417	635.083	396.750	279.000	71.583	8.958	0.625	0.0
5	AC Saturation	0.698	0.701	0.705	0.708	0.712	0.716	0.72	0.724	0.728	0.731	0.735	0.739	0.742	0.745	0.748	0.75	0.752	0.754	0.755	0.7
6	Disp Income	34825	34934	35050	35172	35302	35438	35583	35734	35892	36056	36222	36391	36561	36729	36896	37059	37218	37369	37514	376
7																					

BFSData

Item-specific Explanatory Variables

The item-specific explanatory variable setup is the same as the Global variable setup, except the file includes the attribute fields used in the historic data setup. The first few columns (columns A-D in the example below) list the attributes, the next column (column E below) is labeled Variable and contains the name of the explanatory variable and the remaining columns are labeled with dates and contain the date values (consistent with Forecast Pro, as described above).

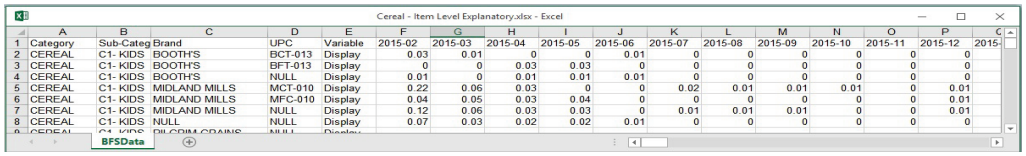

Cereal - Item Level Explanatory.xlsx - Excel

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P
1	Category	Sub-Categ	Brand	UPC	Variable	2015-02	2015-03	2015-04	2015-05	2015-06	2015-07	2015-08	2015-09	2015-10	2015-11	2015-12
2	CEREAL	C1- KIDS	BOOTH'S	BCT-013	Display	0.03	0.01	0	0	0.01	0	0	0	0	0	0
3	CEREAL	C1- KIDS	BOOTH'S	BFT-013	Display	0	0	0.03	0.03	0	0	0	0	0	0	0
4	CEREAL	C1- KIDS	BOOTH'S	NULL	Display	0.01	0	0.01	0.01	0.01	0	0	0	0	0	0
5	CEREAL	C1- KIDS	MIDLAND MILLS	MCT-010	Display	0.22	0.06	0.03	0	0	0.02	0.01	0.01	0.01	0	0.01
6	CEREAL	C1- KIDS	MIDLAND MILLS	MFC-010	Display	0.04	0.05	0.03	0.04	0	0	0	0	0	0	0.01
7	CEREAL	C1- KIDS	MIDLAND MILLS	NULL	Display	0.12	0.06	0.03	0.03	0	0.01	0.01	0.01	0	0	0.01
8	CEREAL	C1- KIDS	NULL	NULL	Display	0.07	0.03	0.02	0.02	0.01	0	0	0	0	0	0

BFSData

The examples shown above use a row format spreadsheet. Explanatory variables may also be imported using all other formats supported for input files (e.g., transaction spreadsheet, ODBC, etc.). ODBC requires the same layout as a transaction spreadsheet. An example of item-specific transaction explanatory variables, "*Cereal- Item Level Explanatory Transactions.xlsx*," is provided with the software and shown below.

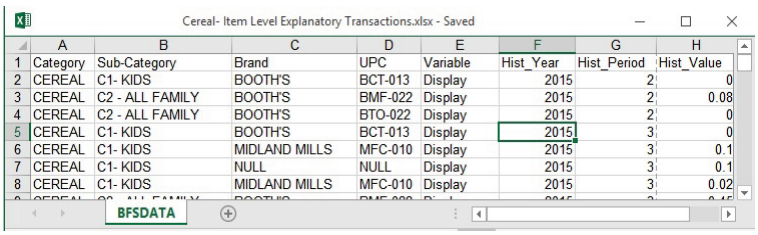

Cereal- Item Level Explanatory Transactions.xlsx - Saved

	A	B	C	D	E	F	G	H
1	Category	Sub-Category	Brand	UPC	Variable	Hist_Year	Hist_Period	Hist_Value
2	CEREAL	C1- KIDS	BOOTH'S	BCT-013	Display	2015	2	0
3	CEREAL	C2 - ALL FAMILY	BOOTH'S	BMF-022	Display	2015	2	0.08
4	CEREAL	C2 - ALL FAMILY	BOOTH'S	BTO-022	Display	2015	2	0
5	CEREAL	C1- KIDS	BOOTH'S	BCT-013	Display	2015	3	0
6	CEREAL	C1- KIDS	MIDLAND MILLS	MFC-010	Display	2015	3	0.1
7	CEREAL	C1- KIDS	NULL	NULL	Display	2015	3	0.1
8	CEREAL	C1- KIDS	MIDLAND MILLS	MFC-010	Display	2015	3	0.02

BFSDATA

Notice that the attribute fields are in the first few columns. The attribute columns are followed by a column for the explanatory variable name. This column must be labeled *Variable*. The last 3 columns, *Hist_Year*, *Hist_Period* and *Hist_Value*, are number fields identifying the year, period and values respectively.

Custom Filter Fields

You can define custom filter fields which can be included in the report views and filtered like any other fields. They can also be included in the numeric output file.

Examine the file *123 Bakery - Filters.xlsx*.

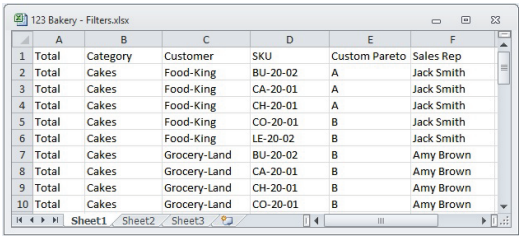

123 Bakery - Filters.xlsx

	A	B	C	D	E	F
1	Total	Category	Customer	SKU	Custom Pareto	Sales Rep
2	Total	Cakes	Food-King	BU-20-02	A	Jack Smith
3	Total	Cakes	Food-King	CA-20-01	A	Jack Smith
4	Total	Cakes	Food-King	CH-20-01	A	Jack Smith
5	Total	Cakes	Food-King	CO-20-01	B	Jack Smith
6	Total	Cakes	Food-King	LE-20-02	B	Jack Smith
7	Total	Cakes	Grocery-Land	BU-20-02	B	Amy Brown
8	Total	Cakes	Grocery-Land	CA-20-01	B	Amy Brown
9	Total	Cakes	Grocery-Land	CH-20-01	B	Amy Brown
10	Total	Cakes	Grocery-Land	CO-20-01	B	Amy Brown

Sheet1 Sheet2 Sheet3

Notice that the attribute fields are present. In our example, columns E and F contain the filter fields.

If you will be using a table or query to hold the filter fields you'll want to examine the *Filters* table in the file *123 Bakery - ODBC.mdb*.

Filters : Table

ItemId0	ItemId1	ItemId2	ItemId3	Description	TextValue
Total	Cakes	Food-King	CO-20-01	Custom Pareto	B
Total	Cakes	Food-King	CH-20-01	Custom Pareto	B
Total	Cakes	Food-King	CA-20-01	Custom Pareto	B
Total	Cakes	Food-King	BU-20-02	Custom Pareto	C
Total	Muffins	Grocery-Land	COR-12-11	Sales Rep	Bob White
Total	Muffins	Grocery-Land	BLU-12-11	Sales Rep	Bob White
Total	Muffins	Grocery-Land	APP-12-11	Sales Rep	Bob White
Total	Muffins	Food-King	COR-12-11	Sales Rep	Bob White
Total	Muffins	Food-King	BLU-12-11	Sales Rep	Bob White
Total	Muffins	Food-King	APP-12-11	Sales Rep	Bob White

Record: 37 of 86

Notice that the attribute fields are present. *Description* is a text field containing the filter name. *TextValue* is a text field containing the filter's value.

Formatting rules

In a spreadsheet, a header row that includes the names of the filter fields is required.

The filters can be assigned at any level.

If an item has no associated filters, it does not need to be included in the spreadsheet, table or query.

Definitions: Custom Allocations

A custom allocation approach allows you to dictate how a group-level forecast is allocated to its component series. This is appropriate when the breakdown between the two levels is known.

To use a custom allocation you must define the allocation proportions in a file, table or query and read them into Forecast Pro using the Definitions row in the Data Manager.

Examine the file *Running Shoes – Custom Allocation.xlsx*.

Running Shoes - Custom Allocation.xlsx

	A	B	C	D	E	F	G
1	Category	Gender	Style	Color	Style ID	Size	CALL
2	Running	Mens	Mercury Cloud	Black	MC-6-M-E1	7.0	0.8
3	Running	Mens	Mercury Cloud	Black	MC-6-M-E1	7.5	1.1
4	Running	Mens	Mercury Cloud	Black	MC-6-M-E1	8.0	2.6
5	Running	Mens	Mercury Cloud	Black	MC-6-M-E1	8.5	4.1
6	Running	Mens	Mercury Cloud	Black	MC-6-M-E1	9.0	6.6
7	Running	Mens	Mercury Cloud	Black	MC-6-M-E1	9.5	8.3
8	Running	Mens	Mercury Cloud	Black	MC-6-M-E1	10.0	10.6
9	Running	Mens	Mercury Cloud	Black	MC-6-M-E1	10.5	12
10	Running	Mens	Mercury Cloud	Black	MC-6-M-E1	11.0	13.1
11	Running	Mens	Mercury Cloud	Black	MC-6-M-E1	11.5	9.9

Sheet1

Notice that the attribute fields are present. In our example, column G contains the allocation proportions.

If you will be using a table or query to hold the allocation proportions you'll want to examine the *CustomAllocation* table in the file *Running Shoes.mdb*.

CustomAllocation : Table

ItemId0	ItemId1	ItemId2	ItemId3	ItemId4	ItemId5	Description	NumValue
Running	Womens	Mercury Cloud	Black	MC-1-W-B1	5.0	CALL	0.4
Running	Womens	Mercury Cloud	Black	MC-1-W-D1	5.0	CALL	0.4
Running	Womens	Mercury Cloud	Black	MC-2-W-A1	5.0	CALL	0.4
Running	Womens	Mercury Cloud	Bone	MC-1-W-B1	5.0	CALL	0.4
Running	Womens	Mercury Cloud	Bone	MC-1-W-D1	5.0	CALL	0.4
Running	Womens	Mercury Cloud	Bone	MC-2-W-A1	5.0	CALL	0.4
Running	Womens	Mercury Cloud	White	MC-1-W-B1	5.0	CALL	0.4
Running	Womens	Mercury Cloud	White	MC-1-W-D1	5.0	CALL	0.4
Running	Womens	Mercury Cloud	White	MC-2-W-A1	5.0	CALL	0.4
Running	Womens	Mercury Cloud I	Black	MCP-1-W-B1	5.0	CALL	0.4
Running	Womens	Mercury Cloud I	Black	MCP-1-W-D1	5.0	CALL	0.4

Record: 1 of 530

Notice that the attribute fields are present. *Description* is a text field containing the keyword CALL to identify the entry as an allocation proportion. *NumValue* is a number field with field size double containing the allocation proportion.

Formatting rules

In a spreadsheet, a header row that uses the keyword CALL to identify the allocation proportions column is required.

The allocation proportions can be assigned at any level.

The allocations do not need to sum to 100 (i.e., they do not need to be percentages), they are simply relative proportions. For example, if a group had two component series, A and B, and you assigned A a weight of 2 and B a weight of 1, then series A would get a two thirds allocation and series B would get one third.

Definitions: Integer Rounding Quantities

Forecast Pro allows you to display *integer forecast rows* for the statistical and/or final forecasts on the override grid. These rows allow you to define an integer quantity to round to (e.g., 1, 50, 1000, etc.) and then round the forecasts using an algorithm that that keeps track of the "remainder" (rounded amount) and applies it to the next period's forecast value. This is particularly useful for low-volume forecasts or when there are large minimum order quantities, where traditional rounding can result in biased forecasts.

When using integer forecasts, the user defines the default rounding quantity in the Advanced Controls tab of the **Settings>Options** dialog box. If you wish to assign item-specific rounding quantities you can do so by defining them in a file, table or query and reading them into Forecast Pro using the Definitions row in the Data Manager.

Examine the file *123 Bakery – Definitions – Integer Rounding Quantities.xlsx*.

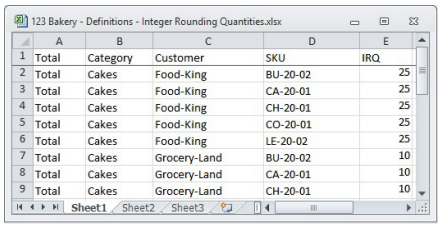

123 Bakery - Definitions - Integer Rounding Quantities.xlsx

	A	B	C	D	E
1	Total	Category	Customer	SKU	IRQ
2	Total	Cakes	Food-King	BU-20-02	25
3	Total	Cakes	Food-King	CA-20-01	25
4	Total	Cakes	Food-King	CH-20-01	25
5	Total	Cakes	Food-King	CO-20-01	25
6	Total	Cakes	Food-King	LE-20-02	25
7	Total	Cakes	Grocery-Land	BU-20-02	10
8	Total	Cakes	Grocery-Land	CA-20-01	10
9	Total	Cakes	Grocery-Land	CH-20-01	10

Sheet1 Sheet2 Sheet3

Notice that the attribute fields are present. In our example, column E contains the rounding quantities.

If you will be using a table or query to hold the rounding quantities you'll want to examine the *IRQ* table in the file *123 Bakery - ODBC.mdb*

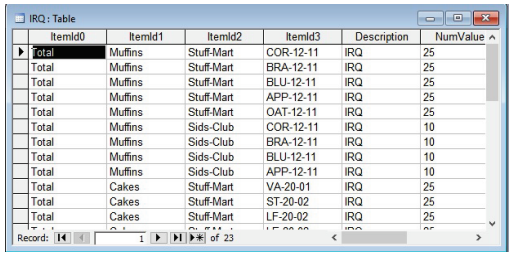

IRQ : Table

ItemId0	ItemId1	ItemId2	ItemId3	Description	NumValue
Total	Muffins	Stuff-Mart	COR-12-11	IRQ	25
Total	Muffins	Stuff-Mart	BRA-12-11	IRQ	25
Total	Muffins	Stuff-Mart	BLU-12-11	IRQ	25
Total	Muffins	Stuff-Mart	APP-12-11	IRQ	25
Total	Muffins	Stuff-Mart	OAT-12-11	IRQ	25
Total	Muffins	Sids-Club	COR-12-11	IRQ	10
Total	Muffins	Sids-Club	BRA-12-11	IRQ	10
Total	Muffins	Sids-Club	BLU-12-11	IRQ	10
Total	Muffins	Sids-Club	APP-12-11	IRQ	10
Total	Cakes	Stuff-Mart	VA-20-01	IRQ	25
Total	Cakes	Stuff-Mart	ST-20-02	IRQ	25
Total	Cakes	Stuff-Mart	LF-20-02	IRQ	25

Record: 1 of 23

Notice that the attribute fields are present. *Description* is a text field containing the keyword IRQ (Integer Rounding Quantity) to identify the entry as a rounding quantity. *NumValue* is a number field with field size double containing the rounding quantity.

Formatting rules

In a spreadsheet, a header row that uses the keyword IRQ to identify the rounding quantity column is required.

The rounding quantities can only be assigned at the end-item level.

If you wish to use the default rounding quantity defined in Forecast Pro for a given item, it does not need to be included in the spreadsheet, table or query.

Override Files

In addition to entering forecast overrides interactively using the Forecast Overrides view, Forecast Pro allows you to import overrides from other Forecast Pro projects, from Excel spreadsheets (both row and transaction-style) and from databases using ODBC.

The sample file shown below is installed into the Forecast Pro input directory and is named *123 Bakery – Overrides - Rows.xlsx*

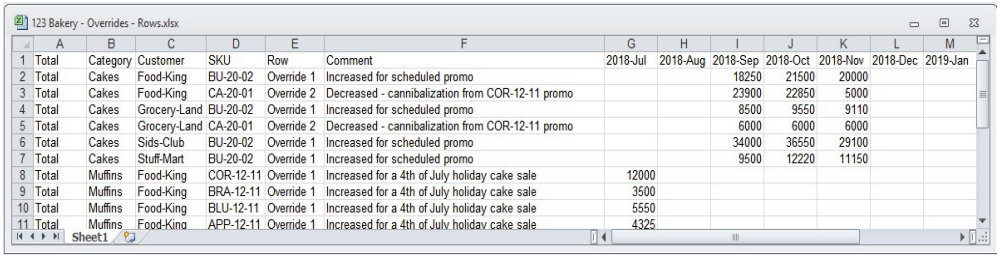

	A	B	C	D	E	F	G	H	I	J	K	L	M
1	Total	Category	Customer	SKU	Row	Comment	2018-Jul	2018-Aug	2018-Sep	2018-Oct	2018-Nov	2018-Dec	2019-Jan
2	Total	Cakes	Food-King	BU-20-02	Override 1	Increased for scheduled promo			18250	21500	20000		
3	Total	Cakes	Food-King	CA-20-01	Override 2	Decreased - cannibalization from COR-12-11 promo			23900	22850	5000		
4	Total	Cakes	Grocery-Land	BU-20-02	Override 1	Increased for scheduled promo			8500	9550	9110		
5	Total	Cakes	Grocery-Land	CA-20-01	Override 2	Decreased - cannibalization from COR-12-11 promo			6000	6000	6000		
6	Total	Cakes	Sids-Club	BU-20-02	Override 1	Increased for scheduled promo			34000	36550	29100		
7	Total	Cakes	Stuff-Mart	BU-20-02	Override 1	Increased for scheduled promo			9500	12220	11150		
8	Total	Muffins	Food-King	COR-12-11	Override 1	Increased for a 4th of July holiday cake sale	12000						
9	Total	Muffins	Food-King	BRA-12-11	Override 1	Increased for a 4th of July holiday cake sale	3500						
10	Total	Muffins	Food-King	BLU-12-11	Override 1	Increased for a 4th of July holiday cake sale	5550						
11	Total	Muffins	Food-King	APP-12-11	Override 1	Increased for a 4th of July holiday cake sale	4325						

The first columns (A through D in our example) define the attributes and name of the item (i.e., the hierarchy information). The next column defines the override row to which the changes should be applied, followed by a column for comments (which, naturally, populates the comment field). Finally, the overrides themselves appear from left to right, starting with the first forecast period. If a specific cell is blank, no override will be imported for that period.

Please note that you may use formula overrides (formulas must begin with an equal sign) as well as number overrides.

If you plan to import your overrides using a transaction style spreadsheet, examine the sample file *123 Bakery – Overrides - Transactions.xlsx* located in the Forecast Pro input directory.

If you plan to import your overrides using ODBC, examine the *Overrides* table in the sample Access database file *123-Bakery-ODBC.mdb* located in the Forecast Pro input directory.

Overrides : Table

ItemId0	ItemId1	ItemId2	ItemId3	FC_Period	FC_Year	Override_Row	Override_Formula	Override_Created	Override_Comment
Total	Cakes	Grocery-Land	BU-20-02	8	2018	2	=STAT*.8	5/19/2017 4:17:48 PM	G-Land to discontinue cakes at 40% of store locations
Total	Cakes	Grocery-Land	BU-20-02	9	2018	2	=STAT*.7	5/19/2017 4:17:48 PM	G-Land to discontinue cakes at 40% of store locations
Total	Cakes	Grocery-Land	BU-20-02	10	2018	2	=STAT*.6	5/19/2017 4:17:48 PM	G-Land to discontinue cakes at 40% of store locations
Total	Cakes	Grocery-Land	BU-20-02	11	2018	2	=STAT*.6	5/19/2017 4:17:48 PM	G-Land to discontinue cakes at 40% of store locations
Total	Cakes	Grocery-Land	BU-20-02	12	2018	2	=STAT*.6	5/19/2017 4:17:48 PM	G-Land to discontinue cakes at 40% of store locations
Total	Cakes	Grocery-Land	BU-20-02	1	2019	2	=STAT*.6	5/19/2017 4:17:48 PM	G-Land to discontinue cakes at 40% of store locations
Total	Cakes	Grocery-Land	BU-20-02	2	2019	2	=STAT*.6	5/19/2017 4:17:48 PM	G-Land to discontinue cakes at 40% of store locations
Total	Cakes	Grocery-Land	BU-20-02	3	2019	2	=STAT*.6	5/19/2017 4:17:48 PM	G-Land to discontinue cakes at 40% of store locations
Total	Cakes	Grocery-Land	BU-20-02	4	2019	2	=STAT*.6	5/19/2017 4:17:48 PM	G-Land to discontinue cakes at 40% of store locations

Record: 9 of 114

Notice that the attribute fields are present.

FC_Period and *FC_Year* are number fields identifying the forecast date for the override.

Override_Row can either be a number field identifying the row number or a text field identifying the row name.

Override_Formula is a text field containing either a numeric value for the override or a formula (formulas must begin with an = sign).

Override_Created is a date/time field identifying when the override was created. If the entry is blank the override report will list the date/time that the override was loaded into the project.

Override_Comment is a text field containing the comment you wish to associate with the override. If the entry is blank, no comment will be associated with the override.

Formatting rules

In a spreadsheet, a header row is required. The Row and Comment fields must be present and labeled accordingly.

In a spreadsheet, the dates in the header row need to be either in an Excel format, or in a Year-Period format, e.g. 2017-July, or 2016-52.

Overrides can be assigned at any level.

The override row can be specified by either the row number or the row name.

Overrides can be entered as values or formulas.

If an item has no overrides, it does not need to be included in the spreadsheet, table or query.

Helper Files

Helper variables are used in conjunction with event modeling, forecast by analogy and the weighting transformation.

Usually event schedules are created interactively using the Event Manager. However, if you wish, you can create helper variables containing the event schedules in Excel or in a database and import them into your project using the Data Manager.

Forecast by analogy requires a helper variable containing the analogy series. Using a weighting transformation requires a helper variable containing the weights. Depending on your needs, you can either create these helper variables using the Create Helper Variable dialog box (accessed via **Helper variable>Create** on the Navigator's context menu) or you can create them in Excel or in a database and import them into your project using the Data Manager.

Helper variables are created in an identical fashion to your input data. Please note the following:

1. Attribute fields are not used since helper variables are associated with items on the Navigator using modifiers.

2. Description, Starting Year, Starting Period, Periods per Year and Periods per Seasonal Cycle are all required fields.

3. Their span should cover both the historic and forecast period.

Several sample helper files are shipped with Forecast Pro including, *"Beer - Helpers.xls."* and the table *"Helpers"* in the sample Access database file *Beer.mdb*.

Modifier Files

In addition to entering forecast modifiers interactively on the Navigator, Forecast Pro allows you to import modifiers from other Forecast Pro projects and from Excel spreadsheets.

Importing modifiers from Excel requires that you set up the spreadsheet in a specific format. The sample file shown below is installed into the Forecast Pro input directory and is named *123 Bakery - Modifiers.xlsx.*

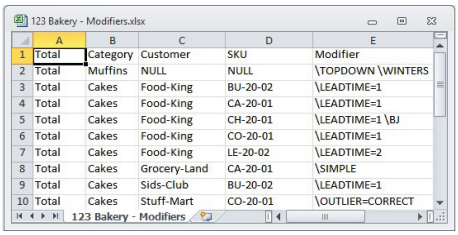

123 Bakery - Modifiers.xlsx

	A	B	C	D	E
1	Total	Category	Customer	SKU	Modifier
2	Total	Muffins	NULL	NULL	\TOPDOWN \WINTERS
3	Total	Cakes	Food-King	BU-20-02	\LEADTIME=1
4	Total	Cakes	Food-King	CA-20-01	\LEADTIME=1
5	Total	Cakes	Food-King	CH-20-01	\LEADTIME=1 \BJ
6	Total	Cakes	Food-King	CO-20-01	\LEADTIME=1
7	Total	Cakes	Food-King	LE-20-02	\LEADTIME=2
8	Total	Cakes	Grocery-Land	CA-20-01	\SIMPLE
9	Total	Cakes	Sids-Club	BU-20-02	\LEADTIME=1
10	Total	Cakes	Stuff-Mart	CO-20-01	\OUTLIER=CORRECT

123 Bakery - Modifiers

The first columns (A through D in our example) define the attributes and name of the item (i.e., the hierarchy information). The next column defines the modifier(s) to import.

If you will be using a table or query to hold the modifier fields you'll want to examine the *Modifiers* table in the file *123 Bakery - ODBC.mdb*.

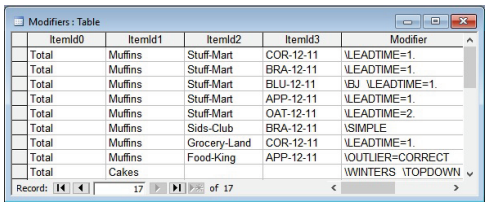

Modifiers : Table

ItemId0	ItemId1	ItemId2	ItemId3	Modifier
Total	Muffins	Stuff-Mart	COR-12-11	\LEADTIME=1.
Total	Muffins	Stuff-Mart	BRA-12-11	\LEADTIME=1.
Total	Muffins	Stuff-Mart	BLU-12-11	\BJ \LEADTIME=1.
Total	Muffins	Stuff-Mart	APP-12-11	\LEADTIME=1.
Total	Muffins	Stuff-Mart	OAT-12-11	\LEADTIME=2.
Total	Muffins	Sids-Club	BRA-12-11	\SIMPLE
Total	Muffins	Grocery-Land	COR-12-11	\LEADTIME=1.
Total	Muffins	Food-King	APP-12-11	\OUTLIER=CORRECT
Total	Cakes			\WINTERS \TOPDOWN

Record: 17 of 17

Notice that the attribute fields are present. *Modifier* is a text field containing the modifier(s) you wish to apply.

Formatting rules

In a spreadsheet, a header row is required. The Modifier field must be present and labeled accordingly.

The modifiers can be assigned at any level.

Multiple modifiers can be assigned in a single cell or data field.

If an item has no modifiers, it does not need to be included in the spreadsheet, table or query.

Chapter 3

Saving Your Work

Forecast Pro TRAC can save eight different types of output.

Forecast Projects allow you to save your forecasting session so that you can return later and pick up where you left off or to share the session with others. The forecast project saves the data, forecasts, overrides, Hot List, modifiers etc. as they currently exist in memory. If you update your input data files after saving a project, when you reopen the project you'll have the option of either generating new forecasts using the updated data or restoring the conditions present when the project was saved (i.e., not recalculating the forecasts using the new data).

Numeric Output Files can contain time series output such as history, forecasts, fitted values and confidence limits as well as statistical output such as the forecasting model specification and summary statistics. These files can be saved in text, spreadsheet, ODBC or XML formats. You can include output for either all items forecasted or just for the Hot List. Generally speaking, if you will be importing the forecasts into another application these are the files you will want to use.

Formatted Forecast Reports can be saved to Excel. You can save these reports for the currently displayed item, all items on the Hot List or all items forecasted. The Excel workbook will include a separate worksheet for each item containing a graph (optional) and the on-screen Forecast Report. These

reports are convenient if you wish to present the forecasts and/or models used to colleagues.

Tracking Reports can be saved to Excel. You can save these reports for the currently displayed item, all items on the Hot List or all items forecasted. The Excel workbook will include a separate worksheet for each item containing a graph (optional) and the on-screen Tracking Report. These reports are only available if the database includes at least one archived forecast for the given item.

The *Item Report* can be viewed on-screen and saved to Excel. This report can contain Pareto information, statistical output such as the forecasting model specification and summary statistics as well as time series output such as history, forecasts, fitted values and confidence limits.

The *Override Report* can be viewed on-screen and saved to Excel. This report lists items where overrides and/or comments have been made.

The *Exception Report* can be viewed on-screen and saved to Excel. If you are using filters to identify exceptions, the report will only include the filtered items.

The *Outlier Report* can be viewed on-screen and saved to Excel. This report lists items where outliers have been detected and/or corrected.

The following sections discuss these output files in more detail.

Forecast Projects

Forecast Projects allow you to save your forecasting session so that you can return later and pick up where you left off or to share the session with others. The forecast project saves the data, forecasts, overrides, Hot List, modifiers etc. as they currently exist in memory. If you update your input data files after saving a project, when you reopen the project you'll have the option of either generating new forecasts using the updated data or restoring the conditions present when the project was saved (i.e., not recalculating the forecasts using the new data).

Opening projects

To save a forecast project you select **File>Save** or click the blue Save Project icon ( ). To open a forecast project you select **File>Open** or click the blue Open Project icon ( ).

If you open a project and data files have not changed (i.e., the time and date stamps match), the project will restore the conditions present when the project was saved allowing you to continue your work.

If you open a project and the data files are not available, the project will restore the conditions present when the project was saved—thus you can review projects, make adjustments, etc. even if you do not have access to the data files. This will often be the case when sharing projects with colleagues.

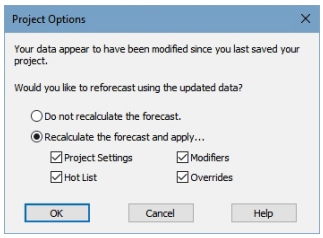

If you open a project and data files have changed, a dialog box will appear allowing you to either generate new forecasts using the updated data or restore the conditions present when the project was saved (i.e., do not recalculate the forecasts using the new data).

Project files

The project normally consists of a single zipped file with the extension .FPZip. The zipped file contains three separate data files.

The Project Settings File (*.FPProj) is an XML file containing the project settings and Data Manager definitions.

The TRAC Database File (*.db) is a SQLite database file that contains information pertaining to the current forecast period and all previously archived forecasts.

The Project Snapshot File (*.fcb) is a binary file that saves the forecast components that are currently in memory (e.g., the forecasts, summary statistics, etc.) so that the session can be restored without having to read in the data, recalculate the forecasts, etc. This allows you to return to your session where you left off and to share the session with others.

Zipped vs. unzipped projects

Forecast Pro allows you to save projects in an unzipped form. This is usually not recommended unless there is a specific issue with your operating environment that makes the unzipped format preferable.

In networked environments, working with zipped projects usually offers better performance. When you open a zipped project, Forecast Pro copies the zipped file from the network drive to a local workspace, unzips the project files and opens the project—allowing you to work on your project locally, even though it was saved to a network drive. When you resave your zipped project, Forecast Pro will zip the updated local project files and save the updated .FPZip file to the network drive.

Forecast Pro defaults to saving zipped projects. To save and open unzipped projects, you select FPProj as the file type in the **File>Save as** and **File>Open** dialog boxes.

Numeric Output Files

Numeric Output Files can contain time series output such as history, forecasts, fitted values and confidence limits as well as statistical output such as the forecasting model specification and summary statistics. These files can be saved in text, spreadsheet, ODBC or XML formats. You can include output for either all items forecasted or just for the Hot List. Generally speaking, if you will be importing the forecasts into another application these are the files you will want to use.

You can specify the format and content of Numeric Output File using **Settings>Numeric Output Design**. The operation of the Numeric Output Design dialog box is described in the Command Reference section of this manual.

You can preview the contents of the current Numeric Output File using the Numeric Output view. You can open this view using **View>Numeric Output** or using the yellow Preview Numeric Output icon ( ).

To save a Numeric Output file for all items forecasted, select **Project>Export>Full Numeric Output** or click the purple Save Numeric Output icon ( ).

To save a Numeric Output file for the current Hot List you select **Project>Export>Hot List Numeric Output** or use the **Save Hot List Numeric Output** option on the Hot List's context menu.

Formatted Forecast Reports

Formatted Forecast Reports can be saved to Excel. You can save these reports for the currently displayed item, all items on the Hot List or all items forecasted. The Excel workbook will include a separate worksheet for each item. These reports are convenient if you wish to present the forecasts and/or models used to colleagues.

Each report contains the information found in the current Forecast Report view and (optionally) the graph found in the current graph view.

You can specify the format and content of Formatted Forecast Report using **Settings>Forecast Report Design** and **Settings>Graph Settings**. The operation of the Forecast Report Design and Graph Settings dialog boxes are described in the Command Reference section of this manual.

To save a Forecast Report file for all items forecasted, you select **Project>Export>Full Forecast Report** or click the purple Save Formatted Forecast Report icon ( ).

To save a Forecast Report file for the current Hot List you select **Project>Export>Hot List Forecast Report** or use the **Save Hot List Forecast Reports** option on the Hot List's context menu.

You can print a forecast report for the currently displayed item using **File>Print**, the blue Print icon or the **Print this page** option on the Forecast Report view's context menu. You can also save an Excel report for the

currently displayed item using the **Save this page to Excel** option on the Forecast Report view's context menu.

Tracking Reports

Tracking Reports can be saved to Excel. You can save these reports for the currently displayed item, all items on the Hot List or all items forecasted. The Excel workbook will include a separate worksheet for each item. These reports are only available if the database includes at least one archived forecast for the given item.

Each report contains the information found in the current Forecast Report view and (optionally) the graph found in the current graph view.

You can specify the format and content of Tracking Reports using **Settings>Tracking Report Settings**. The operation of the Tracking Report Settings dialog box is described in the Command Reference section of this manual.

To save a Tracking Report file for all items forecasted, select **Project>Export>Full Tracking Report** or click the purple Save Forecast Tracking Report icon.

To save a Forecast Report file for the current Hot List, select **Project>Export>Hot List Tracking Report** or use the **Save Hot List Tracking Report** option on the Hot List's context menu.

You can print a tracking report for the currently displayed item using **File>Print**, the blue Print icon or the *Print this page* option on the Tracking Report view's context menu. You can also save an Excel report for the currently displayed item using the **Save this page to Excel** option on the Tracking Report view's context menu.

Item Reports

*Item Report*s can be viewed on-screen and saved to Excel. These reports can contain Pareto information, statistical output such as the forecasting model specification and summary statistics as well as time series output such as history, forecasts, fitted values and confidence limits.

To open an Item Report you select **View>Item Report** or click the yellow View Item Report icon ( ). To save an Item Report you select **Project>Export>Item Report** or use the view's context menu.

You can specify the format and content of the report using **Settings>Item Report Design**. The operation of the Item Report Design dialog box is described in the Command Reference section of this manual.

Override Reports

Override Reports can be viewed on-screen and saved to Excel. These reports list the items where overrides and/or comments have been made.

To open an Override Report you select **View>Override Report** or click the yellow View Override Report icon ( ). To save an Override Report, select **Project>Export>Override Report** or use the view's context menu.

You can specify the format and content of the report using **Settings>Override Report Design**. The operation of the Override Report Design dialog box is described in the Command Reference section of this manual.

Exception Reports

*Exception Report*s can be viewed on-screen and saved to Excel. These reports list the items where your forecast error or some other performance metric has fallen outside of an acceptable range. Exception reporting reduces the need for manual review of your forecasts and allows you to focus on the items where human attention is most needed.

To open an Exception Report you select **View>Exception Report** or click the yellow View Exception Report icon ( ). To save an Exception Report you select **Project>Export>Exception Report** or use the view's context menu.

Forecast Pro provides a wide array of exception reports, some of which monitor the current forecasts and others which monitor your archived forecasts. You can specify the format and content of the report using **Settings>Exception Report Settings** or selecting **Exception Report Settings** on the Exception Report context menu. The operation of the Exception Report

Settings dialog box is described in the Command Reference section of this manual.

Outlier Reports

Outlier Reports can be viewed on-screen and saved to Excel. These reports list the items where outliers have been detected and/or corrected.

To open an Outlier Report you select **View>Outlier Report** or click the yellow View Outlier Report icon (). To save an Outlier Report you select **Project>Export>Outlier Report** or use the view's context menu.

The content and format of the report is controlled using the Outlier Report's context menu.

Chapter 4

Using the Data Manager

The Data Manager provides a centralized place to define all of the input data for Forecast Pro and keep it up to date.

Row Types

You invoke the Data Manager by clicking the red Data Manager icon ( ) or selecting **Operations>Data Manager**.

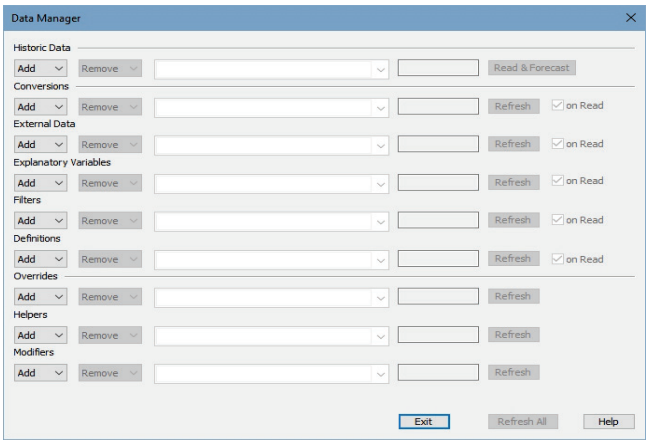

The Data Manger consists of eight rows each of which allows you to read in a different type of data.

In addition to the *Historic Data* (which is required), Forecast Pro's Data Manager allows you to import other information to assist you in creating and working with the forecasts. Specifically, you can import:

Conversions Factors. Conversion factors allow you to display the history and forecasts in alternative units of measure (e.g., dollars, kilos, etc.).

Explanatory Variables. Variables that you may wish to include in a dynamic regression model may be read in with an explanatory variable file.

External Data. Forecast Pro allows you to import external data rows (such as alternative forecasts, open orders, etc.) into the forecast override grid.

Filter Fields. Filter fields can be included in report views and filtered like any other fields. They can also be included in the numeric output file.

Definitions. Definition files contain values to use in conjunction with custom forecast allocations and item-level integer rounding quantities.

Overrides. In addition to entering forecast overrides interactively using the Forecast Overrides view, Forecast Pro allows you to import them.

Helper Variables. Helper variables are used in conjunction with event models, forecast by analogy models and weighting transformations.

Modifiers. In addition to specifying forecast modifiers interactively on the Navigator, Forecast Pro allows you to import them.

Row Controls

Each row has a number of controls associated with it allowing you to specify data sources, remove data sources and read the data.

Add drop down. The Add drop down is used to select data sources. You first specify the type of data source you will be adding (Excel, MLT or ODBC) and then the program will prompt you to specify the file, table or query. The Overrides, Helpers and Modifiers rows also allow you to specify a Forecast Pro project as the data source. If you wish to specify multiple data sources for a given row, you may. The first time you add an ODBC data source you will be prompted to choose an ODBC driver and connect to the database.

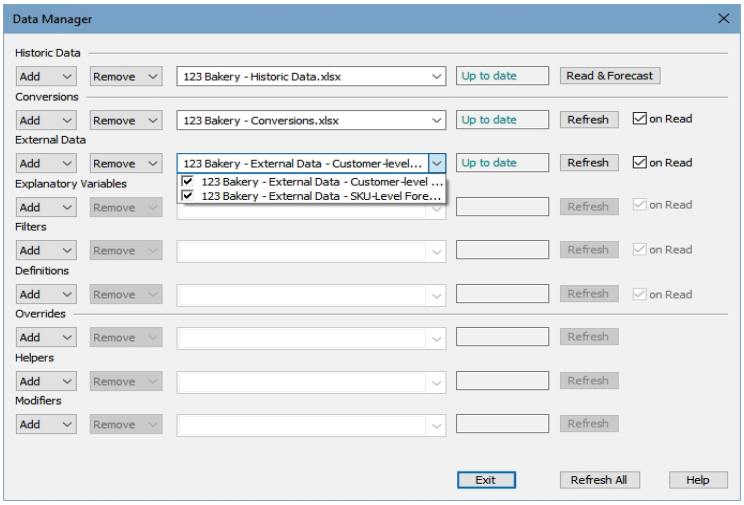

Selected data sources drop down. After you add one or more data sources to a given row, the filenames, table names, query names, or project names will appear on the selected data sources drop down display. If you open this drop down, you will see a list of the selected data sources along with a checkbox to the left of each one. This check box is used to mark the source as active or inactive. When you read the data, only the active data sources will be read.

Remove drop down. The remove drop down is used to remove data sources from the row. You have the option of either removing all specified data sources or only the inactive (unchecked) ones. After you use the Remove drop down you will need to refresh the row to actually delete the data from the current project.

Status box. The status box displays the current status for the data sources specified on the row. “Refresh pending” means that the current row specifications do not match the project’s current state and you need to refresh the row to sync them with the project.

Read & Forecast button. The Historic Data row includes a button labeled “Read & Forecast”. Clicking the button will read in the historic data and generate the forecasts. It will also refresh (i.e., read or re-read) any data rows where the “on read” option is selected.

Refresh button. The refresh button is used to read (or re-read) the data for the specified row.

On Read checkbox. Four of the rows include an “on Read” checkbox. If this option is selected Forecast Pro will automatically refresh the row whenever you read in the historic data. The on Read option is useful when the associated data will be changing each forecast period.

Notice that the Overrides, Helpers and Modifiers rows do not have on Read options and therefore can only be refreshed by clicking their Refresh buttons. This is due to the fact that overrides, modifiers and helper variables can all be modified interactively using Forecast Pro, and therefore refreshing them automatically from the data source may have unintended consequences (e.g., overriding your manual changes, etc.).

Refresh All button. The Refresh All button will refresh all data rows with the exception of the Historic Data. Note that it will refresh all rows that have defined data sources even if their status is up to date.

Chapter 5

Working with Hierarchies

This chapter discusses the strategies that Forecast Pro uses to produce consistent forecasts for hierarchical data.

Defining the Hierarchy

Product data can almost always be organized into several levels of aggregation. Suppose that the lowest-level forecasts you need to generate represent SKUs (Stock Keeping Units). SKU-level forecasts are often needed to support production planning and inventory control. The corporation might aggregate these SKUs first into products and then into product lines for marketing and sales. These might be aggregated further into geographical regions for the benefit of top management.

For the purposes of this discussion, we will be referring to *end-items* and *groups.* End-items are the lowest-level data in your hierarchy. In the above example, each SKU would be an end-item. Groups are aggregations of end-items. In the above example, products, product lines and geographical territories would all be examples of groups.

Forecast Pro allows you to define product hierarchies and create one set of self-consistent forecasts. It allows top-down, bottom-up or user-defined reconciliation, seasonal adjustment based upon aggregate data and model selection at the aggregate level.

It is *not* necessary that the end-item histories begin and end at the same time. Thus group-level data may consist of end-items that have been retired or replaced by new end-items. Obsolete end-items will contribute to the group-level history but will not themselves be forecasted. You will notice that the starting and ending dates for the overall historic data consist of the starting date for the oldest end-item and the ending date for the newest. Forecasts will be prepared for all end-items and groups that are "alive" at the end of the data set. Those whose histories terminate before that time are considered dead—they contribute to the historic aggregates (and therefore influence aggregate forecasts), but they are not themselves forecasted.

How to organize your data into hierarchies was discussed in the *Setting Up Your Historic Data* chapter.

Reconciling the Hierarchy

If no group-level reconciliation modifiers have been specified, a bottom-up approach will be used to reconcile the forecasts. The procedure operates as follows.

First Forecast Pro prepares forecasts for each and every group and end-item.

Then Forecast Pro recomputes the group-level forecasts by aggregating the constituent forecast (bottom up). The original group-level forecasts are replaced, but the width of their confidence limits are retained and re-centered on the new forecasts.

Three group-level modifiers can be used to specify alternative reconciliation approaches. These modifiers can only be used on group-level data.

\TOPDOWN directs Forecast Pro to perform top-down reconciliation for all members of the indicated group.

This procedure begins by preparing forecasts for each and every group and end-item.

Then it proportionally adjusts the nested lower-level end-item and group forecasts to sum to the \TOPDOWN group forecast. If the \TOPDOWN group is itself nested within larger groups, the forecasts for the larger groups are

computed by the bottom-up approach. (An example of top-down reconciliation is given below.)

You cannot define a \TOPDOWN group as a member of another \TOPDOWN group.

\CALL (Custom Allocation) directs Forecast Pro to use user-defined proportionalities to allocate the group-level forecast down to the next level.

You must define the proportions in a file, table, or query and read it into Forecast Pro using the Definitions row in the Data Manager. Consult the *Setting Up Your Optional Data* chapter for details.

The procedure begins by creating forecasts for the specified group and then using the user-defined proportional factors to allocate the group forecast down one level. Any additional lower levels are adjusted via a top-down procedure unless they too have a **\CALL** defined.

\INDEXES directs Forecast Pro to calculate the seasonal indexes at the indicated group level and use them to deseasonalize the histories for the nested end-items and groups. The deseasonalized series are then forecasted and reseasonalized. This approach works well when the end-items share the same seasonality. It allows one to deal with end-items whose histories are too short to extract seasonality directly. You cannot use \INDEXES if the group level includes zeros or negative numbers. You cannot nest \INDEXES groups within each other. Model parameters are reestimated for each end-item and nested group of the \INDEXES group.

The following screenshot illustrates the use of the \TOPDOWN modifier.

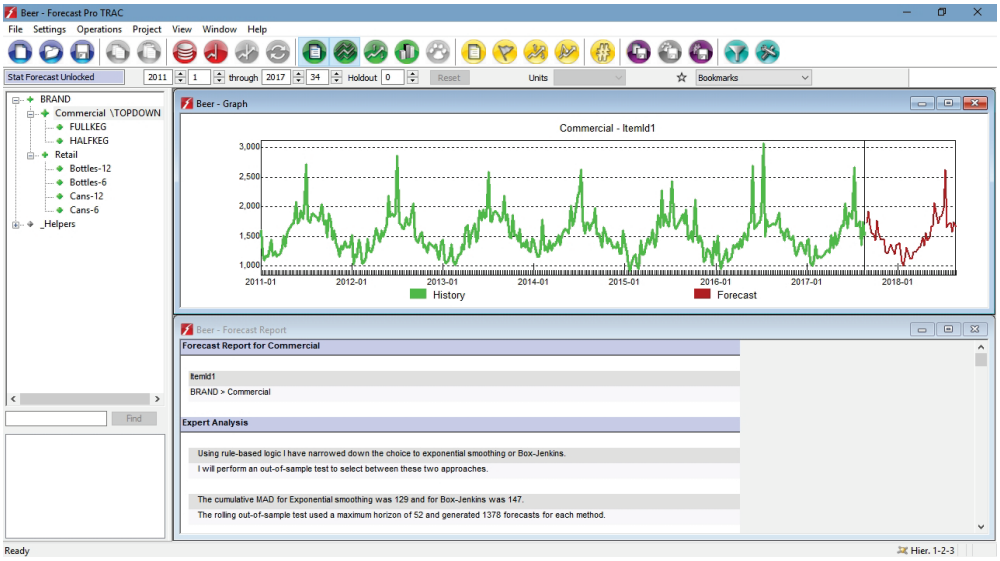

In the example above, \TOPDOWN directs Forecast Pro to compute forecasts and confidence limits for the groups Commercial and Retail at the group level. The end-item forecasts and confidence limits are multiplicatively adjusted so that the end-item forecasts sum correctly to the group level forecasts. Forecasts for the group Brand are obtained by bottom-up aggregation of the groups Retail and Commercial.

Note the assumptions behind this approach. We have assumed that the nested end-items are statistically similar, so that we can obtain better forecasts of the groups Retail and Commercial by forecasting them as a unit. However, Retail and Commercial show distinctly different patterns. Thus the best forecasts of Brand are obtained by direct aggregation of the two member groups.

Placing the \TOPDOWN flag on Brand rather than Retail and Commercial directs Forecast Pro to reconcile the entire hierarchy using a top-down approach. Usually, in a hierarchy involving three or more levels of aggregation, you will want to place the \TOPDOWN modifiers somewhere near the middle.

Two things happen as we move from item-level data upwards through the hierarchy. First, the effects of irregularity decrease through aggregation. Aggregate data series are typically more stable and reliably forecastable than disaggregate data. Second, statistical models become more complex and

distinct from each other. Thus the signal-to-noise ratio increases as we move upwards through the hierarchy.

Now consider the case where both the \TOPDOWN and \INDEXES keywords are used for Retail and Commercial. In this case, Forecast Pro extracts seasonal indexes for these groups at the group level. It then deseasonalizes each end-item's history using the indexes from the parent group, forecasts the resulting nonseasonal data, and then reseasonalizes the forecasts. The forecasts are then reconciled using a top-down procedure as already described.

A note about negative values

In a multiple-level problem, it is assumed that the data are basically nonnegative—the aggregate level data *must* be non-negative, but the item level data can include a few negatives.

Negative sales figures are relatively common since many companies register returns as negative sales. Furthermore, returns are often accumulated on the books and taken as occasional "hits" on sales. While this may make good accounting sense, it raises forecasting problems and will certainly decrease accuracy. Ideally, companies would either distribute returns to the months in which the products were shipped or maintain returns as an entirely different historic record. Nevertheless, listing sales as negative is common and must be dealt with in some way.

Beyond an expected decrease in accuracy for forecasts from a product with negative sales, negatives pose a technical problem for product hierarchies. Multiplicative seasonal indexes cannot be extracted from a nonpositive series, and additive indexes cannot be disaggregated to nested products. Thus, if the group level data contain negatives, Forecast Pro will use an additive index model and seasonal disaggregation will not be allowed (i.e., the \INDEXES flag will be ignored). In the case of positive group level data with negative values in the constituent data, the multiplicative adjustment of negative values may have a different effect than one expects—a seasonal index of 1.5 to a sales figure -100 yields the value -150.

Top-down adjustment is also problematic when negative *forecast* values are involved. Therefore, Forecast Pro clips negative forecasts to zero if you have specified top-down disaggregation, regardless of how you have set *Allow negative forecasts* in the Options dialog box in the **Settings** menu.

Chapter 6

Using Forecast Modifiers and Model Specification Dialog Boxes

By default, Forecast Pro will automatically select a forecasting model for each item on the Navigator using expert selection. The expert selection option works extremely well and is the method of choice for the majority of Forecast Pro users. The expert selection algorithm is described in the *Forecast Pro Statistical Reference Manual*.

Alternatively you can dictate the models and/or forecasting options using *forecast modifiers*. Forecast modifiers are added to items on the Navigator using the Navigator's context menu often with the assistance of a model specification dialog box.

This chapter documents the available modifiers and the operation of the model specification dialog boxes.

Model Specification Modifiers and Dialog Boxes

Supported modifiers and their associated dialog boxes are listed below grouped by model type.

Expert Selection

Model>Expert Selection (no modifier). If an item on the Navigator does not contain a modifier then Forecast Pro will use its expert selection algorithm to automatically select the appropriate forecasting method. If an item on the Navigator does contain one or more modifiers and you select *Model>Expert Selection,* the modifiers will be removed and Forecast Pro will use its expert selection algorithm to automatically select the appropriate forecasting method.

Exponential Smoothing

Selecting **Model>Exponential Smoothing>Custom** on the Navigator context menu invokes the Exponential Smoothing Model Settings dialog box shown below.

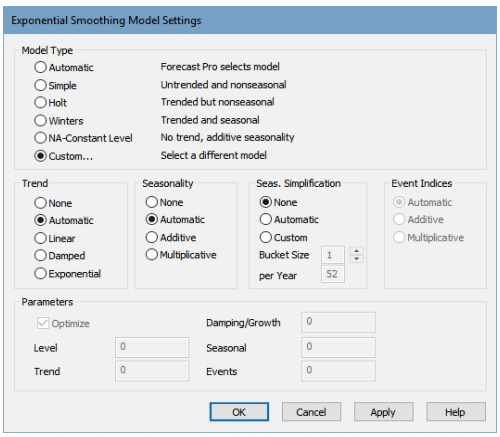

The Model Type section of the dialog box allows you to specify Automatic selection (i.e., Forecast Pro will automatically select the form of the exponential smoothing model to apply), one of the standard models or Custom. If you select Custom, you will need to specify the form of Trend and Seasonality to use.

If you are forecasting data where the periods per cycle is greater than 13 (e.g., weekly data) you can use the Seasonal Simplification options to reduce the number of seasonal indices to use. *None* instructs Forecast Pro to not use seasonal simplification. *Automatic* allows Forecast Pro to automatically select the number of indices using out-of-sample testing. *Custom* allows you to specify the number of indices by specifying the “bucket size” (i.e., how many

periods to combine into each seasonal index) or by setting the number of indices to use per year. Note that seasonal simplification cannot be used in conjunction with event models.

The Event Indices section allows you to specify the form of the event indices (additive, multiplicative or let Forecast Pro decide). This section will only be available if you have specified an event schedule. Furthermore, when events are used in conjunction with a seasonal model, the form of both the event indices and seasonal indices must match—so the Event Indices section will also be greyed out if Seasonality is set to anything other than None.

The Parameters section allows you to set the smoothing weights to specific values. In general this is not a recommended practice (allowing Forecast Pro to optimize the smoothing weights is recommended). If you do wish to set the smoothing weights, you must specify a specific form for the trend and seasonality.

When you click the OK button Forecast Pro will apply the appropriate modifier on the Navigator and build the specified model. The exponential smoothing modifiers are:

Model>Exponential Smoothing>Auto: *\EXSM.* Use the automatic fitting exponential smoothing model.

Model>Exponential Smoothing>Simple: *\SIMPLE.* Use the simple exponential smoothing model.

Model>Exponential Smoothing>Holt: *\HOLT.* Use the Holt exponential smoothing model.

Model>Exponential Smoothing>Winters: *\WINTERS.* Use the Winters exponential smoothing model.

Model>Exponential Smoothing>NA-Constant Level: *\NA-CL.* Use the Non-trended Additive Seasonality Constant Level exponential smoothing model.

Model>Exponential Smoothing>Custom: *\EXSM=XY.* Use a custom exponential smoothing model with trend type X (N=no trend, L=linear trend, D=damped trend, E=exponential, *=Forecast Pro decides), seasonality type Y

(N=nonseasonal, M=multiplicative seasonal, A=additive seasonal, *=Forecast Pro decides) and optimized smoothing weights.

Model>Exponential Smoothing>Custom: *\EXSM=XY(A,B,C,D).* Use a custom exponential smoothing model with trend type X (N=no trend, L=linear trend, D=damped trend, E=exponential, *=Forecast Pro decides), seasonality type Y (N=nonseasonal, M=multiplicative seasonal, A=additive seasonal, *=Forecast Pro decides) and user defined smoothing weights (A=level, B=trend, C=damping/growth, D=seasonal).

Model>Exponential Smoothing>Custom: *\SS.* Use Forecast Pro's automatic identification procedure to determine whether to use seasonal simplification and the appropriate bucket size.

Model>Exponential Smoothing>Custom: *\SS=n.* Use seasonal simplification with bucket size equals n.

Box-Jenkins

Selecting **Model>Box-Jenkins>Custom** on the Navigator context menu invokes the Box-Jenkins Model Settings dialog box shown below.

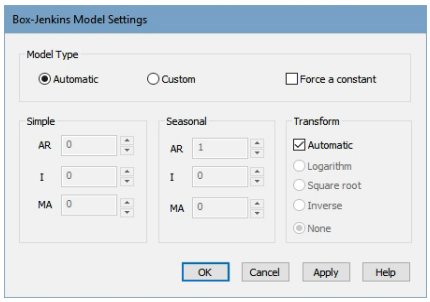

The Model Type section of the dialog box allows you to specify Automatic (i.e. Forecast Pro will automatically determine the ARIMA model orders) or Custom (i.e. you will explicitly select the model orders to use). In addition, there are options to include a constant term in the model and to apply a power transformation.

When you click the OK button, Forecast Pro will apply the appropriate modifier on the Navigator and build the specified model. The Box-Jenkins modifiers are:

Model>Box-Jenkins>Auto: *\BJ.* Use the automatic fitting Box-Jenkins model.

Model>Box-Jenkins>Custom: *\ARIMA(p,d,q).* Use a non seasonal Box-Jenkins model with model orders p, d and q.

Model>Box-Jenkins>Custom: *\ARIMA(p,d,q)*(P,D,Q).* Use a seasonal Box-Jenkins model with model orders p, d, q, P, D and Q.

Model>Box-Jenkins>Custom: *\CONST.* Include a constant intercept in the Box-Jenkins model.

Discrete Data

Model>Discrete: *\DISCRETE.* Use a simple exponential smoothing model and base confidence limits on one of the discrete distributions (Poisson or negative binomial). This option is used to obtain better estimates of the confidence limits for low volume integer series (typically with many zeros).

Intermittent Data

Model>Intermittent: *\INTER.* Use the Croston's intermittent data model. The Croston's model is designed for data sets where the demand for any given period is often zero and the exact timing of the next order is not known.

Very Simple Models

Several of the menu options available under **Model>Very Simple Models** invoke the dialog box below.

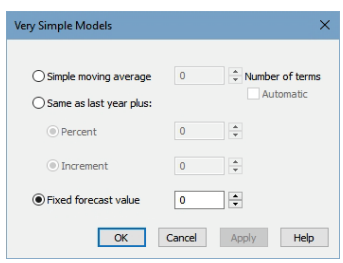

Simple Moving Average allows you to either set the number of observations to use in the moving average or have Forecast Pro calculate the number of terms to use automatically.

Same as last year sets each forecasted value to equal the value for the same period last year (either history or forecast as the case may be). The *Percent* option sets each forecasted value to equal the value for the same period last year plus the specified percentage. The *Increment* option sets each forecasted value to equal the value for the same period last year plus the specified increment.

Fixed forecast value sets the forecast to the specified value.

Selecting **Model>Very Simple Models>Fixed Forecast Value After** invokes the dialog box shown below that sets the forecast to the specified value after the specified date. This option is often used when you plan to discontinue an item at a specific time and wish to set the forecast to zero.

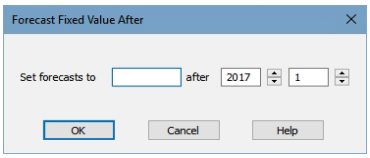

The modifiers associated with very simple models are:

Model>Very Simple Models>Simple Moving Average: *\SMA=n.* Use an n-term simple moving average. If the Automatic option is selected, Forecast Pro will choose the number of terms to use and the forecast modifier will be displayed as *\SMA*.

Model>Very Simple Models>Same as Last Year>No Change: *\SALY.* Set each forecasted value to equal the value for the same period last year (either history or forecast as the case may be).

Model>Very Simple Models>Same as Last Year>Plus Percent: *\SALYP=n.* Set each forecasted value to equal the value for the same period last year plus the specified percentage (*n*).

Model>Very Simple Models>Same as Last Year>Plus Increment: *\SALYI=n.* Set each forecasted value to equal the value for the same period last year plus the specified increment (*n*).

Model>Very Simple Models>Same as Last Year>Plus Delta Percent: *\SALYDP.* Set each forecasted value to equal the value for the same period last year adjusted by the percentage change between the value for the same period last year and the same period two years ago. That is, the model assumes that sales will grow at the same percentage rate as last year.

Model>Very Simple Models>Same as Last Year>Plus Delta Increment: *\SALYDI.* Set each forecasted value to equal the value for the same period last year adjusted by the unit change between the value for the same period last year and the same period two years ago. That is, the model assumes that sales will grow by the same number of units as last year.

Model>Very Simple Models>Fixed Forecast Value: *\FIXED=n.* Set each forecasted value to *n*.

Model>Very Simple Models>Fixed Forecast Value After: *\FIXEDAFT=(Year, Period, n).* Set each forecasted value which falls after the year and period specified to *n*.

Curve Fitting

Models>Curve Fit>Automatic: *\CFIT.* Use the curve fitting routine that best fits the historic data.

Models>Curve Fit>Straight Line: *\CFIT=LINE.* Fit a straight line to the data set.

Models>Curve Fit>Quadratic: *\CFIT=QUAD.* Fit a quadratic curve to the data set.

Models>Curve Fit>Exponential: *\CFIT=EXPO.* Fit an exponential curve to the data set.

Models>Curve Fit>Growth Curve: *\CFIT=GROW.* Fit a growth curve (S-curve) to the data set.

Bass Diffusion

The Bass model is most often used to forecast first time purchases of new-to-world products.

The model tries to capture the adoption rates of two types of users—innovators and imitators. Innovators are early adopters of new products and are driven by their desire to try new technology. Imitators are more wary of new technology—they tend to adopt only after receiving feedback from others.

The Bass model uses two coefficients to quantify the adoption rates. The *Coefficient of Innovation,* referred to in the literature as "p", controls the adoption rate for the innovators. The *Coefficient of Imitation,* referred to in the literature as "q" controls the adoption rate for the imitators.

If you have 5 or more historic data points, these coefficients can be fit to the data. To build a Bass model with fewer than 5 data points you must set the values for these coefficients along with the total number of potential adopters. Typically, "p" is between 0 and 0.1, while "q" is between .3 and .5.

Selecting **Model>Bass Diffusion** on the Navigator context menu invokes the Bass Diffusion Model Settings dialog box shown below.

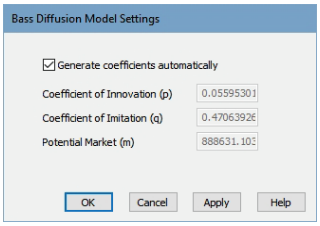

If you have five or more data points, the *Generate coefficients automatically* option allows you to instruct Forecast Pro to automatically fit the model parameters to the data. If you do not select this option then you'll need to supply the model coefficients and potential market size.

The modifiers associated with the Bass diffusion model are:

Model>Bass Diffusion: *\BASS.* Use the Bass diffusion model with automatically generated coefficients and potential market size.

Model>Bass Diffusion: *\BASS(p,q,m)*. Use the Bass diffusion model with user defined parameters p, q and m.

By Analogy

The concept behind the Forecast by Analogy model is very simple. The approach is sometimes referred to as a "looks like" analysis. The model assumes that the initial sales over the first few time periods for a new product will follow the same pattern as the initial sales of an existing, similar product as it launched. To use this method, you must supply the launch profile or "analog series" in the form of a helper variable. This launch profile is simply sales for the analog product, starting with the launch period and extending for as many time periods as deemed relevant.

If the product has not yet launched (i.e., there are no historic data available), then you must also supply an estimate of the initial sales over a specific period of time (the "launch total" over the "launch horizon"). Forecast Pro will then create the forecast by proportionally allocating the launch total over the launch horizon using the analog series to define the proportions.

If any historic demand for the new product is available, you may either supply the launch total and launch horizon or allow Forecast Pro to calculate the forecasts by fitting the existing data to the analog series.

Selecting **Model>By Analogy** on the Navigator context menu invokes the Forecast by Analogy Model Settings dialog box shown below.

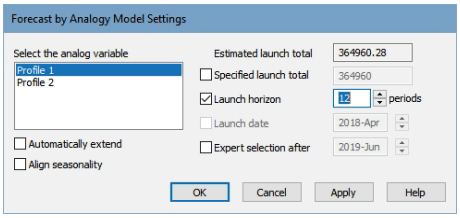

The *Select the analog variable* box displays the currently available helper variables and allows you to select the one you'd like to use as the analog series. If historic demand exists, the *Estimated launch total* displays the launch total calculated using the current analog variable and the historic demand. The *Specified launch total* allows you to input a specific launch total to use. The *Launch horizon* allows you to specify the length of the launch

horizon. The default launch horizon equals the length of the history, if any, plus the current forecast horizon.

The *Launch date* option is only available if the product has not yet launched (i.e., there is no historic demand for the product). It allows you to specify when the product will come on line, allowing you to set up the forecasting model in advance of the launch date.

The *Automatically extend* option will automatically forecast the analog variable if it does not cover the full launch horizon. The *Align seasonality* option will (if necessary) adjust the beginning period of the analog series so that it aligns with the launch date's period. For example, if the analog series started in January and the launch date corresponded to March, then selecting this option would result in using the analog series' first March observation as the first analog series data point (thus aligning the seasonal patterns of the analog series and the new product).

Expert selection after is used to specify a date where Forecast Pro should stop using the Forecast by Analogy method and switch to using expert selection. For example if you set *Expert selection after* to January 2019, Forecast Pro would use the Forecast by analogy approach until your historic data included February 2018.

The modifiers associated with Forecast by Analogy are:

Model>By Analogy: *\ANALOG=_X.* Use the by analogy model with automatic calculation of the launch total. _X is the name of the time series containing the analogy series.

Model>By Analogy: *\ANALOG=_X(A,B).* Use the by analogy model with a specified launch total of A and a specified launch horizon of B.

Model>By Analogy: *\ANALOG=_X(A,B,C,D).* Use the by analogy model with a specified launch total of A, a specified launch horizon of B and a specified launch date with a starting year of C and starting period of D. This option is only available when the series being forecasted has no historic data.

Model>By Analogy: *\ANALOGEX, \ANALOGAL, \ANALOGD.* There are three suffixes that may appear after the keyword ANALOG. EX appears if the *Automatically extend* option is selected. AL appears if the *Align seasonality*

option is selected. D appears if the *Expert selection after* option is selected. If all three options are selected, all 3 suffixes will appear after the keyword ANALOG.

Component Model

The custom component model includes each of the components found in a custom exponential smoothing model, specifically, level, trend, seasonality and events. In a standard smoothing approach, final values of each component are estimated from the data and used to assemble the forecasts. In a custom component model you have the ability to either let Forecast Pro estimate the final value for a given component or to customize the values to be used.

Custom component models are useful in a variety of situations including (1) customizing the trend for longer-term forecasts, (2) customizing the seasonal pattern for short data sets and (3) defining the impact of future events that have not occurred historically.

Selecting **Model>Component** on the Navigator context menu invokes the Component Model Settings box shown below.

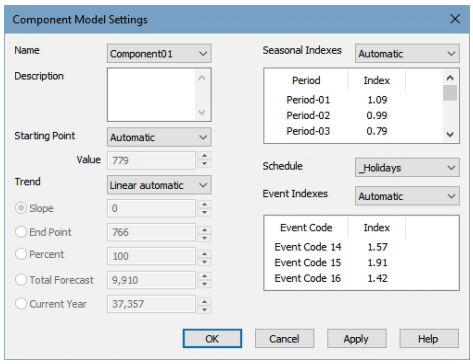

Name: Forecast Pro names each of your component model specifications and saves them in the project's data base. These named specification sets provide a convenient way to apply the same component model specifications to multiple items on the Navigator. The *Name* drop down allows you to select previously defined specification sets, create new ones, save the current set using a different name and delete the current set.

Description: The description field allows you to enter a description for the current specification set.

Starting Point: The Starting Point drop down allows you to set the starting point for the forecasts. *Automatic* instructs Forecast Pro to estimate an exponential smoothing model for the data and set the starting point to the final level. *Custom* allows you to enter a specific value to use as the starting point. *Final Historic* sets the starting point to the final historic observation.

Trend (None and Linear): The trend drop down allows you to specify the form of the trend you wish to forecast. All trends will begin at the starting point. *None* will yield a flat-line trend. *Linear automatic* instructs Forecast Pro to estimate an exponential smoothing model for the data and set the trend to the final trend estimated for the model (this will either be linear or flat-line). *Linear custom* will default to the linear automatic value but the slope can then be modified using the radio buttons (discussed below).

Trend (Bent): You can modify a linear trend by "bending" it up or down. When you bend a trend downward you decrease the per-period-slope increment resulting in a trend that dampens out as you forecast forward. When you bend a trend upward you increase the per-period-slope increment resulting in a trend that accelerates as you forecast forward. *Bent automatic* will use the linear automatic trend as the baseline which can then be bent using the radio buttons (discussed below). *Bent custom* will default to using the linear automatic trend as the baseline but allows you to change the slope for the baseline trend (using the Slope radio button) as well as allowing you to bend the baseline trend using the radio buttons (discussed below).

Trend (radio buttons): The trend radio buttons are used to customize the slope. When the Component Model Settings dialog box is open, the graph includes the current committed trend. The *Slope* button allows you to modify the slope (per-period change) directly. The *End Point* allows you to set the value of the trend at the last forecast point. The difference between the trend's starting point and ending point can be thought of as the "vertical rise." *Percent* allows you to specify the ending point of a bent trend model as a percentage of the vertical rise. *Total Forecast* allows you to set a value for the sum of the total forecast. *Current Year* allows you to set a value for the sum of the current year's historic + forecast values.

Seasonal Indexes: The Seasonal Indexes drop down allows you to specify the multiplicative indexes you wish to use to forecast the seasonality. *None* will yield a nonseasonal forecast. *Automatic* instructs Forecast Pro to estimate an exponential smoothing model for the data and set the seasonal indexes to the final seasonal indexes estimated for the model. *Custom* allows you to edit the currently displayed values. *Helpers* allow you to import the indexes from a helper variable. If the helper variable contains more than one year's worth of indexes, Forecast Pro will import the final year's values.

Schedule: Schedule allows you to specify an event schedule to use.

Event Indexes: Event Indexes allows you to specify the form and values for the event indexes. *None* will not apply event indexes to the forecast. *Automatic* instructs Forecast Pro to estimate an exponential smoothing model for the data and set the event indexes to the final event indexes estimated for the model. *Custom additive* allows you to edit the currently displayed values and apply them as additive indexes. Please note that the displayed values may correspond to a multiplicative model and, if so, should be changed to values that reflect incremental units driven by each event code. *Custom multiplicative* allows you to edit the currently displayed values and apply them as multiplicative indexes.

The following modifier is associated with the custom component model.

Model>Component: *\COMP=name.* Use a custom component model with the specifications defined in *name*.

Dynamic Regression

Dynamic regression produces a forecast based on the forecasted item's history (like univariate methods) and that of the explanatory variables (e.g., product promotion, advertising, demographic variables, macroeconomic indicators, etc.).

Selecting **Dynamic Regression>Manage** or **Models>Dynamic Regression** on the Navigator context menu invokes the Dynamic Regression dialog box shown below.

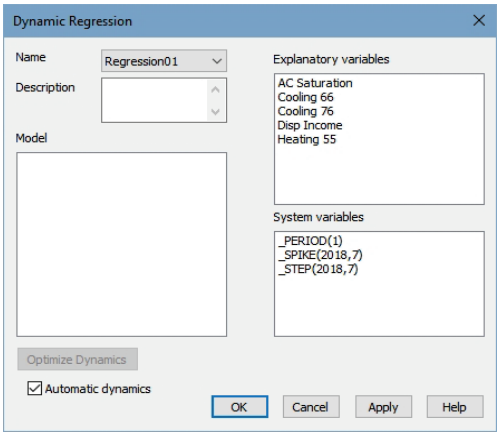

Name: Forecast Pro names each of your dynamic regression model specifications and saves them in the project's data base. These named specification sets provide a convenient way to apply the same dynamic regression model specifications to multiple items on the Navigator. The *Name* drop down allows you to select previously defined specification sets, create new ones, save the current set using a different name and delete the current set.

Description: The description field allows you to enter a description for the current model selection.

Explanatory variables: Explanatory variables allow you to capture how the dependent variable changes in response to external variables. The Explanatory variables field lists all explanatory variables (both global and item-specific) read in through the Data Manager. Please see the Explanatory Variables section of "Setting up your Optional Data" for details on the differences between global and item-specific explanatory variables.

System variables: There are two type of terms that can be included in a dynamic regression models—explanatory variables and dynamic terms. Explanatory variables are defined above. Dynamic terms capture how the dependent variable changes in time and can help forecast trends and seasonal patterns in a similar fashion to extrapolation methods such as Box-Jenkins and exponential smoothing.

The System variables field list dynamic terms as well as well as system terms for defining additional explanatory variables in Forecast Pro.

Each of these system terms starts with an underbar to identify it as a Forecast Pro defined variable. These terms include the following:

> _PERIOD(n) creates a dummy variable for period n. The dummy variable will have a value of 1 for the nth period in the seasonal cycle and a value of 0 for all periods. For example, if you are using monthly data, _Period(3) will have a value of 1 for every March and a value of 0 for all other time periods.
>
> _SPIKE(year, period) is used to create a dummy variable for a single period. For example, for a monthly data set, _SPIKE(2017,7) has a value of 1 for July 2017 and a value of 0 for all other periods. This function can be useful when modeling outliers.
>
> _STEP(year, period) is used to create a variable that consists of all zeroes prior to a specified time period and all ones thereafter. For example, for a monthly data set, _STEP(2017,7) will have a value of 0 prior to July 2017 and a value of one for July 2017 and thereafter. This function can be useful when modeling level shifts.

The Systems variables list box may also include the following dynamic terms:

> _CONST is used to include a constant term in the model. A constant term has a value of 1 for all time periods.
>
> _TREND is used to include a linear trend term in the model.
>
> _DEPENDENT[-n] is used to include a lagged dependent variable in the model. For example, if you are modeling the item *Sales*, DEPENDENT[-1] is *Sales* for the prior period.
>
> _AUTO[-n] is used to create a Cochrane-Orcutt model which includes a variable equal to the fitted model error from n periods earlier. For example, when you add _AUTO[-1] to the model, the program uses the fitting error from time t-1 to help predict the fitting error for time t.

Forecast Pro offers an Automatic dynamics mode whereby you specify the explanatory variables you wish to use and the program automatically determines the dynamic terms to include. When you use Automatic dynamics, the dynamic terms are not visible in the Dynamic Regression dialog box.

When you double click on an item in the Terms list box, it is added to the Model box. When you add _STEP or _SPIKE, you will be asked to specify a date to use before it is added. Once you have added a variable to the Model box, it can be edited using the context menu as described in the *Model* section below.

Model: The Model box lists all terms that you have added to the model. To the left of each item is a checkbox to indicate if the item should be included in the regression. If this box is checked, the term will be included in the model. If the box is not checked, the term will be tested in the dynamic regression test batteries but will not be included in the model. If Automatic dynamics is selected, the dynamic terms will not appear in the Terms and Model boxes, but an optimized set of dynamics terms will be included in the model.

Each term in the Model box can be adjusted using the Model box's context menu. Right click on the item in the Model box to activate the Model box's context menu.

> *Clear* will remove all terms from the dialog box.
>
> *Remove* will remove just the selected item from the Model box.
>
> *Lag* allows you to adjust the specified lag. You can use the lag selection to create lagged explanatory variables or to adjust the lags for the _AUTO and _DEPENDENT terms.
>
> *Transform* allows you to use a logarithm, square root or inverse transform on the selected variable. Transform is not available for dummy variables or if the transform cannot be applied to the selected variable (e.g. logarithm for items with zero values).
>
> *Edit* allows you to change the specified period for the dummy variables (_PERIOD, _STEP and _SPIKE).

Optimize Dynamics (button): This button will replace any existing dynamic terms in the model with a set of dynamic terms selected by Forecast Pro. Note that the optimal model dynamics are selected for the item currently selected on the Navigator and the specified non-dynamic terms in the Model box. Once the optimal dynamics are added—they are simply part of the specified

model—they will not be automatically updated if you subsequently change the model.

Automatic Dynamics (check box): Forecast Pro offers an Automatic dynamics mode whereby you specify the explanatory variables you wish to use and the program automatically determines the dynamic terms to include. Automatic Dynamics is selected by default.

There are important (but subtle) differences between a model specification that uses the Automatic Dynamics option and a model specification that used the Optimize Dynamics button to determine the dynamics.

1. If you apply the same model specification to multiple items on the Navigator, the model specification that uses automatic dynamics will select the dynamic terms independently for each item. The model specification built using the Optimized Dynamics button will simply apply the specified terms to each item.

2. If you update your project with new data and reforecast, the model specification that uses automatic dynamics will re-select the dynamic terms using the new data (potentially changing the selected terms). The model specification built using the Optimized Dynamics button will simply apply the previously specified terms to each item.

Note that the dynamic regression model specification can also be edited with the context menu in the Forecast Report view. You are able to right click on a term in the model details section of the report to activate the context menu and remove an item. The Variable Test Specification battery and Dynamics Test Battery sections of the report list explanatory variable and dynamic to consider adding to the model. To add a term to the model, right click on the term you wish to add and select Add term from the Forecast Report context menu.

In order to generate forecasts, Forecast Pro needs values for each explanatory variable for all periods in the forecast horizon. These values may be provided in the explanatory variables file, or you may use Auto Extend if you want Forecast Pro to use expert selection to generate forecasts for any explanatory variable that does not have values provided for all periods in the forecast horizon. Forecast Pro will provide values only for those periods where none

was provided. By default, Auto Extend is not selected. To select Auto Extend, go to the Advanced Controls tab on the Options menu.

Event Model Modifiers

Event adjustment models extend exponential smoothing by allowing you to adjust for events like sales promotions, strikes or simply for unexplained outliers. Multiple events are used to account for promotions of different types or sizes, different calendar effects like Easter and Independence Day weeks, or any other "out of the ordinary" occurrence.

To build an event model, you must tell the program when events of each type occur. To do this, you must construct an event schedule which classifies each period by event type (0=no event, 1=event of type 1, 2=event of type 2, etc.). You can create event schedules interactively using the Event Manager or create them outside of Forecast Pro and import them using the Helpers row in the Data Manager. When you create an event schedule interactively, Forecast Pro will create a helper variable with this schedule.

For examples of constructing event variables and building event models, consult *Building Event Models* in the tutorial section. For details on using the Event Manager consult *The Event Manager* in the command reference section. For statistical details, consult the section on exponential smoothing in the *Forecast Pro Statistical Reference Manual*.

The following modifiers are associated with event models.

Events>Select: *\EVENT=_X.* Use an event model with the helper _X.

Model>Exponential Smoothing>Custom: *\EXSM=XYZ.* Use an exponential smoothing model with trend type X (N=no trend, L=linear trend, D=damped trend, E=exponential, *=Forecast Pro decides), seasonality type Y (N=nonseasonal, M=multiplicative seasonal, A=additive seasonal, *=Forecast Pro decides) and event type Z (M=multiplicative, A=additive, *=Forecast Pro decides). This modifier can only be used in conjunction with \EVENT=.

Model>Exponential Smoothing>Custom: *\XYZ(A,B,C,D,E).* Use an exponential smoothing model with trend type X (N=no trend, L=linear trend, D=damped trend, E=exponential, *=Forecast Pro decides), seasonality type Y

(N=nonseasonal, M=multiplicative seasonal, A=additive seasonal, *=Forecast Pro decides), event type Z (M=multiplicative, A=additive, *=Forecast Pro decides) and user defined smoothing weights (A=level, B=trend, C=damping/growth, D=seasonal, E=event). This modifier can only be used in conjunction with \EVENT=.

The Weighting Transformation Modifier

The weighting transformation is most commonly used to deseasonalize your variables using externally supplied seasonal weights, or to normalize the data for trading day effects (e.g., 4-4-5 calendars, number of working days per month, etc.).

To use the weighting transformation you must create a *helper* variable containing the weights. Helper variables names *must* start or end with an underbar. Helper variables are not forecasted and their values are not included in group totals. Helper variables are used in conjunction with event models, by analogy models and weighting transformations.

The following modifier is used to specify the weighting transformation.

Weights>Select: *\WGT=_X.* Use a weighting transformation. _X is the name of the time series containing the weights.

The procedure divides each value of the specified time series by the corresponding value (weight) in _X. It then forecasts the deweighted variable and multiplies the forecasts by their corresponding weights. The weighting variable _X must span the entire history and forecast period for each variable to be forecasted.

Outlier Detection/Correction Modifiers

When you generate forecasts, Forecast Pro uses the current settings in the Outliers tab on the Options dialog box to determine the default behavior for detection and correction of outliers. There are three choices:

None. Do not detect nor correct for outliers.

Detection only. Detect and report outliers but do not correct for them (i.e., base the forecasts on the uncorrected values).

Detection and correction. Detect and correct outliers (i.e., base the forecasts on the corrected values).

You can override this project level-setting for individual items using the forecast modifiers listed below. Forecast Pro only detects outliers for end items (i.e., non-group level data), so the modifiers cannot be used for groups.

Outliers>Default: *(no modifier).* Revert back to the default project-level setting (specified on the Outliers tab on the Options dialog box) for the specified item. This option is used to remove any of the outlier modifiers (listed below) from the Navigator.

Outliers>Off: *OUTLIER=OFF.* Do not detect nor correct outliers for the specified item.

Outliers>Detect: *OUTLIER=DETECT.* Detect and report outliers for the specified item, but do not correct for them (i.e., base the forecasts on the uncorrected values).

Outliers>Correct: *OUTLIER=CORRECT.* Detect and correct outliers for the specified item (i.e., base the forecasts on the corrected values).

Reconciliation Modifiers

If no group-level reconciliation modifiers have been specified, a bottom-up approach will be used to reconcile the forecasts. The procedure operates as follows.

First Forecast Pro prepares forecasts for each and every group and item.

Then, Forecast Pro recomputes the group-level forecasts by aggregating the constituent forecast (bottom up). The original group level forecasts are replaced, but their confidence limits are retained and proportionately adjusted.

Three group-level modifiers can be used to specify alternative reconciliation approaches. These modifiers can only be used on group-level data.

Top down: *TOPDOWN.* Directs Forecast Pro to perform top-down reconciliation for all members of the indicated group.

This procedure begins by preparing forecasts for each and every group and item.

Then it proportionally adjusts the nested lower level item and group forecasts to sum to the \TOPDOWN group forecast. If the \TOPDOWN group is itself nested within larger groups, the forecasts for the larger groups are computed using a bottom-up approach.

Custom Allocation: \CALL. Directs Forecast Pro to use user-defined proportionalities to allocate the group-level forecast down to the next level.

You must define the proportions in a file, table or query and read it into Forecast Pro using the Definitions row in the Data Manager. Consult the *Setting up Your Optional Data* chapter for details.

The procedure begins by creating forecasts for the specified group and then using the user-defined proportional factors to allocate the group forecast down one level. Any additional lower levels are adjusted via a top-down procedure unless they too have a \CALL defined.

Indexes: \INDEXES. Directs Forecast Pro to calculate the seasonal indexes at the indicated group level and use them to deseasonalize the histories for the nested items and groups. The deseasonalized series are then forecasted and as a final step reseasonalized. This approach works well when the items share the same seasonality. It allows one to deal with items whose histories are too short to extract seasonality directly. You cannot use \INDEXES if the group level includes zeros or negative numbers. You cannot nest \INDEXES groups within each other. Model parameters are reestimated for each item of the \INDEXES group.

Miscellaneous Custom Modifiers

The modifiers listed below can be entered from the keyboard via **Custom Modifier** on the Navigator's context menu.

Custom Modifier: *LOWER=l*. Set the lower confidence limit equal to the value *l*. *l* must be between 0.1 and 50.0 inclusive.

Custom Modifier: *UPPER=u*. Set the upper confidence limit equal to the value *u*. *u* must be between 50.0 and 99.9 inclusive.

Custom Modifier: *CONF(l,u)*. Set the lower confidence limit equal to the value *l* and the upper confidence limit equal to *u*. *l* must be between 0.1 and 50.0 inclusive and *u* must be between 50.0 and 99.9 inclusive.

Custom Modifier: *LEADTIME=n*. Set the lead time for the safety stock to *n*. The specified lead time appears highlighted in the Safety Stock section of the Forecast Report and determines the values written to the Numeric Output file. Fractional entries (e.g., 1.5) are permitted.

Custom Modifier: *TIMEFENCE=n*. Set the length of the time fence for the specified item to *n*. Be aware that if you use the *TIMEFENCE* modifier to create a time fence shorter than the global time fence, it will shorten the time fence for all associated parent groups of the item.

Custom Modifier: *DISCONT*. Treat the specified item as discontinued. Discontinued items are not forecasted but their historic demand does go into group-level aggregates.

Custom Modifier: *DISCONTAFT(Year, Period)*. Treat the specified item as discontinued after the specified date.

Custom Modifier: *START(Year, Period)*. Only use historic data after the specified date to estimate the specified item's forecast model.

Custom Modifier: *POW=key*. Use a Box-Cox power transformation. Key may equal LOG, SQRT, INVERSE or any decimal number from -1.0 to 1.0. If you enter \POW=AUTO, Forecast Pro will test your data and implement an appropriate Box-Cox transformation (or none at all). Consult the *Forecast Pro*

Statistical Reference Manual for a discussion of the Box-Cox power transformations.

Chapter 7

Command Reference

This chapter presents a command overview, a description of the user interface, a detailed description of each menu command and instructions on how to drive the program from the command line.

Operations Overview

Generating forecasts in Forecast Pro consists of executing the following steps:

1. Prepare the database. This is performed outside of the program and entails selecting the type of data format to use and preparing the files. Creating your data files is discussed in the first two chapters of the Reference section.

2. Specify the default options. This is accomplished using the Settings menu. Options set by the user will apply to the current forecast document and can also be saved to the ForecastProTRAC.ini file and used as defaults for all future forecast documents using the “Set as Default” button. Thus, you need only set the options when you use Forecast Pro for the first time, or when you wish to change the current settings.

3. Define the historic data and other inputs. This is accomplished by specifying the appropriate entries in the Data Manager.

4. Read in the data and prepare the forecasts. This is accomplished by clicking the *Read & Forecast* button in the Data Manager. By default, Forecast Pro will automatically create forecasts using its expert selection algorithm and reconcile the forecasts using a bottom-up approach.

5. Save the results. This can vary depending upon your needs, but usually entails saving a numeric output file containing the forecasts and saving a forecasting project so that you can update it the next forecast period rather than creating a new one from scratch.

In addition to the above required steps, there are many optional operations that you can perform. Including:

6. View the forecasts and other information graphically.

7. View on screen reports.

8. View forecasts and other information numerically in the override grid and enter any desired changes and comments.

9. Customize how the forecasts are generated by specifying that specific forecasting models and reconciliation techniques be used.

The remaining sections of this chapter document the above forecasting procedure in more detail.

The Forecast Pro Interface

This section describes the Forecast Pro interface.

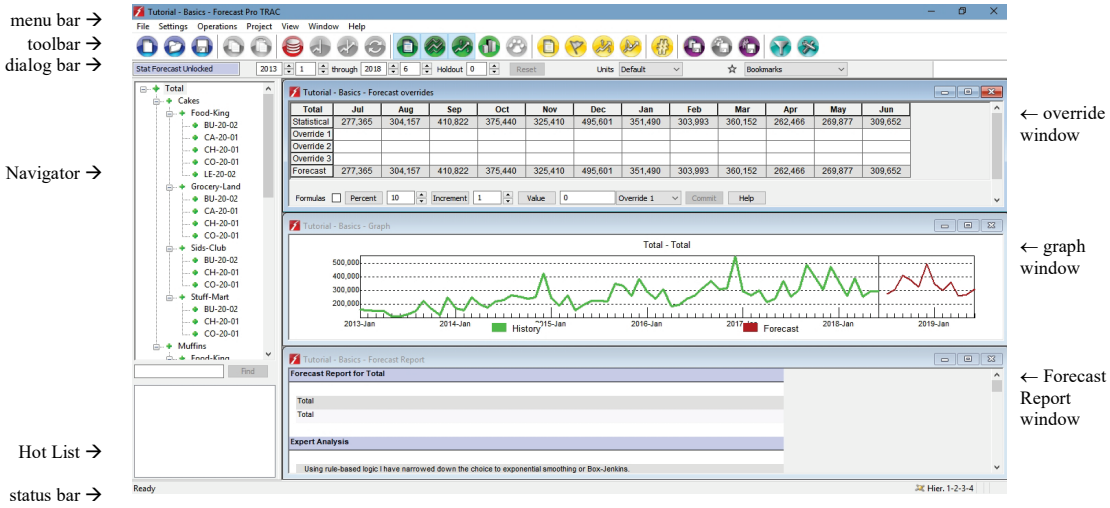

The Forecast Pro window consists of the following seven parts.

Menu bar

Toolbar

Dialog bar

Navigator

Hot List

9 view windows (three are shown above)

Status bar

The Menu bar

The menu provides access to all commands. A complete listing of all menu items and their associated dialog boxes is presented later in this chapter.

The Toolbar

The icons on the toolbar provide a quick access to the most commonly used commands. Here are the icons and their functions.

New Project. Closes the current forecast project allowing you to start a new one.

Open Project. Opens an existing forecast project.

Save Project. Saves the active forecast project.

Copy. Copies the current selection to the Windows clipboard.

Paste. Inserts the contents of the Windows clipboard into the current insertion point.

Data Manager. Accesses the Data Manager allowing you to read in the historic data you wish to forecast, read in any optional data you'd like to use and generate the forecasts.

Read Data. Reads in the Historic data specified in the Data Manager and any optional data with the "on Read" option selected.

Forecast. Generates the statistical forecasts. This option is only available when the historic data have been read and the forecasts need to be manually updated.

Recalculate. If the *Manual override mode* and/or the *Manual forecast mode* is selected on the Performance tab on the Settings dialog box, *Recalculate* will update the forecasts and outstanding override calculations.

View Forecast Report. Opens and closes the Forecast Report view.

View Graph. Opens and closes the Graph view.

Override Forecasts. Opens and closes the Override view, allowing you to adjust the forecasts for the current selection.

Advanced Diagnostics. Opens and closes the Advanced Diagnostics view, allowing you to view the Fitted Error, Autocorrelation Function, Error Autocorrelation Function, Partial Autocorrelation Function or a cumulative graph.

View Tracking Report and Graph. Opens and closes the Tracking Report view and Tracking Report graph.

View Item Report. Opens and closes the Item Report view.

View Exception Report. Opens and closes the Exception Report view.

View Override Report. Opens and closes the Override Report view.

View Outlier Report. Opens and closes the Outlier Report view.

Preview Numeric Output. Opens and closes the Preview Numeric Output view. This view previews the contents and format of the Numeric Output file.

Save Formatted Forecast Report. Allows you to save Formatted Forecast Reports to disk. The Excel file will contain a report for every item forecasted. You can save Formatted Forecast Reports for the items on the current Hot List only using the Hot List's context menu.

Save Forecast Tracking Report. Allows you to save the Tracking Reports to disk. The Excel file will contain a report for every item for which at least one archived forecast is available. You can save Tracking Reports for the items on the current Hot List only using the Hot List's context menu.

Save Numeric Output. Allows you to save a Numeric Output file to disk. The file will contain output for every item forecasted. You can save a Numeric Output file for the items on the current Hot List only using the Hot List's context menu.

Apply Filters. Turns on and off report filtering mode. Filtering mode allows you to filter and sort the item, exception, override and outlier reports as well as the preview numeric output view.

Options. Opens the Options dialog box allowing you to change the settings for the current forecast project.

The Dialog bar

The dialog bar is used to set the span of the data, to define a holdout sample, to change the display units, to manage Bookmarks and to select Bookmarks.

Defining the fit set. All time series methods begin by fitting the coefficients of a model to historic data. You can alter the beginning year and beginning period to specify the first point of the fitting sample and/or the ending year and ending period to specify the last point.

The dialog bar automatically selects the largest fit set for which data exists. Normally, you will want to accept the defaults. However, you may want to ignore earlier data if you distrust their relevance.

Please note that the fit set selected on the Dialog bar is applied to all forecasted items, unless the *\START(Year,period)* modifier is used to specify a different start date for an individual item.

Defining a holdout sample. Edit the holdout sample box to specify a number of data points to be withheld from the end of your data set. If you withhold 2 or more points, Forecast Pro automatically generates out-of-sample evaluation statistics for any models that you build. If you do not want to perform out-of-sample testing, set the holdout sample to zero.

Changing Units. The Units drop down box is used to change the display units. This box will only be available if you have defined conversion factors for your current data set.

Managing Bookmarks. Click on the Bookmark star (☆) to Add, Manage, Rename, Remove or Refresh bookmarks or import Bookmarks from another project. Please see details in the Working with Bookmarks section of the Advanced Navigation tutorial.

Selecting Bookmarks. Click on the Bookmark drop down to view and select bookmarks for the open project.

The Navigator

The Navigator is the primary way to select an item to view in the Forecast Report, Graph, Override and Tracking Report windows. After the data have been read, the Navigator displays the available time series in a tree structure. Selecting an item on the Navigator will automatically display all relevant information in the open views.

The Navigator also features color-coded icons to allow you to spot items that contain overrides and/or comments. A red icon indicates that the item contains an override and/or comment. A green icon indicates that the item does not contain an override or comment. A yellow icon on a group indicates that at least one item further down that branch of the tree contains an override and/or comment. A light blue icon (applicable only if *Manual forecast mode* is selected on the Performance tab of the Options dialog box) indicates that an item's forecast needs to be updated with the Forecast ( ) or Recalculate ( ) buttons on the toolbar. A dark blue icon (applicable only if *Manual override mode* is selected on the Performance tab of the Options dialog box) indicates that an item's override calculations need to be updated with the Recalculate ( ) button.

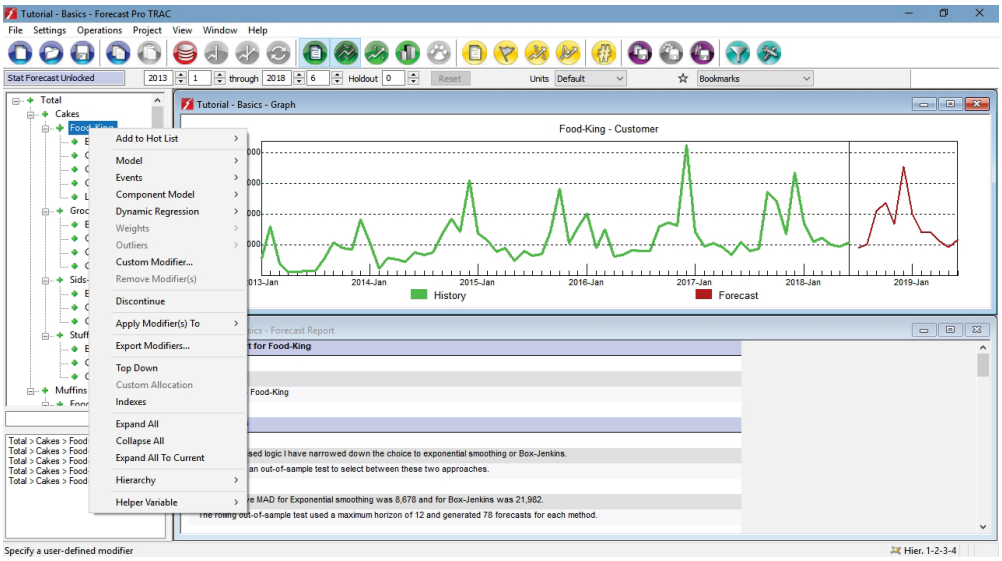

The Navigator's context menu provides a convenient way to build a Hot List, apply forecast modifiers, control the display of the Navigator tree and rearrange the hierarchy. Most of the options are self-explanatory, however a few warrant some explanation.

The options under **Add to Hot List** are used to copy items onto the current Hot List (the Hot List is described in the next section). *Children, Parents* and *Siblings* refer to one level down on the current Navigator branch, one level up on the current branch and the same level on the current branch respectively. On the Analysis submenu, *Overrides* refer to items with direct overrides and

Affected refers to items with indirect overrides (i.e., items where a direct override elsewhere in the hierarchy changed the item's forecast).

The **Model**, **Events**, **Component Model, Dynamic Regression, Weights**, **Outliers**, **Top down, Custom Allocation** and **Indexes** options are all used to specify forecast modifiers. The most commonly used modifiers can be applied directly using the menu options. The less commonly used modifiers can be entered from the keyboard using the **Custom Modifier** option. A list of all supported modifiers and their function is found in the *Using Forecast Modifiers and Model Specification Dialog Boxes* chapter of this manual.

The **Hierarchy** option is used to "shuffle" (rearrange) the hierarchy displayed on the Navigator. Reading data into Forecast Pro establishes the default hierarchy. If the default hierarchy contains 4 levels, the notation 1-2-3-4 is used to represent the default state. **Hierarchy>Promote** and **Hierarchy>Demote** redefine the hierarchy by moving the currently selected level "up" or "down" one level. For example, promoting level 3 of a 4-level default hierarchy, would change the displayed hierarchy from 1-2-3-4 to 1-2-4-3. **Hierarchy>Move to Top** promotes the currently selected level to the highest level, **Hierarchy>Move to Bottom** demotes the currently selected level to the lowest level and **Hierarchy>Restore Default** converts the currently displayed hierarchy back to the default hierarchy.

The **Helper Variable>Create** option is used to access the Create Helper Variable dialog box. This utility allows you to create helper variables based on your current Navigator selection. It is explained in the *Create Helper Variable Dialog Box* section of this chapter.

If you select a helper variable on the Navigator, you can use the **Helper Variable>Delete** option to delete it from your project.

The Hot List

Placing items on the Hot List allows you to efficiently navigate, work with and report on a subset of the items listed on the Navigator. Items can be added to the current Hot List by dragging from the Navigator or by using the **Add to Hot List** option on the Navigator's context menu.

When you select an item on the Hot List, Forecast Pro will immediately select the item on the Navigator and update the affected views.

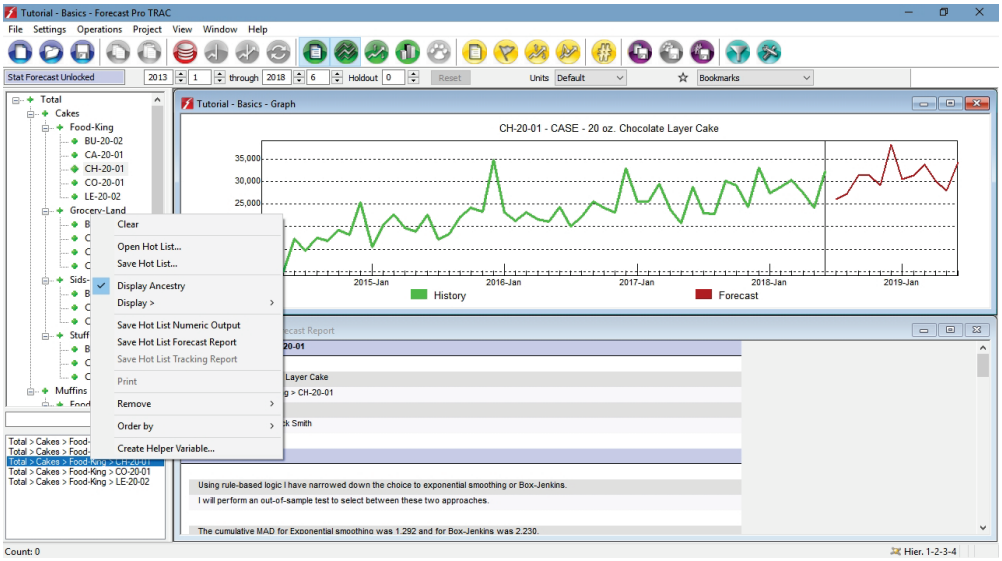

The Hot List's context menu allows you to change how items are displayed, save numeric output files and Formatted Forecast Report files for the current Hot List items and remove items from the current Hot List. The **Open Hot List** option allows you to open a Hot List which had been saved to Excel or import a Hot List from another project. The **Save Hot List** option allows you to export the current Hot List to Excel. The ability to import and export Hot Lists to Excel allows you to create multiple Hot Lists for a project and switch amongst them.

View Windows

There are eleven view windows available in Forecast Pro. Six of them, Forecast Report, Graph, Override Forecasts, Advanced Diagnostics, Tracking Report and Tracking Graph, are context specific—meaning that the contents displayed in these windows will match the current Navigator selection and update automatically as you move about the Navigator. The context-specific views all have green icons.

The remaining five views, Item Report, Exception Report, Override Report, Outlier Report and Preview Numeric Output, are global (i.e., their contents pertain to all items forecasted and are independent of the current Navigator selection). The global views all have yellow icons.

A description of each of the views appears below.

The Status bar

As you use Forecast Pro the status bar displays relevant information.

The Forecast Report View

The Forecast Report view is a scrollable text display containing information about the forecasts and how they were generated. In addition to providing a convenient way to view information while using Forecast Pro, the Forecast Report view is also used to design the text section of the formatted forecast reports that are output to Excel.

The Forecast Report Design dialog box allows you to customize the Forecast Report view. This dialog box is invoked by selecting **Settings>Forecast Report Design** or by selecting **Forecast Report Design** from the Forecast Report view's context menu.

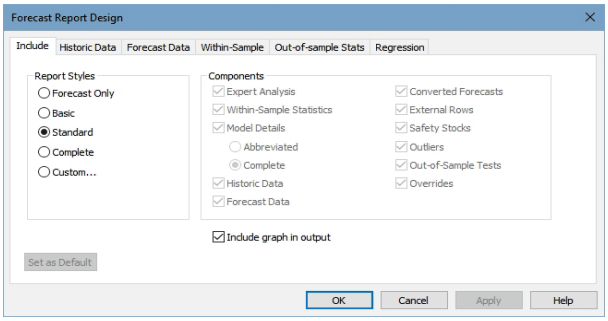

You can either select one of the four standardized report styles or define a custom report. *Include graph in output* can be used in conjunction with any style report and specifies whether to include or omit a graph when saving a Formatted Forecast Report to Excel. The settings for the saved graph may be specified in the Graph view, as described below. A description of each report style appears below.

Forecast Only is a minimal display consisting solely of the forecasts.

Basic displays a listing of the forecasting model used, a set of the most commonly used within-sample statistics and a forecast display including confidence limits and summary statistics.

Standard is the default display. It includes the expert selection logic (if applicable), the model details including model coefficients, a full set of within-sample statistics, a safety stock display, a listing of detected/corrected outliers (if outlier detection is active), a listing of any overrides and the same forecast display found in the Basic style.

Complete displays everything found in the Standard display as well as a numeric listing of the historic values, fitted values, converted forecasts and an expanded forecast display.

Custom allows you complete control over what is included in the display. The *Historic Data*, *Forecast Data* and *Regression* tabs on the Forecast Design may be used to specify details about the components to include when Custom is selected. Your selections under the Custom option are automatically retained for the project allowing you to switch between the Custom style and one or more of the standard styles without having to reset your custom settings.

Important note: The *Within-sample* and *Out-of-sample Stats* tabs allow you to set the numeric precision for the within-sample and out-of-sample statistics. You also have the ability to set precision for regression model details on the *Regression* tab. These precision settings are used for both Custom and standard report styles.

The Graph View

The graph view is used to display variables and forecasts graphically. In addition to providing a convenient way to view the forecasts while using Forecast Pro, the graph can also be included in the formatted forecast reports that you save to Excel. The format and content of the formatted forecast report's graph will match the current settings for the graph view.

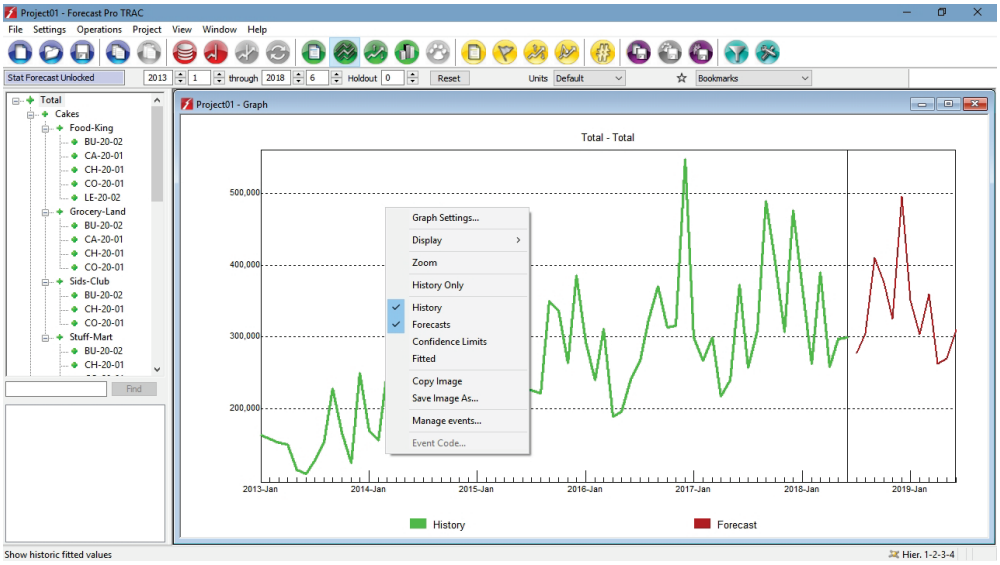

The Graph Settings dialog box allows you to customize the Graph view. This dialog box is invoked by selecting **Settings>Graph Settings** or by selecting **Graph Settings** from the Graph view's context menu.

The Settings dialog box contains three pages or tabs. We will discuss each in turn.

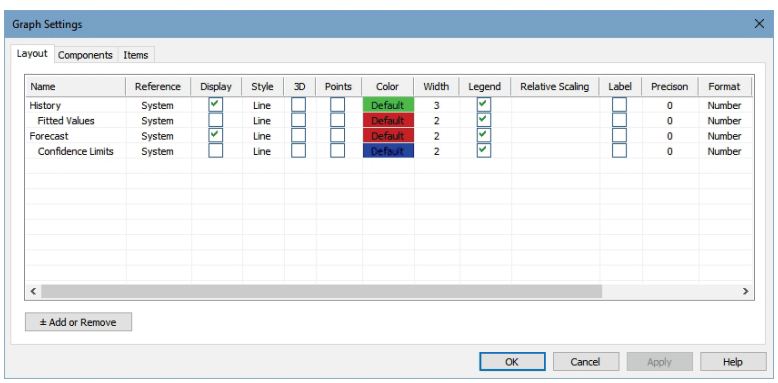

The Layout tab allows you to the select which components to display and format those components as you choose. The History, Fitted Values, Forecast and Confidence Limits components are always available in the Layout tab. When you first start a project, the default graph displays the History and Forecast components.

Other graph components may be added to or removed from the Layout tab using the Add or Remove button. Clicking the Add or Remove button brings up the Add or Remove dialog box shown below. The checkboxes are used to select which components should be included in the Layout tab. The Filter drop down is useful for selecting components by variable type.

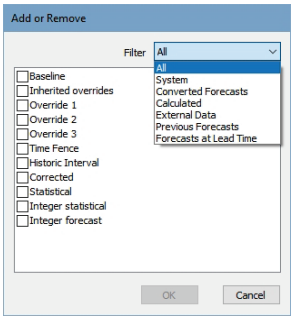

Clicking OK exits the Add or Remove dialog box and updates the list of components in the Layout tab.

You can set specific formatting options for each component.

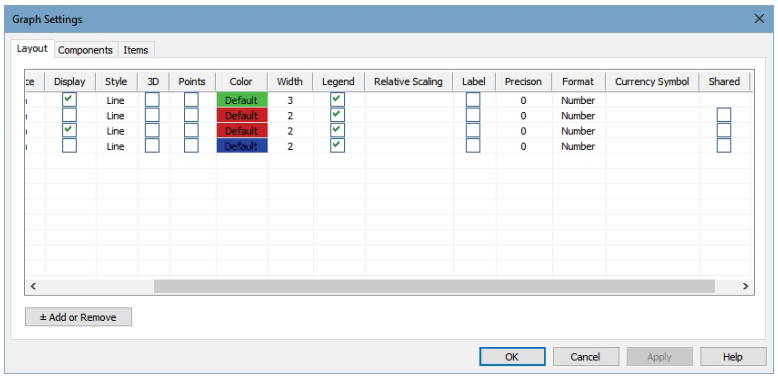

Most of these options are self-explanatory. The Style drop down provides four options: Scatter, Line, Bar or Area. Relative Scaling is only available for Converted Forecasts, Calculated Rows and External Data. Selecting Relative Scaling will display that component on its own scale. This is useful when graphing multiple components with dramatically different values. The Shared option allows you to coordinate color and numeric displays with the Override

view. Checking Shared applies your Graph Setting selections to the Override grid, providing synchronized views.

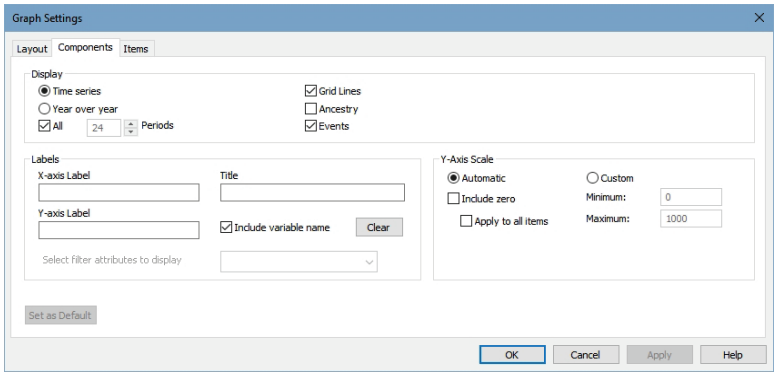

The Components tab includes three sections.

The ***Display*** section allows you to select the graph type.

Time series displays a graph where the y-axis covers both the historic and forecast period. The *All* option will display the entire data set in a nonscrollable display. If the All option is not selected, the *Periods* option allows you to set how many periods should be included in a scrollable display.

Year over year displays a graph where the y-axis is one year long and the data for each year are "stacked" on the display. The *Years* option allows you to specify the number of years to include in the display.

The *Grid Lines* option displays grid lines on the graph. The *Ancestry* option displays the complete ancestry in the variable name (the variable name will appear preceded by any applicable parent group names). The *Events* option marks all historic or forecast points where an event code is being applied. The *Outliers* option marks all historic points where an outlier was detected and/or corrected.

Most of the Display options can also be controlled by bringing up the Display sub-menu on the Graph context menu. Selecting *Zoom* on the Graph context menu will activate the scrollable display.

The ***Labels*** section allows you to define labels for the X and Y axis and add titles to the graph. *Include variable name* will use the variable name followed

by the description as the graph's subtitle. The Clear button erases the current labels.

The ***Y-Axis Scale*** section controls the scaling for the selected item. Please note that these are item-level settings.

Automatic allows Forecast Pro to select the Y-axis scale.

Custom lets you set the minimum and maximum for the Y-axis scale.

Include zero begins the y-axis at zero or at the minimum negative value. If you select *Apply to all items*, the Include zero option will be applied globally. This is the only option in the section that may be applied globally.

The Labels and Y-axis Scale settings are shared by the Year over Year and Time Series graphs. If you prefer to use different settings for the different graph types, Bookmarks are a good way to do so.

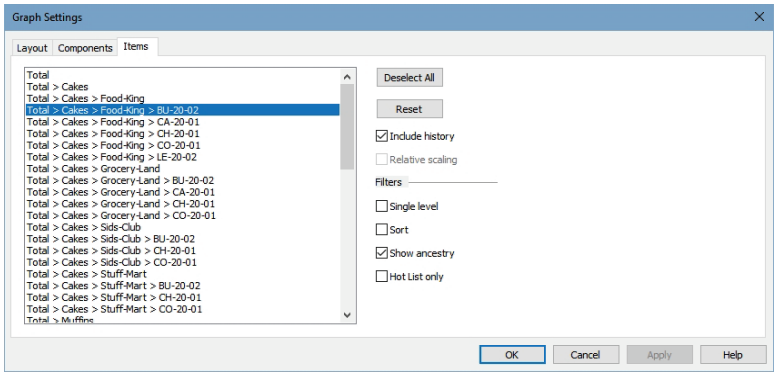

The Items tab allows you to choose up to five variables you want to graph. Simply select the items you want to display and then click the OK or Apply button. *Include history* indicates if the graph should include both the historic values and forecasts. *Relative scaling* indicates if each graphed item should be displayed on its own scale. The *Filter* options can be used to limit which items are displayed for selection.

The Override Forecasts View

The override view is used for several different purposes. It can be used (1) to enter overrides to the statistically generated forecasts (or specified baseline forecasts), (2) to display imported data and calculated rows and (3) to specify the baseline forecasts. We will examine each of these operations in turn.

Working with overrides

The override view is used to enter overrides to the statistically generated forecasts. Changes made at any level of the forecasting hierarchy will automatically reconcile all levels. Refer to the *Understanding How Forecast Overrides are Allocated* chapter for a discussion on how the reconciliation is accomplished.

Up to ten override rows can be used and you can adjust the forecasts using percentages, increments, formulas or by typing in new values. You can also enter comments for any cell on the display.

There is an Overrides tab on the **Settings>Options** dialog box that allows you to control the document-level options, including how overrides should be applied (replacement values vs. incremental adjustments) and the precision of the display. This tab is fully documented in the *Menu Items and Dialog Boxes* section of this Command Reference chapter.

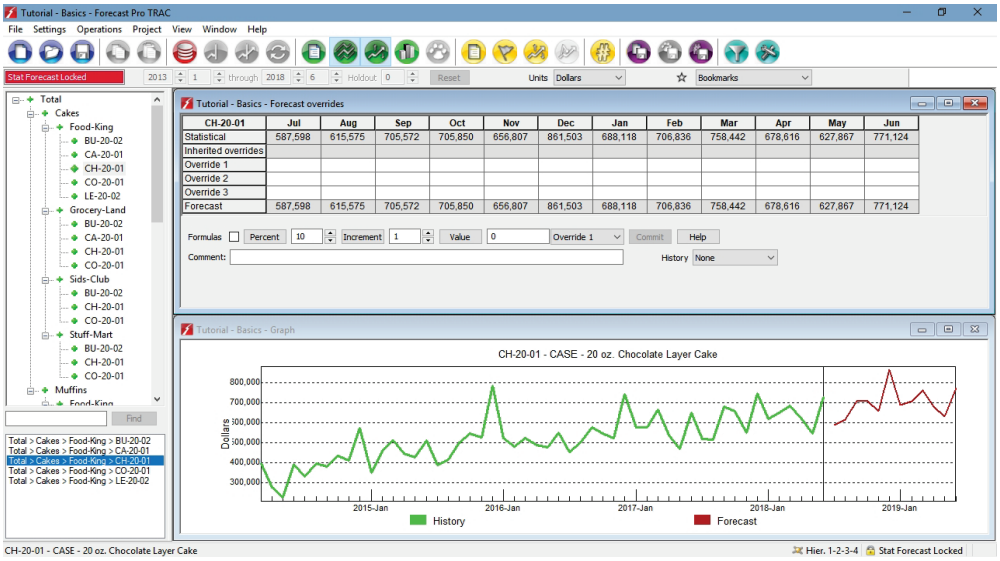

There are six rows in the example shown above.

The top row (currently labeled *Statistical*) shows the Baseline forecast. By default, the Baseline forecast is the Statistical forecast (the statistically based forecasts generated by Forecast Pro) and is labeled as such. If the baseline forecast is changed, the row label will reflect that.

The *Inherited Overrides* row is only present if your project has changed states (i.e., you have either switched units of measure or have rearranged the hierarchy). When you change states it is not always possible to maintain the in-place overrides. This is due to factors such as the item or group that had an in-place override may no longer exist after you change the hierarchy or that changing the units would result in a different forecast allocation across the hierarchy.

To avoid these problems and to insure that the final forecasts will remain consistent regardless of the hierarchy or units you are viewing, Forecast Pro will move the in-place overrides into the Inherited Overrides row prior to performing hierarchy changes and/or conversions.

Override 1, Override 2 and Override 3 contain any overrides that you have entered for the current forecasts. Proposed overrides are displayed in red, committed overrides are displayed in black. The Commit button is used to convert proposed overrides to committed overrides. The *Formula* checkbox

allows you to enter overrides as either formulas or “hard numbers” when using the increment and percent buttons. If the Formula box is checked, using the percent or increment buttons will save a formula. Formulas may be entered directly into the Override grid rows whether the Formula option is selected or not. The override view’s context menu allows you to rename the override row labels and add or subtract rows on the fly.

Forecast contains the current committed forecasts. These are the values that are displayed on the graph, and the values that will be written as “forecasts” in any output files that you save. The Forecast values may differ from the Statistical values if you have committed to overrides for the currently displayed forecasts *or any other forecasts in the hierarchy that would impact the currently displayed forecasts.*

There are three steps in the override process—selecting the point or points to override, entering the override(s) and committing to the override(s).

Points can be selected in the override window or in the graph window. The sum of the values for the points selected will be displayed in the *Value* edit box.

If you want to adjust the selected points by a certain percentage, enter the percentage in the *Percent* edit box and press the button. If you want to adjust each selected point by a certain increment, enter the increment in the *Increment* edit box and press the button. If you want to change the sum of the selected points to a defined value, enter the value in the *Value* edit box and press the button. Adjustments made using the edit boxes will appear in the target override row specified to the left of the Commit button.

You can also enter overrides for a single point by typing a new value into one of the override row cells or by using the right mouse button to drag it to the desired new value on the graph.

After you have entered the desired overrides and comments, click the Commit button to accept the changes and reconcile the hierarchy.

For large, complex hierarchies the reconciliation can take a little time. To minimize the reconciliation time, Forecast Pro supports a manual override mode. When this mode is active, the override window will include a Retain button as well as a Commit Button. The Retain button is used to accept

overrides without reconciling the hierarchy. In this mode you would typically enter and retain all of the overrides you wish to make (i.e., for multiple items) and then click Commit to perform the reconciliation. The control to turn on manual override mode is found on the Performance tab of the **Settings>Options** dialog box.

Row types

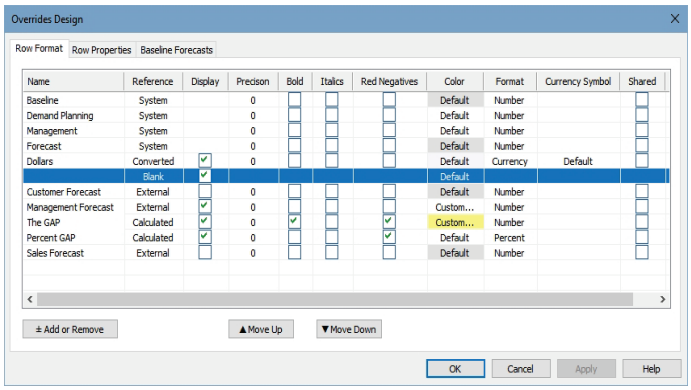

The override grid lets you view and work with different types of rows. The screenshot above displays the Overrides Design dialog box which lists the currently available rows and lets you customize their appearance. Select Overrides Design on the Override view context menu or select **Settings>Override Design** on the menu bar to bring up this dialog box. The current screenshot is taken from the project we created in Lesson 3 and includes a range of row types.

History rows, the Baseline row, Override rows and the Forecast row are all standardized rows which always appear in a specific order at the top of the grid. These rows are always displayed on the grid and on the Row Format tab of the Override Design dialog. The Row Format tab on the Override Design dialog is also where you control the formatting. Let's consider each in turn.

History rows. If you select *Year over Year* on the History drop down, the grid will display the demand history, by year, prior to the Baseline row. If you select *Time Series*, the grid with display both History and Forecast in the current Forecast row and relabel it History/Forecast.

Baseline row. The statistical forecast generated by Forecast Pro is used as the default baseline forecast. (Specifying alternative baseline forecasts is discussed later in this section.) Think of the baseline forecast as the starting point in establishing the final forecast displayed in the Forecast row. If no overrides are made, the final forecast displayed in the Forecast row will equal the Baseline forecast. If you enter overrides, the baseline is adjusted accordingly to establish the final forecast.

Override rows. Operation of the override rows was discussed in detail in the preceding section.

Forecast row. The Forecast Row displays the current forecast. If *Time Series* is selected on the History drop down, this row will be labeled History/Forecast.

In addition to the standardized row types above, Forecast Pro allows you to add Converted rows, External rows and System rows to the grid. To add these rows, click on the Add or Remove button on the Row Format tab on the Override Design dialog box to bring up the dialog box shown below.

You can use the filter drop down to select rows by type or display all rows that may be added to the grid.

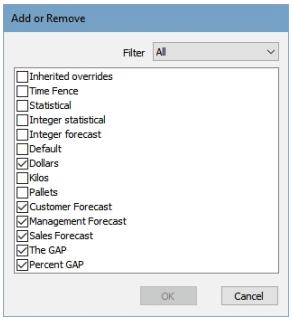

Converted rows. Converted rows display the current forecast in different units of measure. If you read in conversion factors in the Data Manager, corresponding converted rows will be available.

External rows. External rows contain numeric information that you have imported using the Data Manager. Common examples of imported external rows are forecasts generated outside of Forecast Pro such as sales forecasts, or

management forecasts and product information such as current orders or available inventory.

System rows. The System rows include Inherited Overrides, Time Fence, Statistical and Integer forecasts.

Inherited Overrides are described in detail in the beginning of this section.

Time Fences are described in the *Using a Time Fence* section of this Command Reference chapter.

Statistical contains the statistically based forecast generated by Forecast Pro prior to any overrides being applied.

The Integer forecast displays either the final or baseline forecasts using a special form of rounding that keeps track of the "remainder" (rounded amount) and applies it to the next period's forecast value. This is particularly useful for low-volume forecasts or when there are large minimum order quantities, where traditional rounding can result in biased forecasts.

Consider a flat forecast of 0.4 units per month for the next 12 months. If you were to use the ROUND() function to round the forecast to the nearest integer the forecast would become, "0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0". An integer forecast row for the same forecast would be, "0, 1, 0, 1, 0, 0, 1, 0, 1, 0, 0, 1". Integer forecast settings are defined in the Advanced Control tab of the Options dialog box. Details are described in the *Menu Items and Dialog Boxes* section of this Command Reference chapter.

Previous Forecasts. Previous forecasts, a specific type of System row that you can filter on in the Add or Remove dialog box, are simply forecasts that were generated in prior forecast cycles (i.e., that you created in the past and are stored in the Forecast Pro data base). Previous forecasts are useful for tracking how forecasts change from one forecast cycle to the next and measuring forecast accuracy. Forecast Pro archives final forecasts, statistical forecast and baselined forecasts. These archived forecasts are named "Previous Forecast #", "Previous Statistical #" and "Previous Baseline #", where '#' indicates when the forecast was archived. For example, "Previous Forecast 1" displays the final forecasts generated one period ago, while "Previous Statistical 2" displays the statistical forecasts generated 2 months ago. Forecast Pro will

display all archived periods available in the database in the Add or Remove dialog box.

Forecasts at Lead Time. Also a type of System Row, Forecasts at lead time are useful for tracking forecast changes and accuracy for specific lead times. A lead time refers to the number of periods ahead of the forecast origin (the period when the forecast was established) the forecast was made for. Thus, a one-month-ahead forecast would have lead time equals 1, a two-month-ahead forecast would have lead time equals 2, etc.

You can also create Calculated rows and Blank rows in the overrides grid.

Calculated rows. Calculated rows allow you to enter formulas using other rows, arithmetic operators and predefined "tokens" (key words). They are inserted into the grid by selecting the Forecast row (or any row underneath it) right clicking to invoke the context menu and selecting **Add row**. A dialog box will appear allowing you to name the calculated row. Once added, a calculated row can be added or removed from the Row Format tab on the Overrides Design dialog box using the Add or Remove button, as described above. The next section discusses how to enter the formulas and details the supported operators and tokens.

Blank rows. Blank rows can be added to the grid to enhance readability. They are inserted in a similar fashion to calculated rows. To insert a blank row, select the Forecast row (or any row underneath it) right click to invoke the context menu and select **Add row**. A dialog box will appear allowing you to name the row. If you enter a name and click OK the row will be inserted as a calculated row. If you erase the default name and click OK (i.e., process a blank input field) the row will be inserted as a blank row.

Defining calculated rows

You enter formulas by clicking on the cell you wish to modify and typing in the formula. If you wish to edit a previously entered formula, select the cell and press the F2 key. You can copy and paste formulas from cell to cell.

By default, the formulas in a specific calculated row will be calculated at the single level of the hierarchy on which they were entered. Any values you elect to display at other levels will be either aggregations or disaggregations of the

calculated values. *Depending upon your formulas, this may or may not be equivalent to applying the defined formula at the group level.*

If you wish to change the default behavior for a row and calculate the formulas at all levels you can do so using the Row Properties tab (discussed below). Be aware that this option may slow down certain operations including applying overrides.

Formulas are not case sensitive. Formulas are entered in the format:

= expression

The expression can involve any of the following components.

Numbers	Any real number
Operations	+ - * / ^
Parenthesis	()
Comparison operators	= < > <= >=
Functions	Listed below
Tokens	Listed below
Rows	{RowName}
Conversion factors	{#ConversionFactorName}

Note that when a row name is used as part of a definition it must appear in braces (curly brackets).

The conversion factors are a special token that returns the value of the specified conversion factor. The token must appear in braces and have the # symbol prior to the conversion factor name. For example if *Dollars* is a defined conversion factor, then *{#Dollars}* would be the corresponding token.

Token and row references can also include an offset by appending a [-n] to the token or row name. For example:

{rowname}	Returns the value of rowname for the current period.
{rowname}[-1]	Returns the value of rowname for the previous period.
{rowname}[1]	Returns the value of rowname for the next period.

Forecast Pro supports the following functions:

ABS(*n*)	Returns the absolute value of *n*
EXP(*n*)	Returns *e* raised to the nth
LOG(*n*)	Returns the natural logarithm of *n*
LOG10(*n*)	Returns the base-10 logarithm of *n*
SIGN(*n*)	Determines the sign of *n*. Returns 1 if *n* is positive, 0 if *n* is zero and -1 if *n* is negative.
SQR(*n*)	Returns *n* squared
SQRT(*n*)	Returns the square root of *n*
IF(*condition, value1, value2*)	Returns *value1* if *condition* is true or *value2* if *condition* is false
INT(*n*)	Rounds *n* down to the nearest integer
TRUNC(*n*)	Truncates *n* to an integer
EVEN(*n*)	Rounds *n* up to the nearest even integer
ODD(*n*)	Rounds *n* up to the nearest odd integer
ROUND(*n, decimal places*)	Rounds *n* to the specified number of *decimal places*
MROUND(*n, basis*)	Rounds *n* to the nearest multiple of *basis*
ROUNDUP(*n, decimal places*)	Rounds *n* up to the specified number of *decimal places*
ROUNDDOWN(*n, decimal places*)	Rounds *n* down to the specified number of *decimal places*
MROUND(*n, basis*)	Rounds *n* to the nearest multiple of *basis*
MIN(*x, y, ...*)	Returns the smallest value in the specified set of values
MAX(*x, y, ...*)	Returns the largest value in the specified set of values
AVERAGE(*x, y, ...*)	Returns the average of the specified set of values
SUM(*x, y, ...*)	Returns the sum of the specified set of values
MEDIAN(*x, y, ...*)	Returns the median of the specified set of values
MODE(*n*)	Returns the mode (most frequently occurring value) of the specified set of values

NVL(*n, alternative*)	Null value function. If *n* is blank returns *alternative*. If *n* is not blank returns *n*.
ISNUMBER(*n*)	Returns 1 if *n* is a number, 0 if *n* is blank.
ISNULL(*n*)	Returns 1 if is blank, 0 if *n* is a number.

If you use Excel, you may be familiar with these functions. Forecast Pro implements these functions in the same way as Excel with one exception—cell ranges are not yet supported. For example, the following formula will return the sum of the statistical forecast for last period, this period and next period:

=SUM(STAT[-1],STAT,STAT[1])

The following, however, is not supported: =SUM(STAT[-1]:STAT[1])

Forecast Pro supports the following tokens:

FORECAST	Forecast row
PFC#	Forecast from '#' periods prior
PST#	Statistical forecast from '#' periods prior
PBL#	Baseline forecast from '#' periods prior
FCLT#	Forecast at lead time '#'
STLT#	Statistical forecast at lead time '#'
BLLT#	Baseline forecast at lead time '#'
YTD	Returns the calendar year-to-date using the history and forecast values as applicable.
QTD	Returns the calendar quarter-to-date using the history and forecast values as applicable.
STAT	Statistical forecast row
BASE	Baseline forecast row
INTFC	Integer Forecast. Only available if the Integer Forecast is added to the grid in the Override Design dialog.
INTSTATFC	Integer Statistical Forecast. Only available if the Integer Forecast is added to the grid in the Override Design dialog.
OVR1	The first override row

OVR#	Override row “#”
HIST1	The most recent historic row
HIST#	The historic row “#” years prior
HISTBL[-n]	Returns either the history value or the baseline forecast value at -n. Intended for use in formulas where it may point to either a forecast or a historic period.
HISTFC[-n]	Returns either the history value or the forecast value at -n. Intended for use in formulas where it may point to either a forecast or a historic period.
HISTST[-n]	Returns either the history value or the statistical forecast value at -n. Intended for use in formulas where it may point to either a forecast or a historic period.
INHER	Inherited override row.
LOWER	Lower confidence limit
UPPER	Upper confidence limit
LEADTIME	The item’s specified lead time
SAFETY	The safety stock at the item’s specified lead time
SAFETY_ALL	The safety stock corresponding to the lead time of the current cell
DDLT	The demand during the item’s specified lead time
DDLT_ALL	The demand corresponding to the lead time of the current cell
REORDER	The reorder point at the item’s specified lead time
REORDER_ALL	The reorder point corresponding to the lead time of the current cell
MEAN	The mean of the historic data
STDEV	The standard deviation of the historic data
MAD	The within-sample mean absolute deviation
MAPE	The within-sample mean average percent error

SMAPE	The within-sample symmetric mean average percent error
BIC	The Bayesian Information Criterion
{*#Conversion*}	Returns the conversion factor for "Conversion".

Important notes:

10. All of the above functions require that the referenced system row is added in the override grid. For example, the integer system row needs to be in the grid for INTFC to return a value. Referenced system rows can be added using the Add or Remove button in override settings.
11. Tokens that reference baseline or final forecasts are only available in calculated rows.

Formatting rows

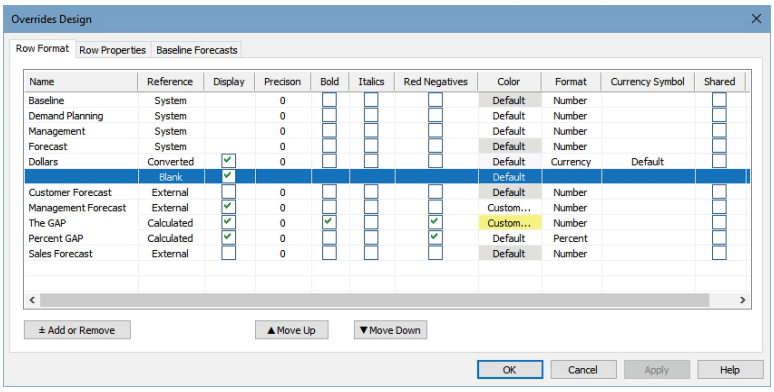

Name	Reference	Display	Precison	Bold	Italics	Red Negatives	Color	Format	Currency Symbol	Shared
Baseline	System		0				Default	Number		
Demand Planning	System		0				Default	Number		
Management	System		0				Default	Number		
Forecast	System		0				Default	Number		
Dollars	Converted	✓	0				Default	Currency	Default	
	Blank	✓					Default			
Customer Forecast	External		0				Default	Number		
Management Forecast	External	✓	0				Custom...	Number		
The GAP	Calculated	✓	0	✓		✓	Custom...	Number		
Percent GAP	Calculated	✓	0			✓	Default	Percent		
Sales Forecast	External		0				Default	Number		

The Row Format tab on the Overrides Design dialog box allows you to control the order in which the rows appear and how they are formatted. Most of the options are self-explanatory, however, we will comment on a few of them.

Only items that have Display checked will be shown on the grid. The core grid rows (Baseline, Overrides, Time Fence and Forecast) are always displayed and do not have checkboxes. If you want to add something to the grid and it is not shown on the Row Format tab, add it by using the Add or Remove button, as described above.

Checking Shared applies your Override settings for that row to the Graph view, allowing for synchronized views.

To change the order of the rows, you select a row you wish to move by clicking on its name and then use the Move Up and Move Down buttons to reposition it.

Clicking the Apply button will immediately apply any changes you have made to the Override View without leaving the Overrides Design dialog box. This is very useful when formatting the display.

The Format option allows you to display the row's values formatted as a number, a percentage, or a currency. If you select Currency, Forecast Pro will default to the currency symbol associated with your operating system's language selection. If you wish to use an alternative symbol, the Currency Symbol drop down can be used to select it. Forecast Pro supports a wide variety of currency symbols, however, if the symbol you'd like to use is not supported; please send an enhancement request to support@forecastpro.com.

Defining row properties

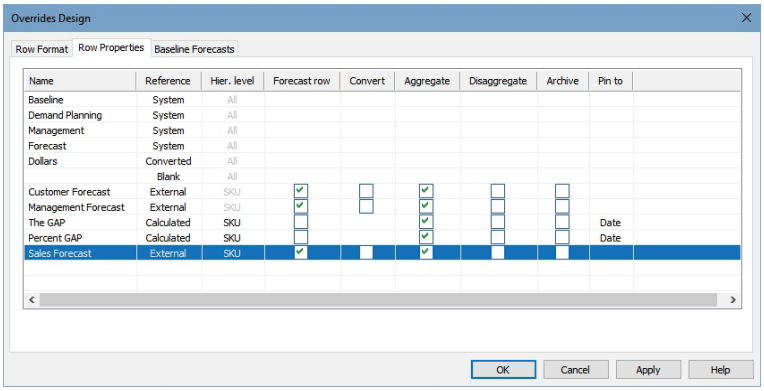

The Row Properties tab on the Overrides Design dialog box allows you to control how certain grid values are calculated and displayed.

The *Hier. level* column displays the level at which the data are defined. For Converted rows this will always be the lowest level since the conversion factors can only be defined at the lowest level. For External rows this column displays the level at which the data were imported. For Calculated rows this

column is a drop down which allows to specify the level to calculate the formulas or to specify that the formulas should be calculated independently at all levels.

The *Forecast row* column allows you to indicate whether the row should be treated as a forecast row. This is important, since only forecast rows can be assigned as baseline forecasts.

The *Convert* column allows you to indicate whether the row should convert to different units when you change the Units displayed in Forecast Pro.

The *Aggregate* column allows you to indicate whether External rows and Calculated rows should display aggregated group totals for levels above the *Hier. Level*.

The *Disaggregate* column allows you to indicate whether External rows and Calculated rows should display disaggregated values for levels below the *Hier. Level*.

The *Archive* column is used to indicate whether the row should be archived in the Forecast Pro database for every forecast origin. This is only recommended if you plan to access the Forecast Pro database independently of Forecast Pro and have a specific need for this information.

The *Pin to* column is used to control where a formula will appear after you update your historic data and generate a new forecast from a different forecast origin. If you select "Date" the formula will be tied to the specific date for which it was originally entered. If you select "Column" the formula will be tied to the column position for which it was originally entered.

Let's illustrate this with an example. Let's say you originally generate a 12-month forecast and the first forecast period is January 2018, so the grid displays 12 columns for the forecasts covering the span of January 2018 through December 2018. You create a calculated row and enter a formula in the 12th column (which corresponds to December 2018). Next month, you update your historic demand, reforecast the project and the grid now displays 12 new forecasts covering the span from February 2018 through January 2019. Where does the formula you entered last month appear in the updated override grid? If you selected Pin to Date, the formula will appear in the 11th column of the display (which now corresponds to December 2018) and the

12th column (which now corresponds to January 2019) will be blank. If you selected Pin to Column, the formula will appear in the 12th column of the display (which now corresponds to January 2019).

Specifying the baseline forecast

By default, Forecast Pro generates a statistical forecast and then allows you to add overrides to this "baseline" forecast to create the final forecast. There may be times where you wish to use an alternative forecast generated outside of Forecast Pro as the baseline forecast rather than the statistical forecast generated by Forecast Pro.

You can set the baseline forecast to use at the end-item level (i.e., the lowest-level on the Navigator) on an item-by-item basis. The baseline forecasts displayed at group levels are always aggregations of the end-item values. If all of the end-items for a given group are using the statistical forecast as the baseline (the default), the group label for the baseline row will read "Statistical." If any end-items for a given group are set to a baseline other than the statistical forecast the group label for the baseline row will read "Baseline."

You can specify the baseline for an individual end-item using the override grid's context menu. Select the forecast row you'd like to use by clicking the row label, invoke the context menu by right clicking the mouse and select **Set as Baseline**.

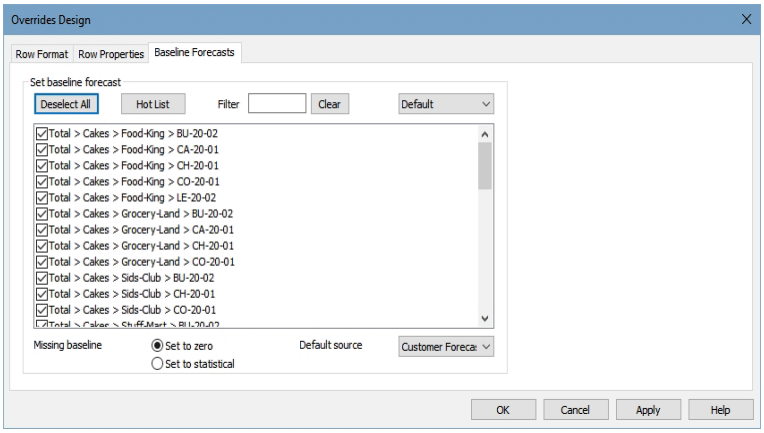

You can use the Baseline Forecasts tab of the Overrides Design dialog box to specify baseline forecasts for groups of end-items.

The baseline selection drop down box (selected as Default in the above image) allows you to select a forecast row to work with. The tab contains a list of all end items with a check mark to indicate whether the baseline matches the current baseline selected. You use the checkboxes and the Apply button to assign the baselines.

The Filter field can also be used to winnow down the displayed list of items. Combining filtering with the use of the Select All button provides an efficient mechanism for assigning the baselines when working with large numbers of end-items.

The Missing baseline setting controls what is displayed if the assigned baseline is partially missing (i.e., the assigned baseline contains values for some but not all of the forecast periods for a given item) or fully missing (i.e., the assigned baseline does not contain any values for a given item). In the latter case, Forecast Pro will display a warning message when you process the request.

The Default source drop down (on the bottom right of the dialog box) allows you to specify the forecast to use as a baseline if no alternative baseline has been specified. Most users will want to leave this selection set to Statistical.

The Advanced Diagnostics View

The advanced diagnostics view provides a variety of statistical displays to assist you in building and diagnosing custom models. To open the advanced diagnostics view click the green Advanced Diagnostics icon ( ).

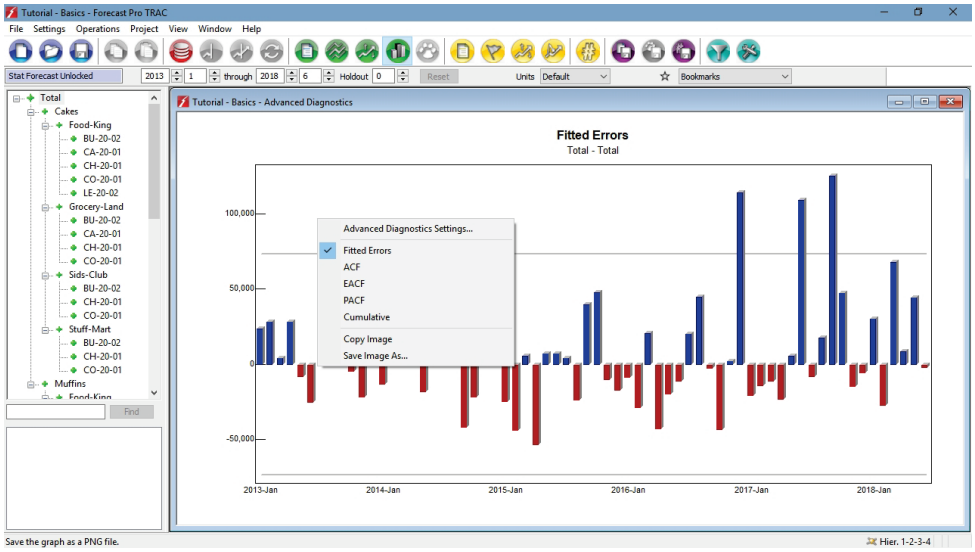

The default view is the *Fitted Errors* graph. Use this display to identify outliers or other problem areas of your data. The dashed horizontal lines mark the 2.5 and 97.5 percentiles. About one error in twenty extends beyond these limits.

You can right click on the *Fitted Errors* graph to bring up the Advanced Diagnostics context menu and select one of the other available diagnostic graphs.

Error ACF displays a graph of the error autocorrelation function. Use the error ACF to make quick assessments of correlational models like Box-Jenkins and dynamic regression. If your statistical model is well specified, then about one autocorrelation in twenty falls outside the limits marked by dashed lines, and the autocorrelations are not patterned. The computation of the error autocorrelation function is covered in the *Forecast Pro Statistical Reference Manual*.

ACF and *PACF* are provided to help you manually identify ARIMA models. The autocorrelation function (ACF) is used to identify differencing, and to identify the number of simple and seasonal MA terms. To identify differencing, select *Advanced Diagnostics Settings* on the context menu.

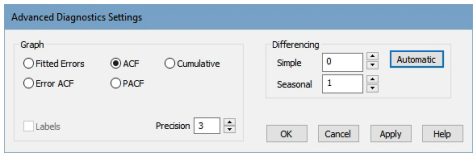

Use the Automatic button to display the differencing values selected by Forecast Pro's automatic Box-Jenkins algorithm. Use the partial autocorrelation functions (PACF) to identify the number of simple and seasonal AR terms. Refer to the *Forecast Pro Statistical Reference Manual* for further information about manual identification of ARIMA models. The *ACF* and *PACF* may also be helpful for manually identifying the dynamic terms in a dynamic regression model.

Cumulative displays a cumulative graph of the currently selected item. This is especially useful when examining a Bass Diffusion model.

The Tracking Report and Tracking Report Graph Views

The Tracking Report view and its associated Tracking Report Graph view allow you to compare previously generated forecasts with what actually happened. To view the Tracking Report, your Forecast Pro database must contain archived forecasts.

To open the Tracking Report and Tracking Report Graph, click the green Tracking Report and Graph icon (). Alternatively, you can open the individual views using **View>Tracking Report** and **View>Tracking Report Graph**. Below is the Tracking Report.

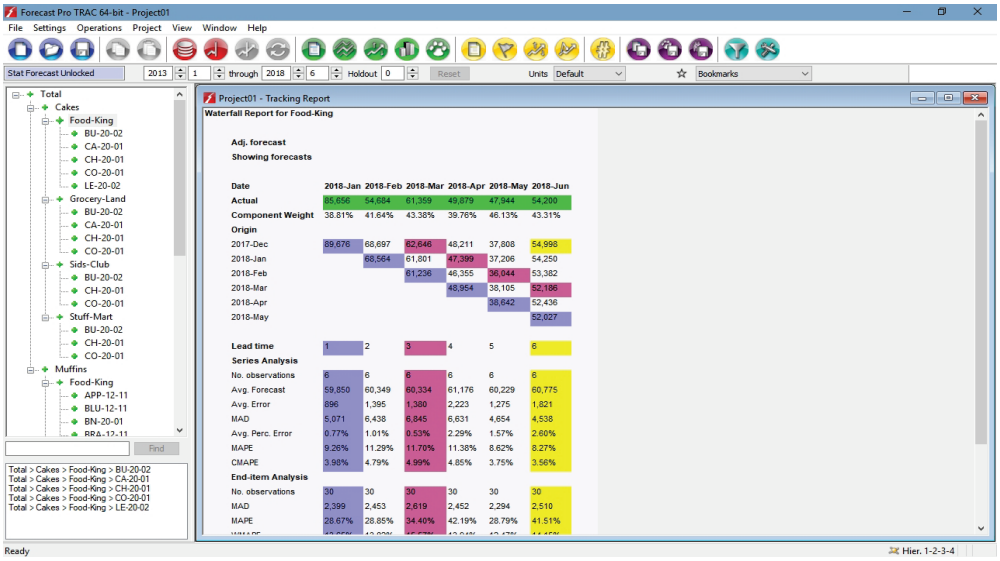

Due to its cascading-like appearance, the tracking report is sometimes referred to as a *waterfall* report. The report compares what we forecasted to what actually happened, therefore it is based on two key elements—the actual demand history and archived forecasts for the periods being analyzed.

In the example above, the actual demand history for the most current six months is shown in the first row with the green shading.

The next row down, displays the *Component Weight*. This displays the percentage that the demand represents of its parent group. Thus in our example, in January of 2018 *Total>Cakes>Food-King* composed 38.81 percent of the demand for its parent group, *Total> Cakes.* This row is only displayed if *Include component weights* is selected in the Layout section of the Tracking Report Settings, as described shortly.

The next 12 rows display the forecasts generated for these periods from different forecast origins. Thus, the row labeled 2017-Dec displays the forecast when the forecast origin was December 2017 and the first forecast period was January 2018. The row labeled 2018-May displays the forecast generated last month when the forecast origin was May 2018 and the first forecast period was June 2018.

The waterfall report allows you to color code *lead* times. A lead time refers to the number of periods ahead of the forecast origin the forecast was made for.

Thus, a one-month-ahead forecast would have lead time equals 1, a two-month-ahead forecast would have lead time equals 2, etc.

In our example, the forecasts for lead time equals 1 are all shaded in blue, the forecasts for lead time equals 3 are all shaded in magenta, the forecasts for lead time equals 6 are all shaded in yellow. The color coding is used on both the tracking report and the tracking report graph. You can control the lead times to color code using the Tracking Report Settings dialog box.

The bottom portion of the report displays cumulative statistics for different lead times.

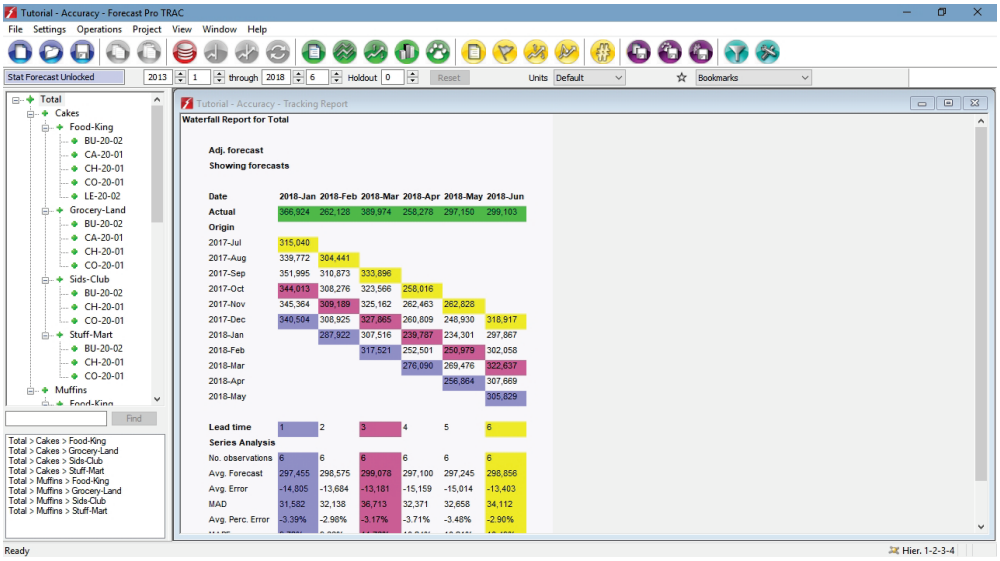

The cumulative statistics for different lead times portion of the report can include up to three sections—Series Analysis, End-item Analysis and Component Analysis. End-item Analysis and Component Analysis are only relevant to group-level data and are not displayed if an end-item is selected. By default, only Series Analysis is shown. Use the Tracking Report settings to turn on End-item and Component Analysis, as described below.

Series Analysis displays statistics for the currently selected time series. If the selection is a group, the statistics are based on the archived group-level forecasts. The CMAPE statistic found in this section is the Component MAPE. This is a weighted MAPE where the weighting factors are the component weights (see discussion above).

End-item Analysis displays statistics based on the archived forecasts for all end-items belonging to the currently selected group. The WMAPE statistic found in this section is a weighted MAPE where the weighting factors reflect the end item's volume. The weights are calculated by dividing each end item's volume by the group's volume.

Component Analysis displays statistics based on the archived forecasts for all component children (i.e., groups and/or end items one level lower in the hierarchy that are components of the currently selected group). The WMAPE statistic found in this section is a weighted MAPE where the weighting factors reflect the component children's volume. The weights are calculated by dividing each component child's volume by the group's volume.

The content and format of the Tracking Report is controlled using the Tracking Report Settings dialog box. This dialog box is invoked by selecting **Settings>Tracking Report Settings** or by selecting **Tracking Report Settings** from the Tracking Report view's context menu.

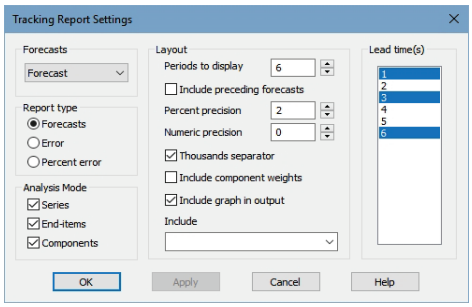

The *Forecasts* section allows you to display either the statistical or adjusted forecasts.

The *Report type* section allows you to display the forecasts, the forecast error or the percent forecast error.

The *Analysis Mode* section allows you to specify which cumulative lead time statistic displays to include.

Most of the items in the *Layout* section are self-explanatory, however, we will comment on a couple of them. The "Periods to display" controls the number of lead times to include in the report. The "Include preceding forecasts"

option displays all archived forecasts that go into the cumulative statistics rather than just the subset in the triangular display.

The *Lead time(s)* box allows you to specify up to four lead times to color code. The selected color coding is used on both the numeric and graphical displays.

The Item Report View

The item report can contain Pareto information, statistical output such as the forecasting model specification and summary statistics as well as time series output such as history, forecasts, fitted values and confidence limits.

To open the Item Report view, select **View>Item Report** or click the yellow View Item Report icon ( ).

If you double-click an item on the Item Report, the Navigator will jump directly to that item. Thus, you can use the Item Report as the equivalent of a Hot List to navigate through the listed items. All of the global report views (yellow icon views) support this kind of navigation.

The content and format of the report is controlled using the Item Report Design dialog box. This dialog box is invoked by selecting **Settings>Item Report Design** or by selecting **Item Report Design** from the Item Report view's context menu.

The Exception Report View

The Exception Report view enables you to quickly find cases where your forecast error or some other performance metric has fallen outside of an acceptable range. Exception reporting reduces the need for manual review of your forecasts and allows you to focus on the items where human attention is most needed.

Forecast Pro provides a wide array of exception reports, some of which monitor the current forecasts and others which monitor your archived forecasts.

Overview

To open the Exception Report view, select **View>Exception Report** or click the yellow Exception Report icon ( ).

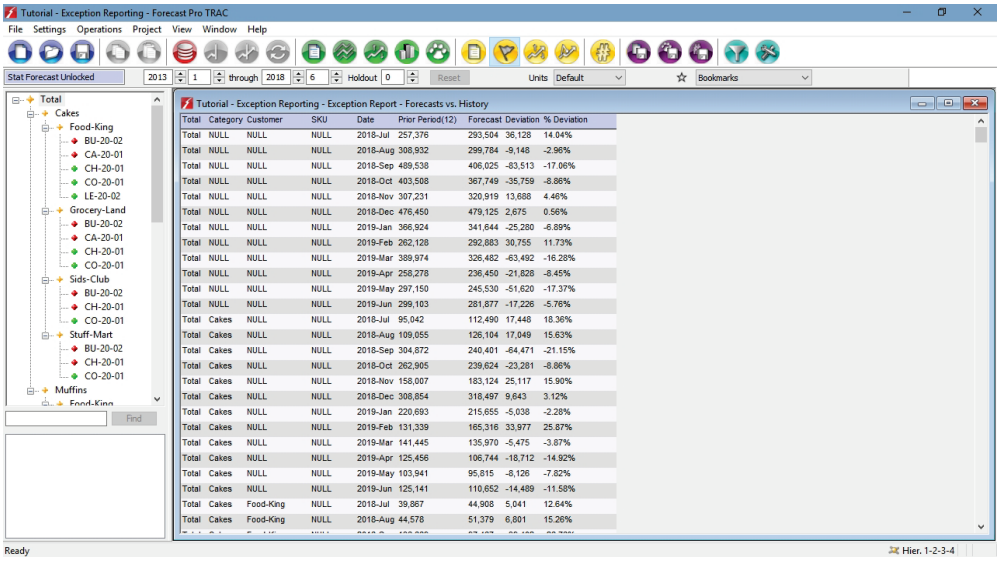

All exception reports are displayed in a format similar to the example shown above. A row is included for every item in the hierarchy. The items can be filtered to show only items that need further examination, as described below. The initial columns identify the item. This is followed by a Date column identifying the exception period. The next two columns display the value being monitored and what it is being monitored against. This is followed by a Deviation column showing the difference and the %Deviation column showing the difference as a percent.

Double-clicking an item on an exception report will cause the Navigator to jump directly to that item. Thus, you can use the exception report as the equivalent of a Hot List to navigate through the listed items. All of the global report views (yellow icon views) support this kind of navigation.

Right clicking on the Exception Report view brings up the Exception Report context menu.

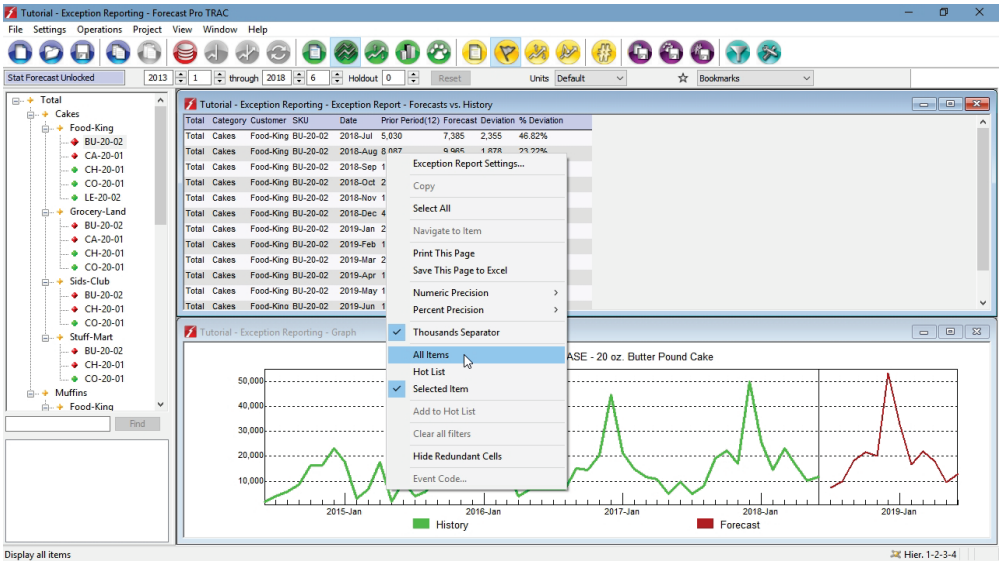

You can select All Items, Hot List or Selected Item to quickly adjust which items should be included in the report.

Exception reports are best used in conjunction with Filters and Bookmarks. Please refer to *Tutorial Lesson 10: Exception Reporting* for an overview of how to use Filters and Bookmarks with Exception reports.

The content and format of the Exception Report is controlled using the Exception Report Settings dialog box. This dialog box is invoked by selecting **Settings>Exception Report Settings** or by selecting **Exception Report Settings** from the Exception Report view's context menu.

There are six exception report types, each of which has a tab on the Exception Report Settings dialog box.

Forecast vs. History

The **Forecast vs. History** report compares forecast values to prior historical values.

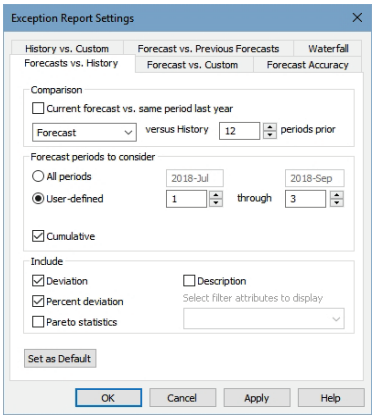

The *Comparison* section at the top of the dialog box is where you define the type of Forecast(s) you want to use (Baseline, Forecast, Statistical) and the number of periods prior you want to compare to. By default, the report shows the current forecast versus the same period last year for all periods in the forecast horizon. Unchecking the "Current forecast vs. same period last year" checkbox allows you to customize the comparison periods.

The *Forecast periods to consider* section allows you to set the number of forecast periods to monitor and to specify if you want to look at each individual period or at the cumulative total across the specified forecast periods. The default setting (cumulative is not checked) monitors all forecast periods individually.

The bottom section, *Include,* allows you to specify which statistics you want to show in the report. By default, the report shows Deviation and % Deviation. If you have included filters in the project (see Custom Filter Fields in the *Setting Up Your Optional Data* section of this chapter), you may add those to the report using the Filter drop down.

Forecast vs. Custom

The **Forecast vs. Custom** report compares the current forecast to any row(s) available in the override grid.

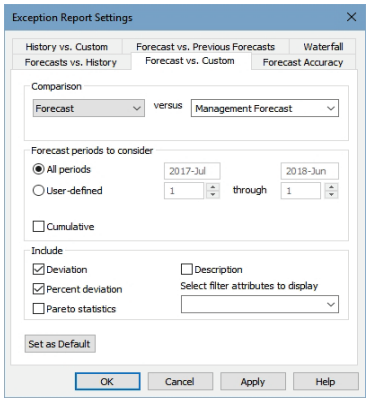

The report shown above (taken from the project *Lesson 10*) compares the current forecast to the Management Forecast (an external row). The comparison drop down menus allow you to change the type of Forecast and the comparison series.

The *Forecast periods to consider* section allows you to set the number of forecast periods to monitor and to specify if you want to look at each individual period or at the cumulative total across the specified forecast periods. The default setting (cumulative is not checked) monitors all forecast periods individually.

The bottom section, *Include,* allows you to specify which statistics you want to show in the report. By default, the report shows Deviation and Percent deviation. If you have included filters in the project (see Custom Filter Fields in the *Setting Up Your Optional Data* section of this chapter), you may add those to the report using the Filter drop down.

Forecast Accuracy

The **Forecast Accuracy** report looks at previously generated forecasts vs. what actually happened.

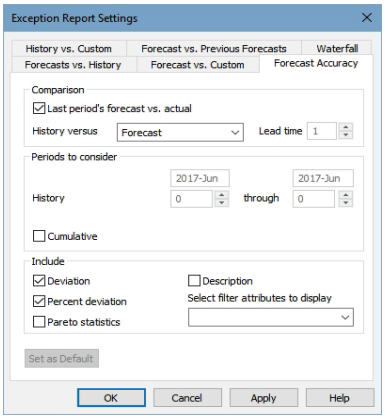

The *Comparison* section at the top of the dialog box is where you define the type of Forecast(s) you want to use (Baseline, Forecast, Statistical) and the *lead* time you want to compare to. A lead time refers to the number of periods ahead of the forecast origin (the period when the forecast is generated) the forecast was made for. Thus, a one-month-ahead forecast would have lead time equals 1, a two-month-ahead forecast would have lead time equals 2, etc. By default, the report compares the forecast we made last period (Forecast(L1)) to what actually happened for the current month only.

The *Periods to consider* section allows you to select how many historic periods to monitor. If *Last period's forecast vs. actual* is selected, this will be shown as the current month (June 2017 in the picture above). Unchecking *Last period's forecast vs. actual* will activate the spinners below the date range. Use the spinners to select the starting and ending historic periods. Selecting "0" will select the most recent historic period, "-1" will select the previous historic period and so on. As you change the spinners, the date above will update to reflect the actual starting and ending periods selected. Check *Cumulative* if you want to look at the total across the specified historical periods.

The bottom section, *Include,* allows you to specify which statistics you want to show in the report. By default, the report shows Deviation and Percent deviation. If you have included filters in the project (see Custom Filter Fields in the *Setting Up Your Optional Data* section of this chapter), you may add those to the report using the Filter drop down.

History vs. Custom

The **History vs. Custom** report compares historical values with any row(s) available in the override grid.

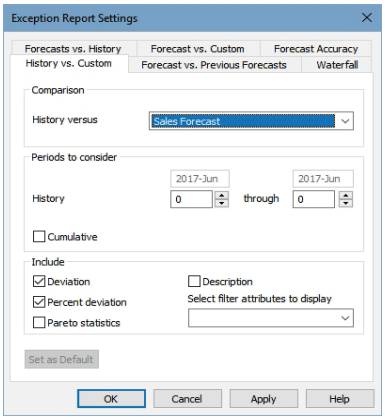

The report shown above (taken from the project *Lesson 10*) compares History to the Sales Forecast (an external variable). The drop down in the *Comparison* section allows you to select which row(s) you wish to compare the history to.

The *Periods to consider* section allows you to select which historic periods to monitor. Use the spinners to select the starting and ending historic periods. Selecting “0” will select the most recent historic period, “-1” will select the previous historic period and so on. As you change the spinners, the date above will update to reflect the actual starting and ending periods selected. Check *Cumulative* if you want to look at the total across the specified historical periods.

The bottom section, *Include,* allows you to specify which statistics you want to show in the report. By default, the report shows Deviation and Percent deviation. If you have included filters in the project (see Custom Filter Fields in the *Setting Up Your Optional Data* section of this chapter), you may add those to the report using the Filter drop down.

Forecasts vs. Previous Forecasts

The **Forecasts vs. Previous Forecasts** report compares the current forecast to previously generated forecasts.

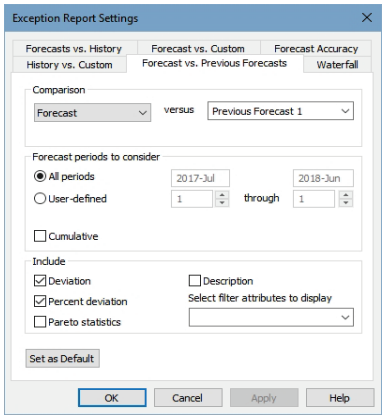

The current report compares the current forecast against the forecast we made last month (Previous Forecast 1). In the *Comparison* section, the left drop down allows you to select which of the current forecasts to consider (Forecast, Baseline or Statistical) and the right drop down allows you to select which previous forecast(s) you want to compare to. Note that you may choose multiple previous forecasts for comparison.

The *Forecast periods to consider* section allows you to set the number of forecast periods to monitor and to specify if you want to look at each individual period or at the cumulative total across the specified forecast periods. The default setting (cumulative is not checked) monitors all forecast periods individually.

The bottom section, *Include,* allows you to specify which statistics you want to show in the report. By default, the report shows Deviation and Percent deviation. If you have included filters in the project (see Custom Filter Fields in the *Setting Up Your Optional Data* section of this chapter), you may add those to the report using the Filter drop down.

Waterfall

The **Waterfall** report provides a complete summary of forecast accuracy across lead times. It doesn't compare items and calculate deviations like the other exception reports. It simply allows you to list statistics from the waterfall reports in a global report that can be sorted and filtered to aid in your review process.

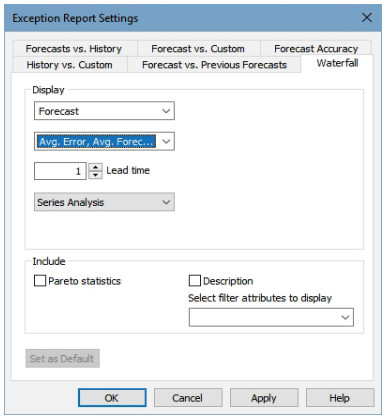

The *Display* section is used to define what to show in the report. The top drop down (currently showing *Forecast*) is used to select which forecasts (Forecast, Baseline and/or Statistical) to include in the report. The drop down below that shows which statistics (Avg. Error, Avg. Forecast and Avg. Perc. Error) to include as additional columns in the report. The Lead time spinner is used to select the lead time for the report. The bottom drop down in the Display section lets you choose the analysis type (Series, End Item-level or Component). Please see *The Tracking Report and Tracking Report Graph Views* in this Command Reference chapter for more details on lead times and analysis types.

The bottom section, *Include,* allows you to specify additional columns to show in the report. If you have included filters in the project (see Custom Filter Fields in the *Setting Up Your Optional Data* section of this chapter), you may add those to the report using the Filter drop down.

The Override Report View

The Override Report view lists items where overrides and/or comments have been made. The display is useful when reviewing the current overrides and also when loading archived override sets.

To open the Override Report view, select **View>Override Report** or click the yellow View Override Report icon ( ).

If you double-click an item on the Override Report, the Navigator will jump directly to that item. Thus, you can use the Override Report as the equivalent

of a Hot List to navigate through the listed items. All of the global report views (yellow icon views) support this kind of navigation.

The content and format of the report is controlled using the Override Report Design dialog box. This dialog box is invoked by selecting **Settings>Override Report Design** or by selecting **Override Report Design** from the Override Report view's context menu.

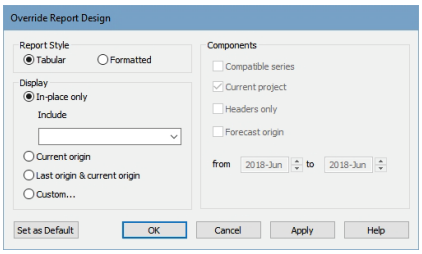

The Display section allows you to control the override sets that will be displayed, to include columns for the Pareto output and to include columns for custom filters (if there are any defined).

In-place only restricts the display to the current in-place override set only. Thus, override sets made in previous states (e.g., in other units and/or hierarchies) will be omitted, as will override sets made for other origins (e.g., previous forecast periods) and override sets associated with other linked projects.

Current origin displays all override sets for the current origin (forecast period). Thus, override sets made in other states (e.g., in other units and/or hierarchies) will be included, but override sets associated with different forecast origins (e.g., previous forecast periods) will be omitted.

Last origin & current origin displays all override sets for the current and immediately previous origin (forecast period). This is often a useful display to view when you have updated the statistical forecasts and are loading override sets from the last forecast period via **Project>Import>Archived Overrides**.

Custom provides complete control over the override sets displayed. Once *Custom* is selected, the *Components* section of the dialog box is activated. *Compatible series* restricts the display to override sets where the state (i.e., units of measure and hierarchy) match the current state. *Headers only* displays the one-line header for each override set, omitting the item-level override

information. The *Forecast origin* option allows you to specify the range of origins to display. Once you check *Forecast origin,* you may use the spinners to set the date range.

The Outlier Report View

The outlier report view lists items where outliers have been detected and/or corrected. To open the outlier report view, select **View>Outlier Report** or click the yellow View Outlier Report icon ( ).

The content and format of the report is controlled using the outlier report's context menu. The context menu also allows you to print the currently displayed report and save it to Excel.

The Outliers tab of the **Settings>Options** dialog box includes an *Include iteration statistics in reports* option. If this option is active the outlier report will include detailed statistics describing the outlier detection. Consult the *Forecast Pro Statistical Reference Manual* for details.

The Numeric Output View

The Numeric Output view displays the contents and format of the currently specified Numeric Output file. When designing the Numeric Output file, it is useful to have the Numeric Output view open.

Selecting **Settings>Numeric Output Design** opens the Numeric Output Design dialog box which is used to specify the format and content of the Numeric Output file (and the Numeric Output view).

The Numeric Output Design dialog box contains both a Set as Default button and an OK button. Clicking the OK button will save your current settings for use with the current project. Clicking the Set as Default button will save your current settings for use with the current project and also save them as the default settings for all new projects.

The Numeric Output Design dialog box contains three pages or tabs. We will discuss each tab in turn.

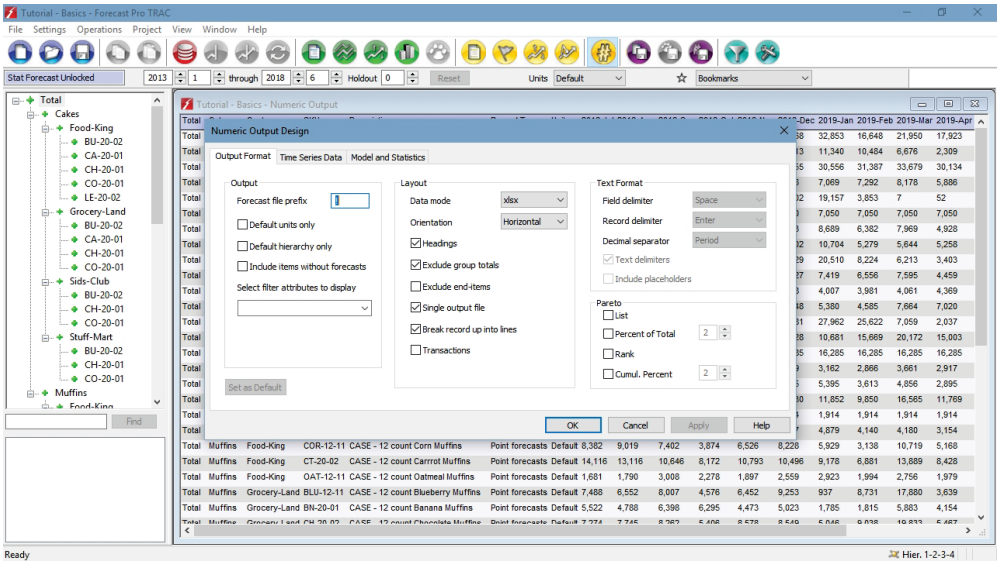

The Output Format tab is used to specify the basic layout for the Numeric Output file.

Forecast file prefix. Specify the character(s) to use as the default prefix when naming the Numeric Output file. If the *Single Output file* option is selected in the Layout section, the default Numeric Output filename will be the prefix followed by the project name followed by "Numeric Output". If the *Single Output file* option is not selected, each output filename will be the prefix followed by the corresponding input filename.

Default units only. If this option is selected, the Numeric Output file will save the forecasts and other components using the default units (i.e., the units defined in the input data source). If this option is not selected, the Numeric Output file save the forecasts and other components using the units of measure currently selected in Forecast Pro.

Default hierarchy only. If this option is selected, the Numeric Output file will save the forecasts and other components using the default hierarchy (i.e., the hierarchy defined in the input data source). If this option is not selected, the Numeric Output file will save the forecasts and other components using the hierarchy currently selected in Forecast Pro.

Include items without forecasts. If this option is selected, the Numeric Output file will include entries for any discontinued items that appear on the Navigator.

If you have included filters in the project (see Custom Filter Fields in the *Setting Up Your Optional Data* section of this Command Reference), you may add those to the Numeric Output using the *Select filter fields to display* drop down.

The Layout section of this tab allows you to specify the default data mode, the orientation (row vs. column), whether or not to include group level output and whether or not to include item-level data.

Single output file. If *Single output file* is selected, a single Numeric Output file containing information for all forecasted items will be written. If *Single output file* is not selected, then a separate Numeric Output file will be written for each input file.

Break record up into lines. If selected, each Numeric Output component specified on the Time Series Data tab (see below) will appear on a separate line of the Numeric Output file.

The Text Format section of this tab allows you to specify the conventions to use when outputting a text file. You will only be able to edit these settings when the data mode is set to Txt (text) or Mlt (text input file format).

The Pareto section of this tab allows you to include the Pareto codes and associated Pareto statistics in the Numeric Output view and output.

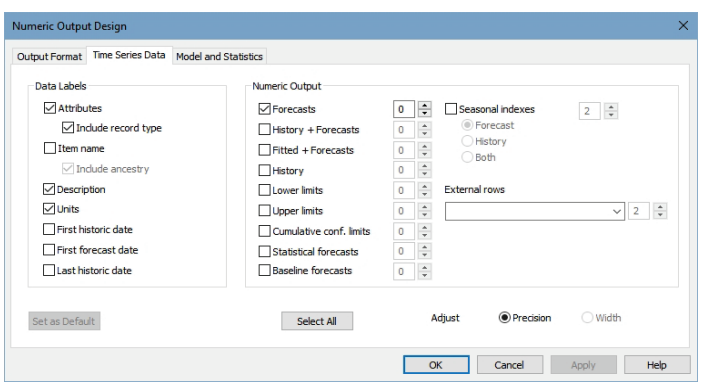

The Time Series Data tab is used to define the data label fields, to select the forecast components to include and to set their precision. If you are outputting a fixed width text file then you may also set the column widths.

Attributes. Include separate attribute fields for each level of the hierarchy. If you used attribute fields in your input data files to define groups then you will most likely want to select this option so that your output file will match your input file. Consult the *Setting Up Your Historic Data* chapter for a discussion of attribute fields.

Include record type. Include a record type field that identifies which forecast component is being output. This is useful if you use the *Break record up into lines* option and are outputting more than one numeric output component.

Item Name. Include a field listing the "name" of the variable as it appears on the Navigator. If the *Include ancestry* option is active the ancestry will be shown as part of the item name. The variable names used for all numeric output components other than "Forecasts" will include a suffix indicating the record type.

Description, Units, First historic date, First forecast date and *Last historic date.* Include fields for the selected options.

The Numeric Output section of the tab allows you to select the Numeric Output components you would like to include.

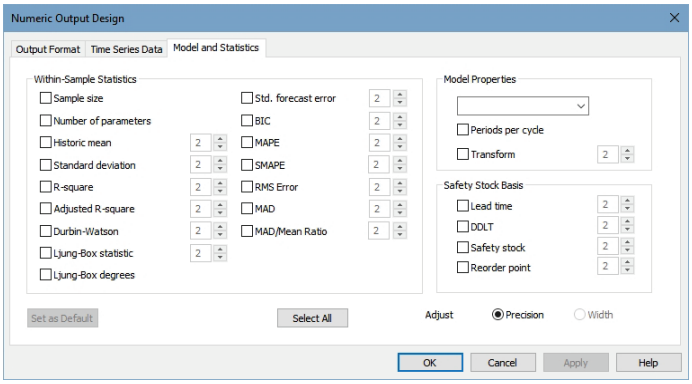

The Model and Statistics tab is used to define the within-sample statistics and model details that you wish to include and to set their precision. If you are outputting a fixed width text file then you may also set the column widths.

Most of the options are self-explanatory. However, we have noted some details below.

Model Properties. The model properties drop down box contains a number of different fields that describe the forecasting model selected. *ModelSpec* is a concise description of the model used. The notation that is used here is also used in the Forecast Report. Consult the *Forecast Pro Statistical Reference Manual* for further details.

Transform. The Box-Cox transformation power. 1 indicates no transform, .5 the square root and 0 the natural log.

The Event Manager

The Event Manager is used to create and modify event schedules. This section documents the options available on the Event Manager dialog boxes.

If you are new to event modeling it is strongly recommended that you work your way through *Building Event Models* in the tutorial section of this manual.

The Event Manager works in conjunction with the graph view and can only be accessed when the graph view is open. You invoke the Event Manager by selecting **Manage events** on the graph view's context menu.

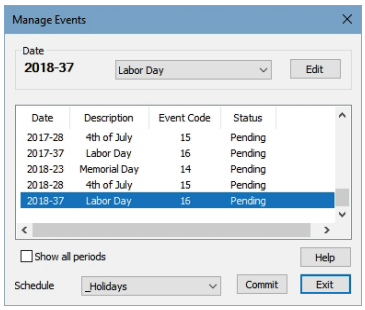

The Manage Events dialog box is shown above. The *Date drop down control* at the top of the box is used to assign event codes to the current schedule. You select points on the graph by clicking on them and then use the drop down to the right of the date to associate the selection with a specific event. Holding down the <Ctrl> key when selecting points on the graph allows you to select multiple points.

The display grid in the middle of the dialog box shows the event code definitions on the current schedule along with their status. *Pending* status indicates that the changes are not yet reflected in the current forecast model. Clicking the Commit button will accept the current changes and build the specified event model.

The *Show all periods* option will include all periods in the display grid regardless of whether or not events have been defined.

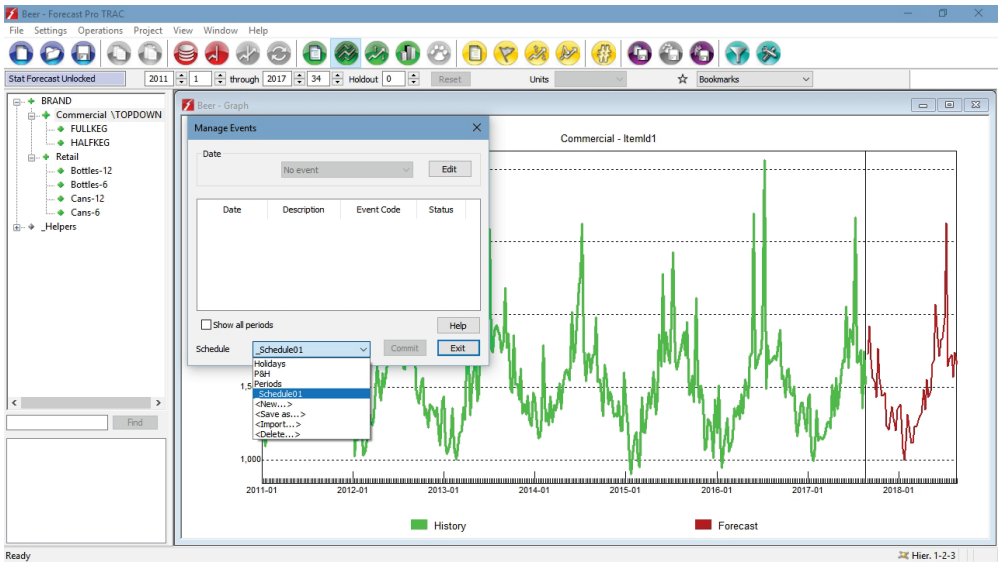

The *Schedule drop down* is used for several purposes. When the drop down is closed it displays the name of the event schedule associated with the current item. Opening the drop down displays the currently available event schedules and allows you to select one. It also provides options to disassociate the currently selected schedule from the currently displayed item (New), rename the current schedule (Save as), import events from other schedules into the current schedule and to delete the current schedule from the Forecast Pro database.

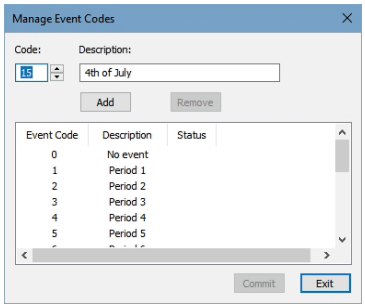

Clicking the *Edit button* invokes the Manage Event Codes dialog box allowing you to define new event codes and edit existing ones. You can select an event code using the spinners or by clicking on it in the display grid. After editing a description you click the Add button to update the display grid. You can remove an event code by selecting it in the display grid and clicking the Remove button. All modifications will have a status of *Pending* until you accept them using the Commit button. The Exit button is used to close the Manage Event Codes dialog box and return to the event manager.

The Create Helper Variable Dialog Box

The Create Helper Variable dialog box is used to create helper variables based on the current Navigator selection.

Helper variables are used in conjunction with event modeling, forecast by analogy and the weighting transformation. They are identified by the use of a preceding or trailing underbar in their name.

Usually event schedules are created interactively using the Event Manager. However, if you wish, you can create helper variables containing the event schedules in Excel or in a database and import them into your project using the Data Manager.

Forecast by analogy requires a helper variable containing the analogy series. Using a weighting transformation requires a helper variable containing the weights. Depending on your needs, you can either create these helper variables using the Create Helper Variable dialog box or you can create them in Excel or in a database and import them into your project using the Data

Manager (see *Helper Files* in the *Setting Up Your Optional Data* section of this Command Reference chapter).

To use the Create Helper Variable dialog box, you first select the item you wish to base the helper variable upon on the Navigator, then right click to invoke the Navigator's context menu and then select **Helper variable>Create**.

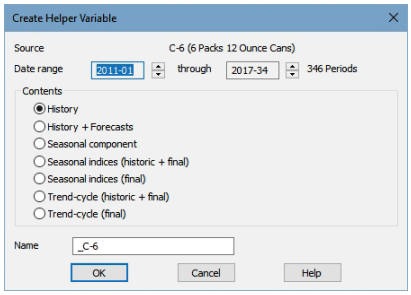

To create the helper variable you select the desired contents, adjust the date range (if necessary) and click OK. Forecast Pro will then create the helper variable and save it in your project.

By default, the helper variable will be named "*_selection*" where *selection* is the name of your current Navigator selection. If you wish to name the helper variable something else you can rename it via the Name field prior to clicking OK.

The *Seasonal component* option is designed for use with the custom component model. The custom component model allows you to import seasonal indexes from a helper file. The Seasonal component option on the Create Helper Variable dialog box estimates a set of seasonal indexes which can then be imported into a custom component model. For example, say you want to build a custom component model for a given SKU that uses the seasonal indexes calculated from the SKU's corresponding product-level group. You would select the product-level group, create the helper variable and then import the indexes from the helper when you build a custom component model for the SKU.

If you wish to delete a helper variable you must first save your project, then select the helper variable on the Navigator, then right click to invoke the Navigator's context menu and then select **Helper variable>Delete**.

Menu Items and Dialog Boxes

This section covers all of the Forecast Pro menu items and their associated dialog boxes.

The File Menu

The file menu is used to manipulate forecast project files, save ODBC connections strings for use in command-line operation, print the currently displayed forecast report and exit Forecast Pro.

File>New is used to open a new forecast project. This task can also be accomplished using the blue New Project icon ( ).

File>Open is used to open an existing forecast project. This task can also be accomplished using the blue Open Project icon ( ).

File>Close is used to close the current forecast project.

File>Save is used to save the current forecast project using the currently specified name. This task can also be accomplished using the blue Save Project icon ( ).

File>Save as is used to name and save the current forecast project.

File>Print is used to print the currently selected view.

File>Print Preview is used to display the currently selected view formatted as it will be printed (see **File>Print** above). Checking the appearance of the output prior to printing can save you trips to the printer.

File>Print Setup is used to change various printing options.

File>Exit is used to exit Forecast Pro.

The Settings Menu

The settings menu provides access to the Options dialog box which allows you to change your project settings and to other dialog boxes which allow you to customize the various Forecast Pro views.

With the exception of **Settings>Options**, which is described below, all of the other menu options lead to dialog boxes that are described in detail in the *Forecast Pro Interface* chapter.

Settings>Options

Selecting **Settings>Options** opens the Options dialog box which is used to change various Forecast Pro settings.

The Settings dialog box contains both a Set as Default button and an OK button. Clicking the OK button will save your current settings for use with the current project. Clicking the Set as Default button will save your current settings for use with the current project and also save them as the default settings for all new projects.

The Settings dialog box contains eight pages or tabs. We will discuss each in turn.

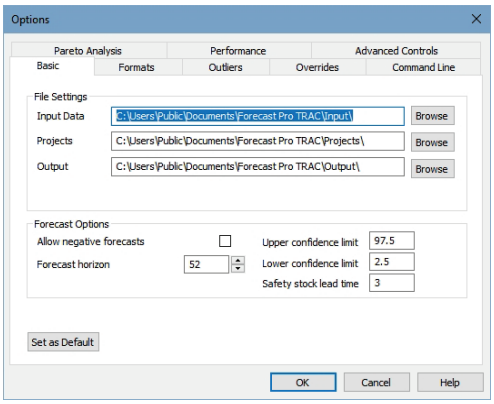

The Basic tab contains frequently accessed settings that every user should understand.

File Settings:

Input Data. Specify the complete path name of the directory where the program should read data.

Projects. Specify the complete path name of the directory where the program should write the project files.

Output. Specify the complete path name of the directory where the program should write all output files (e.g., Numeric Output files, Forecast Reports, etc.).

Forecast Options:

Allow negative forecasts. Most business data are inherently positive, although there are exceptions like telephone connection gain or sales data that include returns as negatives. Users with positive data are understandably disturbed when their forecasts turn negative. This can occur when Forecast Pro captures a downward trend at the end of your historic data. If *Allow negative forecasts* is turned off, Forecast Pro will clip forecasts to zero that would otherwise be negative. If *Allow negative forecasts* is turned on, then Forecast Pro takes no special actions to avoid negative forecasts.

Forecast horizon. Specify how many periods ahead you want Forecast Pro to forecast.

Upper confidence limit. Specify the percentile for the upper confidence limits. The upper confidence limit for a specific item can also be specified using the \Upper= modifier.

Lower confidence limit. Specify the percentile for lower confidence limits. The lower confidence limit for a specific item can also be specified using the \Lower= modifier.

Safety stock lead time. Specify the lead time for the safety stocks. The specified lead time appears highlighted in the Safety Stock section of the Forecast Report and determines the values written to the Numeric Output file. Fractional entries (e.g., 1.5) are permitted.

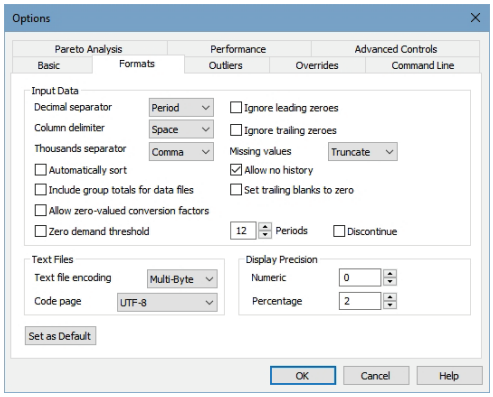

The Formats tab controls the formatting of the input and output.

Input Data:

Decimal separator. Specify whether a period or a comma is used for the decimal point. The program's default is the convention appropriate for the United States, which is the period.

Column delimiter. Specify whether items in an MLT file are separated by spaces, commas or tabs (i.e., are the files space delimited, comma delimited or tab delimited).

Thousands separator. Specify the thousands separator used for input MLT files. The choices are none, comma and period. Forecast Pro does not use thousand separators in output.

Automatically sort. By default, the Automatically sort option is selected and Forecast Pro will merge and sort all historic data sources prior to reading them in. This can be useful when the order of the items in your data sources do not match the desired hierarchy. If the Automatically sort option is not selected, Forecast Pro reads in your historic data sources sequentially and bases the structure of the Navigator's hierarchy on the order of appearance of the items in the data sources.

Include group totals for data files. If this switch is on, Forecast Pro will automatically include a group total for each data file, table or query listed on the Historic Data row of the Data Manager. Including or not including a group total for the data files will yield different hierarchies on the Navigator.

Allow zero valued conversion factors. If this switch is on, Forecast Pro will allow blanks or zeros or missing rows in a conversions file specified in the Data Manager. Forecast Pro will give warnings about the missing data, and changing project units using the Units drop down on the dialog bar is not enabled. However, the converted rows are available in the Overrides Report and Graph views. Blanks will be considered to be zeros, and items with either zero or blank conversion factors will consequently not be included in the totals for those conversion factors. If this switch is not on, Forecast Pro will give an error and not read in a conversions file with blanks, zeros or missing rows.

Zero demand threshold. If you set the zero demand threshold to *"N"* and the Discontinue option is not selected, Forecast Pro's expert selection mode will generate a forecast of zero for any item where the last *N* historic observations are equal to zero. If you set the zero demand threshold to *"N"* and the Discontinue option is selected, Forecast Pro's expert selection mode will discontinue and therefore not forecast any item where the last *N* historic observations are equal to zero.

Ignore leading zeros. If this option is selected, leading zeros prior to the first nonzero data point will be ignored (i.e., the data has not started yet and the zeros are interpreted as placeholders). For a more complete discussion on the handling of missing data and the distinction between a zero value and a missing value, please refer to the Missing Values and Zeros section in the *Setting Up Your Historic Data* chapter.

Ignore trailing zeros. If this option is selected, trailing zeros after the last nonzero data point will be ignored (i.e., the data has ended and the zeros are interpreted as placeholders). This means that the time series would be considered "dead" in multiple-level scripts and not forecasted. In nonmultiple-level scripts the forecasts for the time series would begin after the last nonzero point. For a more complete discussion on the handling of missing data and the distinction between a zero value and a missing value, please refer to the Missing Values and Zeros section in the *Setting Up Your Historic Data* chapter.

Missing Values. Specify how you wish the program to treat missing values. The options are: truncate the data set (i.e. discard all data that precedes the missing value), impute the missing data via a linear interpolation or set missing data points to zero. Regardless of how this option is set, missing

values that appear prior to the first data point are ignored. For a more complete discussion on the handling of missing data and the distinction between a zero value and a missing value, please refer to the Missing Values and Zeros section in the chapter entitled *Setting Up Your Historic Data*.

Allow no history. If this switch is on, Forecast Pro will create a forecast of zero per period for items that have no historic data. If the switch is off, items with no historic demand will be omitted and not appear on the Navigator.

Set trailing blanks to zero. When historic data is read into Forecast Pro, the program establishes a global ending date for the history which corresponds to the most recent observation available *for any item being read in*. If a given item does not have a historic observation for the global ending date, the item is flagged as "dead." Dead items are not forecasted but their demand history goes into the group totals. This can be an issue if your data source codes zeroes as missing (e.g., blank cells in Excel, no transaction in a table or query, etc.). The *Set trailing blanks to zero* option will replace missing trailing observations with zeroes. For a more complete discussion on the handling of missing data and the distinction between a zero value and a missing value, please refer to the Missing Values and Zeros section in the *Setting Up Your Historic Data* chapter.

Text Files:

Text file encoding. Specify the format to use when writing two-byte characters to text files. This option is only of relevance if your data or scripts utilize non-ASCII characters. (Non-Latin based languages such as Japanese, Chinese, Russian, Arabic, etc. use non-ASCII characters.)

Code page. Specify the Windows code page to use when text file encoding is set to multi-byte. The code page specifies the character set to use when writing a multi-byte text file. (Most code pages are designed to support a specific language.)

Display Precision:

Numeric. Specify the precision (i.e., number of decimal places) to use when displaying time series data in the view windows.

Percentage. Specify the precision (i.e., number of decimal places) to use when displaying percentages in the view windows.

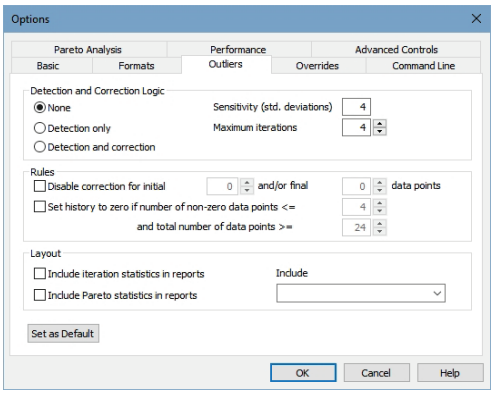

The Outliers tab is used to specify whether or not to detect and/or correct outliers. It also allows you to set the sensitivity settings for outlier detection and to dictate the amount of detail in the outlier report.

Detection and Correction Logic:

None turns outlier detection off. The forecasts will be generated using the uncorrected history.

Detection only will detect outliers and display the suggested corrected values, but the forecasts will be generated using the uncorrected history.

Detection and correction will detect outliers and will automatically use the corrected values when generating forecasts.

Sensitivity (std. deviations) allows you to set the sensitivity of the outlier detection algorithm. If a given fitted error exceeds this threshold and it is the largest error detected during the current iteration it will be flagged as an outlier. Consult the *Forecast Pro Statistical Reference Manual* for details on how outlier detection and correction works.

Maximum iterations allows you to set the maximum number of iterations permitted during outlier detection for a given item. This setting thereby also defines the maximum number of outliers that can be detected for a given item. Consult the *Forecast Pro Statistical Reference Manual* for details on how outlier detection and correction works.

Rules:

Disable correction for initial X and/or final Y data points allows you to disable correction for data points at the beginning or end of an item. The outlier detection algorithm is sometimes less robust for the beginning and ending historic periods. Disabling correction for the initial and ending periods can avoid possible misidentification of outliers in these cases. Consult the *Forecast Pro Statistical Reference Manual* for details.

Set history to zero if number of non-zero data point is <= X and total number of data points >= Y allows you to set the entire history to zero for an item if the number of zeros meets a specified threshold and the total number of historic data periods exceeds a specified threshold.

Layout:

Include iteration statistics in reports allows you to include detailed statistics describing the outlier detection process in the outlier report. Consult the *Forecast Pro Statistical Reference Manual* for details.

The other options in the *Layout* section allow you to include columns for the Pareto output and to include columns for custom filters (if there are any defined).

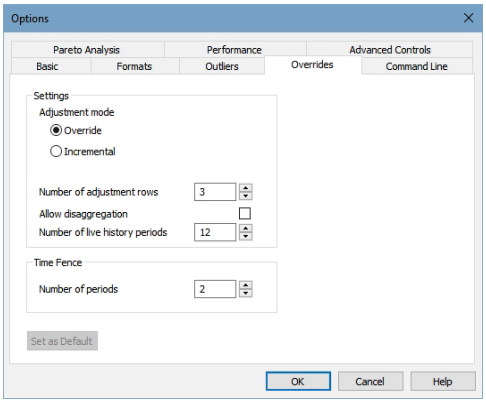

The Overrides tab controls operation and display of the Override view.

Adjustment mode. If this selection is set to “Override” any overrides entered will replace the corresponding baseline forecasts (or overrides entered on a preceding row). If this switch is set to “Incremental” overrides you enter will

be *added* as incremental adjustments to the baseline forecast rather than replacing it. Thus the final forecast will equal the sum of the statistical forecast and all overrides entered. If you attempt to change this setting when overrides exist for the current project, you will receive a warning and all overrides will be erased.

Number of adjustment rows. This setting controls the default number of adjustment rows displayed in the Override view.

Allow disaggregation. If this switch is on, the Override window will include a "Disaggregate to" drop down box, allowing you to disaggregate (i.e., push down) any overrides made at a group level to a lower-level of the hierarchy. The disaggregation is based on a proportional allocation.

Number of live history periods. This setting controls the number of historic periods shown on the override grid when the history display option is set to *Time Series*. When working with large projects this setting can have an impact on performance (i.e., speed).

Time Fence: Number of periods. This option allows you to set the default length for the time fence. This setting can be overridden on an item-by-item basis using the \TIMEFENCE= modifier.

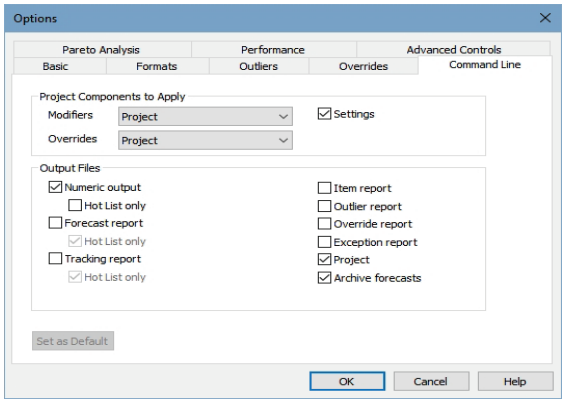

Forecast Pro supports command-line operation allowing you to create forecasts in a "hands-off" mode. For a complete description of how to run the product from the command line consult the *Operating From the Command Line* lesson in the tutorial and the *Command-line Operation* section of this chapter.

The Command Line tab allows you to dictate the options to apply if the current project is run via the command line.

Project Components to Apply:

Modifiers. This dropdown is used to control which modifiers to use when running from the command line. There are four options available: None, Project (default), Data Source and Project and Data Source. If None is selected, expert selection will be used for all items. If Project is selected, Forecast Pro will apply modifiers as specified in the Project's Navigator. If Data Source is selected, Forecast Pro will consider only modifiers specified in the Data Manager. If Project and Data Source is specified, Forecast Pro will apply modifiers specified in the Navigator for items that do not have a modifier in the Data Source and apply Data Source modifiers to items that have them.

Overrides. This dropdown is used to control which overrides to apply when running from the command line. There are four options available: None, Project (default), Data Source and Project and Data Source. If None is selected, all overrides in the project and Data Manager override files will be ignored. If Project is selected, Forecast Pro will apply any matching overrides/comments in the project to the statistical forecasts. If Data Source is selected, Forecast Pro will apply only overrides specified in the Data Manager. If Project and Data Source is specified, Forecast Pro will first load overrides from the Project and then from the Data Source. Thus, the Data source will take precedence over the Project for all override cell values.

Settings. If this option is selected, the command-line run will use all of the settings associated with the project (e.g., all settings in the Settings dialog box, numeric output formatting, etc.). If this option is not selected, the default settings (i.e., the settings that have been saved as the defaults for use with new projects) will be used when running from the command line. Project settings are saved in the project file (*.FPProj). Default settings are stored in the ForecastProTRAC.ini file.

Output Files:

Numeric output. If this option is selected, the command-line run will save the numeric output. You have the option of saving this file for all items forecasted or (if you select Hot List only) just for the items listed on the Hot List.

Forecast Report. If this option is selected, the command-line run will save a formatted forecast report file to Excel. You have the option of saving this file for all items forecasted or (if you select Hot List only) just for the items listed on the Hot List. Keep in mind that these files can potentially be very large.

Tracking Report. If this option is selected, the command-line run will save a forecast tracking report file to Excel. You have the option of saving this file for all items forecasted or (if you select Hot List only) just for the items listed on the Hot List. Keep in mind that these files can potentially be very large.

Item report, Outlier report, Override report and Exception report. If any of these options are selected, the command-line run will save the specified report(s). The format and contents (including all items vs. Hot List only) will match the current settings in the project (or the defaults if you've specified to use them).

Project. If this option is selected, the command-line run will update and save a new project file.

Archive forecasts. If this option is selected, the command-line run will update and save the TRAC database file. You should select this option if you are archiving forecasts for tracking purposes.

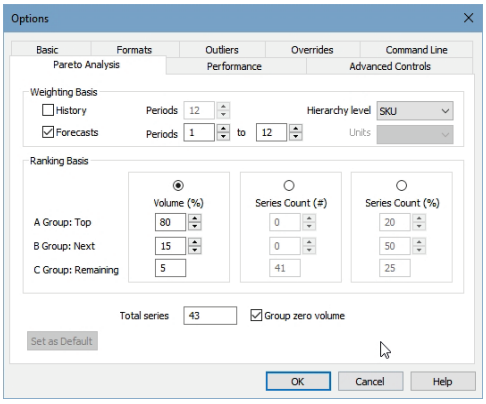

The Pareto Analysis tab allows you to specify how the ABC codes should be assigned.

Weighting Basis:

History and *Forecasts*. These settings allows you to indicate whether the ranking should be performed using the historic demand or the current forecasts. The *Period* spinner(s) allows you to set the span for your selection.

Hierarchy level. The ABC codes are assigned to a single level of the hierarchy. The Hierarchy level drop down allows you to pick the level to use.

Units. If you have defined conversion factors, this option allows you to select the unit of measure on which to base the ABC code assignments. *Default* refers to the unit of measure used for the historic demand data you read into Forecast Pro.

Ranking Basis:

Volume (%). This option allows you to set the codes based upon the cumulative volume of each group. For example, if A is set to 80%, B is set to 15% and C is set to 5%, Forecast Pro will first rank all of the items from highest volume to lowest volume and then sequentially place the top ranked items into the A group until their cumulative volume equals or exceeds 80% of the total volume. It will then sequentially place items into the B group until the cumulative volume of the A and B items equals or exceeds 95% (80% + 15%). The remaining items (5% or less of the cumulative volume) are assigned type C.

Series Count (#). This option allows you to assign specific numbers of items to each group. For example, if you had 1,000 items and you set the A group to 100, the B group to 200 and the C group to 700—the 100 highest volume items would be assigned to group A, the next highest 200 items would be assigned to group B and the remaining 700 would be assigned to group C.

Series Count (%). This option allows you to assign a specified percentage of the total number of items you have to each group. For example if A is set to 20%, B is set to 50% and C is set to 30%, Forecast Pro will first rank all of the items from highest volume to lowest volume and then sequentially place the top ranked items into the A group until the number of A group items equals 20% of the total number of items. It will then sequentially place items into the B group until the number of B group items equals 50% of the total number of

items. The remaining items (30% of the total number of series) are assigned type C.

Group zero volume. Weighting Basis: If this option is on, any items having zero volume are placed into group D.

Total series. This displays the total number of series in the selected hierarchy level.

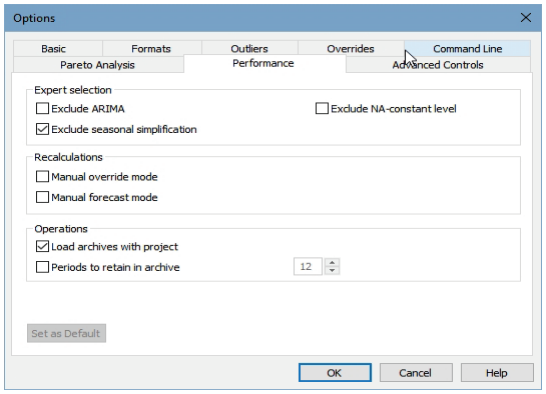

The Performance tab allows you to set expert selection options, select manual recalculation modes and set project operations. Adjusting these settings can improve processing speed and decrease time waiting for project operations and calculations.

Expert Selection:

Exclude ARIMA. If this switch is on Forecast Pro's expert selection mode will not consider Box-Jenkins models. This will speed the forecast generation process considerably.

Exclude seasonal simplification. If this switch is on Forecast Pro's expert selection mode will not consider seasonally simplified forms of exponential smoothing models. If you are forecasting data where the number of periods per cycle is greater than 13 (e.g., weekly data) this will speed the forecast generation process considerably.

Exclude NA-constant level. If this switch is on Forecast Pro's expert selection mode and automatic exponential smoothing mode will not consider the NA-

constant level model. This will speed the forecast generation process considerably.

Recalculations:

Manual override mode. If this switch is on, the override window will include a Retain button as well as a Commit button. The Retain button is used to apply and save overrides without reconciling the hierarchy. This is useful when working with large complex hierarchies where the reconciliation process takes some time. In manual reconciliation mode you would typically enter and retain all of the overrides you wish to make (i.e., for multiple items) and then click the Commit button or the red Recalculate icon to perform the reconciliation once, rather than having to wait for the program to reconcile after each item that you override.

Manual forecast mode. If this option is not selected, Forecast Pro will automatically recalculate the forecast and re-reconcile the hierarchy every time you apply a forecast modifier. For large hierarchies this can be time consuming. If you turn this option on, when you specify a forecast modifier the forecasts will not be updated until you click the red Forecast icon ( ) or the red Recalculate icon ( ). This allows you to specify multiple forecast modifiers without having to wait for the program to recalculate the forecasts in between each specification.

Operations:

Load archives with project. If this option is selected, Forecast Pro will cache the forecast archive information in memory when you open a project. For very large projects this can take some time. If this option is deselected, the forecast archive information will not be cached until the first time you execute an action that requires access to the forecast archive (e.g., opening the forecast tracking report or opening an exception report which uses the archive). Deselecting this option will speed up opening projects and decrease memory usage for individuals who do not need access to the information in the forecast archive.

Periods to retain in archive. This option allows you to limit the number of previous-origin forecasts (i.e., forecasts you generated in prior forecasting periods) to save in the forecast archive. Reducing the number of archived

forecasts will reduce the size of the project's database and speed up program operations.

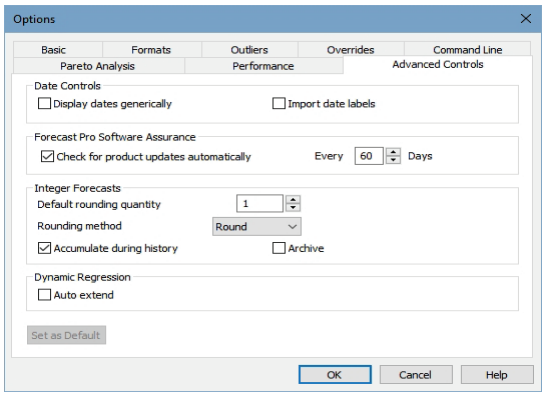

The Advanced Controls tab allows you to set Integer Forecast settings, Date Control features and adjust how frequently Forecast Pro checks for program updates.

Integer Forecasts use a special form of rounding that keeps track of the "remainder" (rounded amount) and applies it to the next period's forecast value. This is particularly useful for low-volume forecasts or when there are large minimum order quantities, where traditional rounding can result in biased forecasts.

Consider a flat forecast of 0.4 units per month for the next 12 months. If you were to use the ROUND() function to round the forecast to the nearest integer the forecast would become, "0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0". An integer forecast row for the same forecast would be, "0, 1, 0, 1, 0, 0, 1, 0, 1, 0, 0, 1".

Integer Forecasts:

Default rounding quantity. Specifies the integer rounding quantity (the basis to round to the nearest multiple of) for Integer forecasts.

Rounding method. This drop down allows you to choose round, round up or round down to the nearest multiple of the default rounding quantity.

Accumulate during history. Selecting this option instructs Forecast Pro to start tracking the remainder during the history rather than starting with the first forecast—this is generally desirable.

Archive. Saves Integer Forecasts to the project's database. Archived Integer forecasts can be accessed through database queries.

Date Controls:

Display dates generically. If this option is selected then monthly labels will be displayed using period numbers rather than the names of the months. For example, "2018-01" will be used rather than "2018-Jan". This option is often used when working with a fiscal calendar where displaying the names of the months could be confusing.

Import date labels. If enabled, Forecast Pro will import date labels from either the helper or historic data source (in the header row or column). If using a helper file, create a helper variable called *_Date* for the date labels. Date labels may be either text or numeric. If the labels are numeric and regularly spaced (time between adjacent periods is always the same), the labels will be extrapolated into future forecast periods. Numeric labels will also be formatted according to the User's locale settings, as specified in Microsoft Windows settings. For ODBC data inputs, date labels must be provided in a helper file. Forecast Pro will prioritize labels in a helper file over those provided in the historic data source. These imported labels will be used in all views instead of the Forecast Pro's standard dates. If *Import data labels* is not checked and a helper file includes a helper variable called *_Date,* Forecast Pro will treat *_Date* like any other variable and not use it for date labels.

Check for product updates automatically. If this option is selected, Forecast Pro will periodically check online for product updates and alert you if they are available. If your support and maintenance contract is up to date, product updates are free.

Dynamic Regression:

Auto extend. If this option is selected, Forecast Pro will generate forecasts for explanatory variables that do not have input values for the forecast horizon. Forecast Pro will use any provided values and only generate forecasts where they are missing. If this option is not selected, a dynamic regression model is selected for a given item and explanatory variable values are not provided for the entire forecast horizon, Forecast Pro will use expert selection instead of the specified regression model for that item.

The Operations Menu

The operations menu is used to perform basic program operations. Most of the actions can also be performed via icons.

Operations>Data Manager

Operations>Data Manager accesses the Data Manger dialog box allowing you to read in and forecast your historic data as well as read any other optional data you wish to work with. This dialog box can also be accessed via the red Data Manager icon ( ). Please read the chapter entitled *Using the Data Manager* for complete details on the Data Manager.

Operations>Read Data

Operations>Read Data reads in the Historic Data defined in the Data Manager from its primary source along with any other inputs where the *on Read* option is selected. This task can also be accomplished using the red Read Data icon ( ).

Operations>Forecast

Operations>Forecast generates statistical forecasts. This option is only available when a historic data has been defined in the Data Manager and the data has been read in. This task can also be accomplished using the red Forecast icon ( ).

Operations>Recalculate

Operations>Recalculate This option is only available when *Manual override mode* or *Manual forecast mode* is selected on the Performance tab of the Settings dialog box. Selecting recalculate will reconcile all forecasts and overrides. This task can also be accomplished using the red Recalculate icon ( ).

Operations>Initialize Archive

Operations>Initialize Archive is used to generate statistical forecasts for previous forecast periods and to insert them into the Forecast Pro database. Essentially this is a form of simulation that allows you to initialize the forecast archive with statistical forecasts.

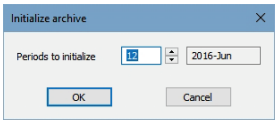

When you select **Operations>Initialize Archive** the Initialize Archive dialog box appears allowing you to specify how many periods you wish to initialize. Forecasts will be generated for the specified period and all subsequent periods up to the current origin.

In the example above, the periods to initialize is set to 12, which corresponds with a forecast origin of June 2016. Thus, the initialization would begin by generating forecasts using data up to and including June 2016 and inserting them into the database. It would then generate forecasts using data up to and including July 2016 and insert them into the database. It would continue this process until all 12 forecasts sets had been generated and inserted into the database.

Operations>Time Fence>Enable

Operations>Time Fence>Enable turns on the time fence for the current project. Once you turn a project's time fence on, *it cannot be turned off!*

Operations>Time Fence>Freeze Next Period

Operations>Time Fence>Freeze Next Period will increase the length of all time fences in the current project by one period for the current forecast origin. If you subsequently, update your data with the next observation and create a new forecast (thus changing the forecast origin) the time fence lengths will revert back to their pre-Freeze-Next-Period settings. This option provides a way of insuring that after the forecasts are finalized someone doesn't inadvertently alter the frozen periods' forecasts prior to the data being

updated, new forecasts being generated and the frozen periods' forecast becoming an officially fenced value.

Operations>Inherit overrides

Operations>Inherit overrides will move all in-place overrides into the Inherited overrides row. This option will be greyed out if there are unsaved changes to the project—so you may have to save the project prior to execution of this command.

Operations>Unlock forecasts

Operations>Unlock forecasts will allow the user to continue to update the statistical forecasts while retaining any inherited overrides that may be in place.

The Project Menu

The project menu is used to import settings, Hot Lists, archived forecasts, archived overrides and calculated rows from other projects, to export (i.e., save to disk) Numeric Output files and report files, to set project passwords, to create and maintain Super Projects and to perform database maintenance.

Project>Import

Project>Import>Project Settings From is used to import project settings from a saved project into the current project. The settings include all selections on the Options dialog box as well as formatting for reports and output files. Because the settings control how data is read into Forecast Pro, you may only import project settings prior to reading in the data.

Project>Import>Hot List From is used to import a Hot List from a saved project or from an Excel file into the current project. Excel Hot List files are created using **Project>Export>Hot List** or using the **Export to** option on the Hot List's context menu. If a Hot List currently exists you will have the option of deleting it prior to importing the new Hot List.

Project>Import>Archived Forecasts From is used to import archived forecasts from a saved project into the current project. When you select this

option a dialog box appears allowing you to specify the origin date for the archived forecast sets that you wish to import.

Project>Import>Archived Overrides is used to load override sets from the current Forecast Pro database. The most common use for this facility is loading override sets made last period after updating the statistical forecasts.

In order for an override from the selected override set to be loaded, the item name and ancestry must match an item in the current project. When you select **Project>Import>Archived Overrides** the Import Archived Overrides dialog box appears.

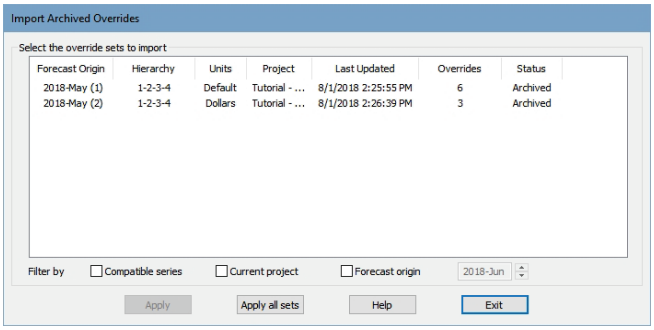

The Import Archived Overrides dialog box allows you to select the override sets to load. You can load an individual override set by highlighting the set and clicking the Apply button. Alternatively, you can load all displayed sets using the “Apply all sets” button. The “Filter by” options allow you to control which override sets to display. Selecting “Compatible series” will eliminate all override sets that do not match the current state (i.e., it will eliminate all override sets where the units of measure and/or the hierarchy structure does not match the current project’s display).

Project>Import>Calculated Rows From is used to import calculated rows from a saved project into the current project.

Project>Import>Bookmarks From is used to import bookmarks from a saved project into the current project.

Project>Export

The **Project>Export** menu allows you to save Numeric Output files and all Forecast Pro report files. Consult the appropriate section of the *Saving Your Work* chapter for details on designing the content and format of these output files.

Project>Set password

Project>Set password is used to password protect your project. When you attempt to open a password-protected project, Forecast Pro will prompt you for the password and only open the project if the correct password is entered.

Project>Control Panel

Project>Control Panel invokes TRAC Control Panel dialog box which is used to view information about the current project's database and to purge unwanted (e.g. obsolete) forecasts and overrides.

Important: To reduce the size of the Forecast Pro database file, you must use the Compact button after purging forecasts and overrides. Purging without compacting results in the records being deleted from the database but does not change the file size.

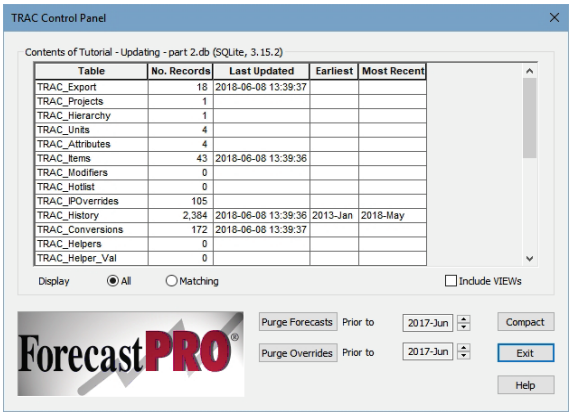

The grid displays the contents of the database.

The Purge Forecasts button will delete all forecasts prior to the specified date.

The Purge Overrides button will delete all overrides prior to the specified date.

The Compact button consolidates the Forecast Pro database file. Use this button to reduce the size of the database file after purging forecasts and/or overrides.

Project>Super Project

Project>Super Project invokes the Super Project Manager. The Super Project Manager is used to consolidate multiple projects into a "super project" or to update a super project.

Setting up and maintaining a super project is reasonably complex. If you are planning on using super projects you should (1) refer to Lesson 8 of the Tutorial section of this manual, *Collaborating With Colleagues* for a discussion of the pros and cons of super projects and an illustration of using the Super Project Manager to create a super project and (2) read the *Working With Super Projects* chapter of the Reference section of this manual carefully. Business Forecast Systems and Forecast Pro distributors also offer consulting services to help you establish your forecasting process.

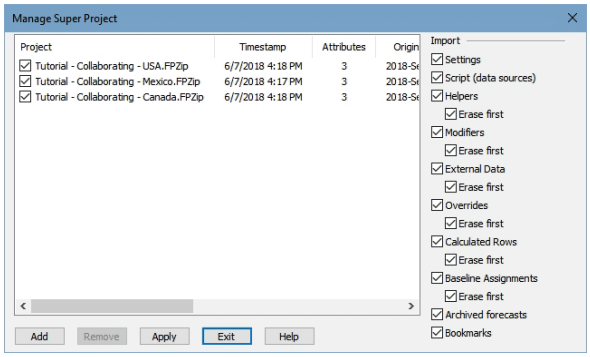

When creating a new super project, you first use the Add button to select the component projects to consolidate. The Import options allow you to specify what you wish to import from the component projects. When first creating a super project you will generally import everything. When updating a super project you may want to pick and choose what to import.

The “Erase first” options are only relevant if you are updating a super project. If selected, they will erase any data corresponding to the specified component prior to loading in the new data.

After you have selected the component projects and the import selections, the Apply button is used to create or update the super project.

The View Menu

The View menu allows you to view reports, display graphs and customize your display.

View>(All available views)

The first ten options on the View menu will toggle the selected view window on and off. This can also be accomplished using the appropriate green or yellow icon.

View>Toolbar

View>Toolbar displays and hides the Toolbar, which includes buttons for some of the most common commands in Forecast Pro. A check mark appears next to the menu item when the Toolbar is displayed.

View>Status Bar

View>Status Bar displays and hides the Status Bar, which describes the progress of the currently running forecast job, or the action to be executed by the selected menu item or depressed toolbar button. A check mark appears next to the menu item when the Status Bar is displayed.

View>Modifiers

View>Modifiers displays and hides the currently selected forecast modifiers on the Navigator.

View>Discontinued items

View>Discontinued items toggles on or off the display of discontinued items. If Discontinued items is not checked, the Navigator will not show discontinued items.

The Window Menu

The Window menu enables you to arrange your open windows in various ways.

Window>Cascade

Window>Cascade arranges all nonminimized windows in an overlapped fashion.

Window>Tile

Window>Tile arranges all nonminimized windows in a tiled (nonoverlapped) fashion.

The Help Menu

Forecast Pro includes extensive on-line help.

Help>Help Topics

Help>Help Topics displays an index of all of the help topics available. Clicking on an entry will jump to the selected topic.

Help>User Guide (PDF)

Help>User Guide (PDF) accesses a pdf version of the *Forecast Pro User's Guide*.

Help>Statistical Reference (PDF)

Help>Statistical Reference (PDF) accesses a pdf version of the *Forecast Pro Statistical Reference Manual*.

Help>Check for Updates

Help>Check for Updates will check online to see if you are running the latest release of the program. If a product update is available you'll receive a message alerting you to this fact with instructions describing how to receive the update. If your support and maintenance contract is up to date, product updates are free.

Help>About Forecast Pro

Help>About Forecast Pro displays the Forecast Pro version number and copyright notice.

Working with Super Projects

There are two primary ways to work with others to establish the final forecasts using Forecast Pro—sequential review of a single stand-alone project and consolidating multiple component projects into a "super project" using the Super Project Manager.

Sequential review of a single stand-alone project is by far the simplest way to work with others to establish the final forecasts. Using this approach the forecaster begins by creating forecasts for all items that need to be forecasted, enters any desired overrides and saves a forecast project—often to a network drive. The project can then be opened by anyone using Forecast Pro (Forecaster or Collaborator licenses) to review, make changes, etc. If there is a third individual who needs to review the project, that person would do so after the second individual has updated the project.

Alternatively, the super project approach allows you to break large forecasting jobs into smaller pieces that can be worked on separately and then consolidated. To implement this approach you first create, forecast and save separate component projects and then use the Super Project Manager to

combine the component projects into a new super project allowing you to view and manipulate the consolidated forecast.

Pros and cons

To illustrate the pros and cons of each approach, let's consider the following example. A company sells their products in three countries—the United States, Canada and Mexico. They employ three different demand planners, each of whom is responsible for forecasting a specific country. After the demand planners have completed their forecasts, they meet with management to review the forecasts and make any required adjustments.

Accomplishing this using a single stand-alone project would require a sequential review process. For example, assume the demand planner for the United States goes first. He or she would create a forecast project which includes all three countries and then review and make adjustments to the United States data and save the project. Next, the planner for Canada would open the project, review and adjust the Canadian data and save the revised project. Next, the planner for Mexico would open the project, review and adjust the Mexican data and save the revised project. At this point, the demand planners' forecasts are complete and the project is ready for the management review meeting.

Clearly, there are some downsides to the above procedure. One is that the demand planners cannot work on their country-level forecasts simultaneously—they have to wait until it is their turn to do so. Another problem is that even though they are working on a single country, the project contains data and forecasts for all three countries, which makes it harder to manage and more prone to error.

The advantage of a forecasting process that uses a sequential review of a single stand-alone project is that is *much* simpler to implement and maintain than one that uses a super project. If it is not essential that your total hierarchy be broken out into smaller pieces, then sequential review of stand-alone projects is the preferred approach.

Using the super project approach with our example, we'd create a separate forecasting project for each country and then consolidate them into a new super project which includes all three countries. The advantage is that the three demand planners can work on their forecasts simultaneously using a

project that only contains their country. When all three are done, the Super Project Manager is used to create a super project containing the consolidated forecast.

Creating a super project

The steps to create a super project are:

1. Create the data files for the component projects.

   The super project's historic data definition will include all of the component projects' data files. Therefore, you need to construct the data files for the component projects so that when they are combined they create the consolidated hierarchy correctly.

2. Select the number of override rows required and other settings.

   To import the component projects' overrides into the super project the number of override rows needs to be consistent across the component projects. To ensure a consistency between the statistical forecasts in the component projects and the super project, you will want to decide on all options relating to the forecasts (e.g., forecast horizon, confidence limits settings, outlier correction settings, etc.) up front and keep them consistent across the component projects.

3. Create the component projects

4. Use the Super Project Manager to create the super project

Prior to creating your first super project you should work through the hands-on section of Lesson 8 of the Tutorial section of this manual, *Collaborating With Colleagues* which walks you through the process of using the Super Project Manager to create a super project.

Updating a super project

How you update your component projects and super project will depend on the details of your forecasting process.

Many users do not modify the forecasts in the super project, but rather just use the super project as a way of viewing the consolidated forecast and saving output files. If this is the case, you may want to simply update the component projects and create a new super project every time new data becomes available.

If you are entering modifications (e.g., overrides and modifiers) in the super project then then the updating process may be a bit more involved. If necessary, you can import overrides and modifiers from the super project into the component projects using the Data Manager. Once the component projects are updated the Super Project Manager allows you to either delete the super project's overrides and/or modifiers prior to importing from the component projects or leave them intact.

As you can see, setting up and maintaining super projects is reasonably complex. Keep in mind that Business Forecast Systems and Forecast Pro distributors also offer consulting services to help you establish your forecasting process.

Using a Time Fence

Time fences are used to "lock" forecast values for a specific number of periods so that they can't be changed via user overrides or by the generation of new statistical forecasts as new data become available. Time fences can be useful in environments where after forecasts have been finalized and acted upon (i.e., production has been scheduled, orders have been placed, etc.) there is a window of time where making changes is not feasible.

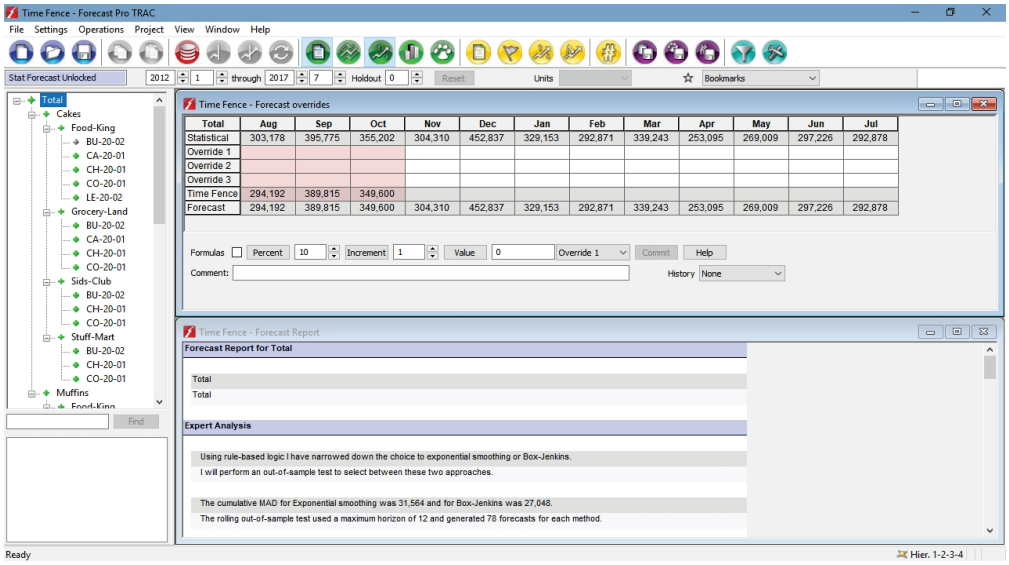

The screen above shows the Forecast overrides view for an item with a 3-month time fence. Notice the second to last row of the grid is labeled Time Fence. The values displayed in this row are the unchangeable fenced values and the same values appear in the final Forecast row.

Notice that the statistical forecast for August (303,178) does not equal the time fence value (294,192). This is because August's time fenced value was established when last month's forecasts were generated.

Let's now consider what happens next month. We update our data with August's sales, open last month's project and generate new statistical forecasts. Our 3-month time fence will now cover August through October. August and Septembers time fenced values will be retained from last month's time fenced values and November's time fenced value will equal last month's final forecast for November.

Time Fence Settings and Options

Turning on the time fence: When you first create a project, the time fence will be turned off. You turn the time fence on by selecting **Operations>Time Fence>Enable**. Once you turn a project's time fence on, *it cannot be turned off!*

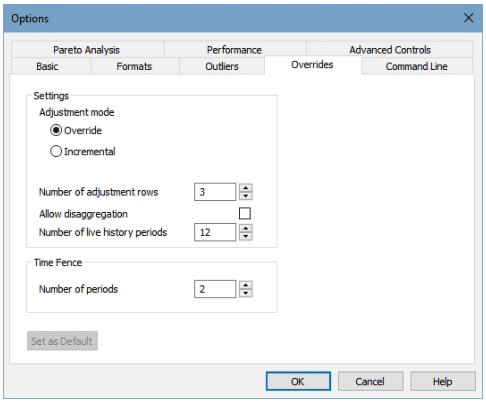

Setting the length of the time fence: The default length for the time fence can be specified on the Override tab of the Options dialog box (shown above).

You can override the global setting on an item-by-item basis using the custom modifier *\TIMEFENCE=n* where *n* is the number of periods you wish to use. Be aware that if you use the *\TIMEFENCE* modifier to create a time fence shorter than the global time fence, it will shorten the time fence for all associated parent groups of the item.

Freezing the next period: Selecting **Operations>Time Fence>Freeze Next Period** will increase the length of all time fences in the current project by one period for the current forecast origin. If you subsequently, update your data with the next observation and create a new forecast (thus changing the forecast origin) the time fence lengths will revert back to their pre-Freeze-Next-Period settings. This option provides a way of insuring that after the forecasts are finalized someone doesn't inadvertently alter the frozen periods' forecasts prior to the data being updated, new forecasts being generated and the frozen periods' forecast becoming an officially fenced value.

Command-line Operation

You can run Forecast Pro noninteractively. When the command line to execute the program includes a project filename (*.FPZip or .FPProj if the project is unzipped) and the /b parameter, Forecast Pro will read in the data, create the forecasts, save all output files and then exit. This feature is particularly useful when you are integrating forecasting with other software systems and need a "hands-off" approach.

To drive Forecast Pro entirely automatically, follow these two steps.

1.) Prepare a valid project and place it in your project directory. The easiest way to prepare this file is to use Forecast Pro interactively to define all desired project settings, historic data, output formats and (optionally) forecast modifiers and overrides, and then save the project using use **File>Save**.

Alternatively, you can generate the file externally but this will require understanding the project file and the (optional) corresponding db file formats.

2.) Drive Forecast Pro by issuing the command

"C:\Program Files\Forecast Pro v5\ForecastProTRAC.exe" Test.FPZip /b

where we assume that the program is located in the directory "C:\Program Files\Forecast Pro v5" and a project named Test.FPZip is located in your current project directory. The "/b" parameter is used to indicate that you wish to run in command-line (batch) mode.

If you want to apply the settings from a bookmark in the project, you may do so by adding "\BOOKMARK=Name" to the command shown above. For example, "\BOOKMARK=Bookmark01" will apply the settings from a bookmark named Bookmark01.

You can issue the command in any number of ways, including:

Select Run from the Start menu and issue the command from there.

Set up a shortcut to issue the command.

Issue the command from within a different application using a system call.

The Command Line tab of the **Settings>Options** dialog box allows you to specify several key command line forecasting options and the output files to save in command-line mode. This tab's functionality was fully described in the *Menu Items and Dialog Boxes* section earlier in this chapter.

Important Notes

1. If there are spaces in the pathname you must include quotation marks as illustrated in the example above.

2. If you do not want the Forecast Pro menu to appear, run the program minimized.

Chapter 8

Understanding How Forecast Overrides are Allocated

The override facility in Forecast Pro allows you to adjust your forecasts. Changes made at any level of the forecasting hierarchy will automatically reconcile all levels. This chapter was written to clarify exactly how the reconciliation is performed and document cases where adjustments cannot be fully implemented due to conflicting overrides or program settings.

The examples in this chapter assume that you are making adjustments to the statistical forecasts, however, it you have defined alternative baseline forecasts or generated an inherited forecast row the overrides will be made to these forecasts using the same logic.

Adjustments Made to a Single Level

In this section we will describe how Forecast Pro reconciles the forecast hierarchy when overrides are made to a single level of the hierarchy.

Important: *If you restrict your overrides to a single level of the hierarchy, the reconciliation process is simple and intuitive. We strongly urge you to limit your overrides to a single level if at all possible.*

Figure 1 below depicts a simple forecast hierarchy prior to any overrides being entered. Notice that at this point the Statistical forecasts equal the committed Forecasts.

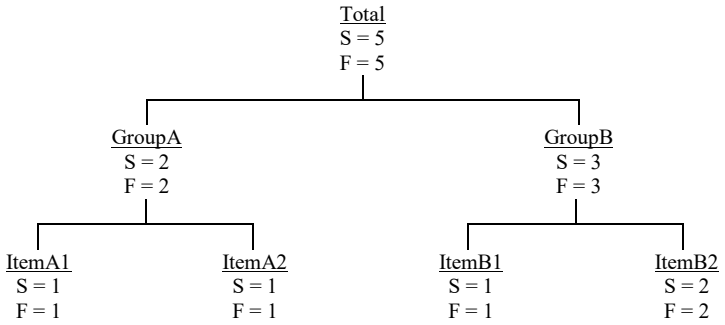

Figure 1.

Let's now modify the above hierarchy by overriding the forecast for ItemA1 to equal 75. This results in the hierarchy shown in Figure 2 below.

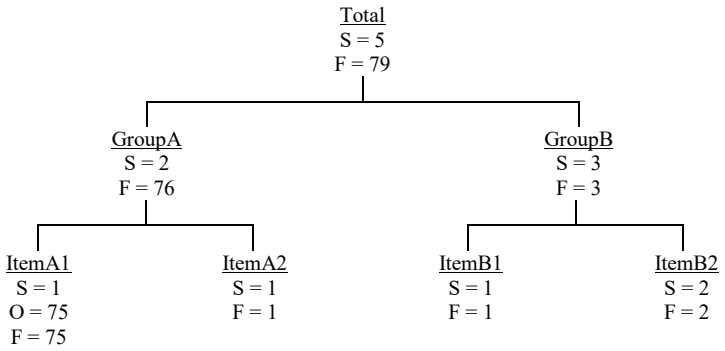

Figure 2.

Notice that at this point the committed Forecasts for GroupA and Total do not equal the Statistical forecasts.

If we had started with the hierarchy depicted in Figure 1 and adjusted the GroupB forecast to 75 (rather than ItemA1) we would generate Figure 3.

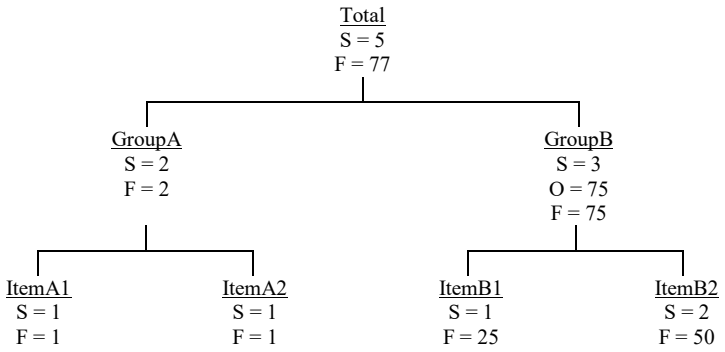

Figure 3.

Notice that the committed Forecasts for ItemB1, ItemB2 and Total do not equal the Statistical forecasts. Notice also that Forecast Pro allocated the GroupB override to ItemB1 and ItemB2 based on the proportions established by their Statistical forecasts.

Adjustments Made to More Than One Level

In this section we will describe how Forecast Pro reconciles the forecast hierarchy when overrides are made to more than one level of the hierarchy.

Important: *There are many ways that this type of reconciliation could be performed—each one leading to different forecasts. If you plan to make overrides at more than one level of the hierarchy, it is essential that you understand how Forecast Pro performs the reconciliation.*

Figure 4 below depicts a forecast hierarchy where ItemA1 has been overridden to equal 75. (It is identical to Figure 2.)

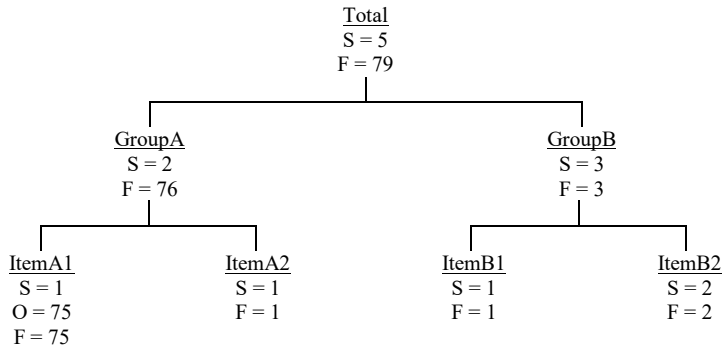

Figure 4.

Let's now modify the above hierarchy by overriding the forecast for Total to equal 475. This results in the hierarchy shown in Figure 5 below.

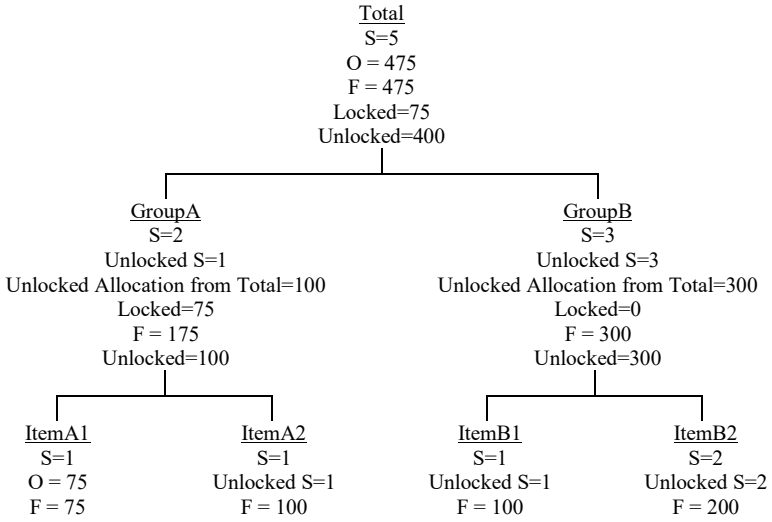

Figure 5.

Making adjustments at more than one level introduces the concept of *locked* and *unlocked* volume and forecasts. Proportional allocations of unlocked volume are made using unlocked forecasts.

Notice that:

1.) When we adjusted Total to 475, 75 units were locked due to the override on ItemA1 and 400 units were unlocked and needed to be allocated downward.

2.) Forecast Pro used the *Unlocked* statistical forecasts for GroupA and GroupB to determine the allocation proportions for the 400 unlocked units from Total.

3.) Forecast Pro did not alter the committed Forecast value for ItemA1 because it had been explicitly overridden (i.e., it was locked). Therefore, the 100 units of unlocked volume from GroupA was allocated exclusively to ItemA2 (the only unlocked item in GroupA).

4) Forecast Pro used the *Unlocked* statistical forecasts for ItemB1 and ItemB2 to determine the allocation proportions for the 300 unlocked units from GroupB.

Let's now modify the above hierarchy by overriding the forecast for Item A2 to equal 75. This results in the hierarchy shown in Figure 6 below.

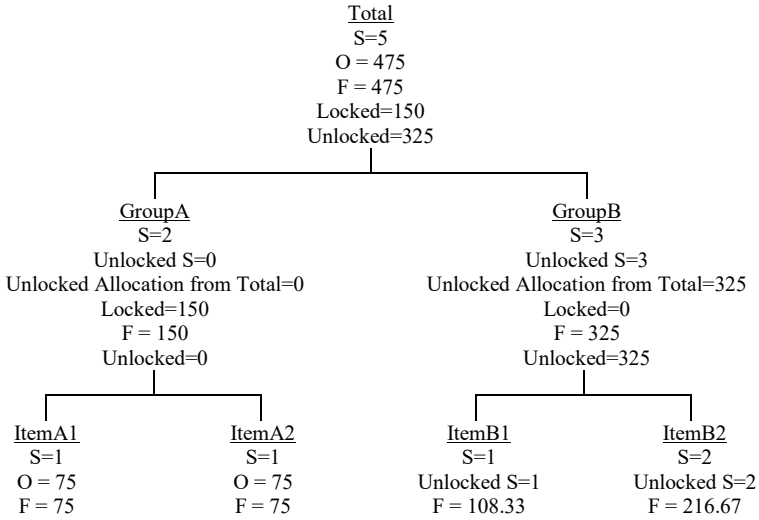

Figure 6.

Notice that GroupA has no unlocked volume so the 325 units of unlocked volume from Total are allocated exclusively to GroupB.

Examples Where Overrides Cannot be Allocated

If you make forecast overrides to more than one level of the hierarchy you can generate situations where the overrides cannot be allocated. In these instances an error message will be displayed and the override will not be applied.

Figure 7 below depicts a forecast hierarchy where ItemA1 and ItemA2 have both been overridden to equal 75 and ItemB1 has been overridden to equal 150.

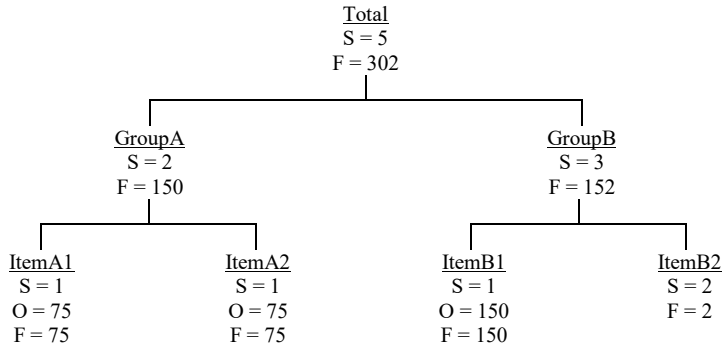

Figure 7.

Suppose you try to override the forecast for GroupA to equal 100. The override could not be applied because both ItemA1 and ItemA2 are “locked”. Forecast Pro will display an error message and reject the override.

Let's now modify the above hierarchy by overriding the forecast for GroupB to equal 50. This results in the hierarchy shown in Figure 8 below.

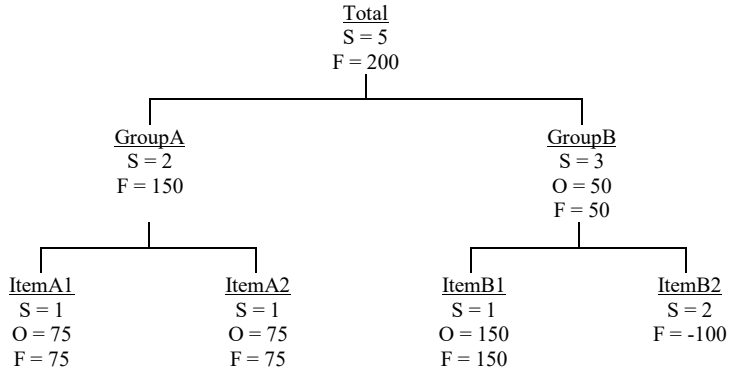

Figure 8.

Notice that the override for GroupB resulted in a negative committed Forecast for ItemB2. This result assumes that the "Allow Negative Forecasts" setting in the Forecast Pro Options dialog box was selected. If this setting was not selected, the override could not be made and Forecast Pro would display an error message and reject the override.

Summation

It should be quite apparent that if you make adjustments to more than one level of a hierarchy the reconciliation becomes quite complex. Although several of the examples presented seemed involved, they were actually quite simple hierarchies consisting of only 3 levels and 4 end-items. Most users of Forecast Pro are working with thousands of items and much more complex hierarchies.

In developing the reconciliation routine for Forecast Pro the authors considered using several different algorithms. In many situations the choice of which algorithm to use would change the forecast values—sometimes significantly. None of the algorithms were "statistically superior" to the others—they were simply different.

Virtually all of the complexity (and differences between algorithms) disappears if you restrict yourself to making forecast adjustments to a single level of the hierarchy. Although we've stated this previously, it's worth repeating:

If you restrict your overrides to a single level of the hierarchy, the reconciliation process is simple and intuitive. We strongly urge you to limit your overrides to a single level if at all possible.

Appendixes

Speeding Up the Program

If decreasing processing time is important for your application, you may wish to investigate some of the following options.

1. Use the 64-bit Version

Forecast Pro can be installed as either a 32-bit or 64-bit application. If you are currently running the 32-bit version, migrating to the 64-bit version will speed up various operations such as opening and saving projects and will also allow you to work with larger data sets.

2. Work Locally

Working with projects stored on a network drive can be substantially slower than working on a project stored locally. Reading input data over the network and saving output files to network drives can also be substantially slower than reading and writing from a local drive.

If you are currently saving your projects to a network drive, saving them to a local drive will improve performance considerably. If you are currently reading your input data from a network drive and/or saving output files to a

network drive, moving them to a local drive will speed up read/write operations considerably.

3. Use Zipped Projects

As we mentioned above, in some networked environments, working with projects stored on a network drive can be substantially slower than working on a project stored locally. To improve performance in these cases, Forecast Pro supports “zipped” projects. When you save a zipped project, Forecast Pro zips the four standard project files described in a previous section of this manual into a single file with the extension .FPZip.

When you open a zipped project, Forecast Pro copies the zipped file from the network drive to a local workspace, unzips the project files and opens the project—allowing you to work on your project locally, even though it was saved to a network drive. When you resave your zipped project, Forecast Pro will zip the updated local project files and save the updated .FPZip file to the network drive.

To save and open zipped projects, you select FPZip as the file type in the **File>Save as** and **File>Open** dialog boxes.

4. Upgrade Your Hardware

Upgrading your hardware will improve performance dramatically. Forecast Pro will run on any Intel processor from the 386 on. Since this product deals with large amounts of floating point data, it makes sense to run it on as fast a machine as possible.

As an illustration, using expert selection and four years of monthly history per item, a Core2/2.2GHz machine forecasted 10,000 items in 39 seconds. The same run took 4 minutes on a Pentium III/600MHz machine and 24 minutes on a Pentium/90MHz.

Having enough RAM is also important—we recommend 2 Gigs if you are running a 32-bit operating system and as much RAM as your budget can afford if you are running a 64-bit operating system.

As you use the program, Forecast Pro is storing and retrieving a great deal of information in memory. If you do not have an adequate amount of RAM

Forecast Pro will resort to using virtual memory (i.e., caching to your hard drive). This is substantially slower than using RAM. Increasing the amount of RAM in your computer will improve performance of all of your applications.

5. Use Manual Override Mode

If you make overrides to large complex hierarchies, the reconciliation can take a little time. To minimize the reconciliation time, Forecast Pro supports a manual override mode. When this mode is active, the override window will include a Retain button as well as a Commit Button. The Retain button is used to accept overrides without reconciling the hierarchy. In this mode you would typically enter and retain all of the overrides you wish to make (i.e., for multiple items) and then click Commit or the red Recalculate icon ( ) to perform the reconciliation. The control to turn on manual override mode is found on the Performance tab of the **Settings>Options** dialog box.

6. Use Manual Forecast Mode

By default, Forecast Pro will automatically recalculate the forecasts and re-reconcile the hierarchy every time you apply a forecast modifier. For large hierarchies this can be time consuming. Forecast Pro allows you turn on manual recalculation using an option of the Performance tab of the **Settings>Options** menu. If you turn this option off, when you specify a forecast modifier the forecasts will not be updated until you click the red Forecast icon ( ) or the red Recalculate icon ( ). This allows you to specify multiple forecast modifiers without having to wait for the program to recalculate the forecasts in between each specification.

7. Reduce the Size of the Forecast Archive

By default, Forecast Pro will archive the last 12 forecasts you've created in the project's database. Forecast Pro allows you to limit the number of periods archived using the *Periods to retain in archive* spinner on the Performance tab of the **Settings>Options** menu. Reducing the number of archived forecasts will result in a smaller project and better performance for many program operations (e.g., opening and saving projects, loading the tracking report, etc.).

8. Optimize Outlier Detection and Correction for Speed

The outlier detection and correction algorithm will slow down the processing. If you wish to use outlier detection and correction but speed is of concern, one option is to initially run the forecasts in the detection only mode to detect the outliers and then turn outlier detection off and use the \OUTLIER=CORRECT modifier to perform the corrections on the detected items.

9. Don't Use .xlsx Files

If you are using xlsx files you might consider switching to csv or xls files. Xlsx files are slower to read than csv and xls files and use over twice as much memory. Also, if you are using large spreadsheet files breaking them up into several smaller files or increasing the amount of RAM available will speed up processing.

10. Optimize Model Selection for Speed

Specifying "Exclude ARIMA" on the Performance tab of the **Settings>Options** dialog box will speed up forecast generation considerably. This option will remove Box-Jenkins models from consideration when using expert selection.

If you are forecasting data where the number of periods per cycle is greater than 13 (e.g., weekly data) you might consider turning on the "Exclude seasonal simplification" option in the Performance tab of the **Settings>Options** dialog box. If this switch is on, Forecast Pro's expert selection mode will not consider seasonally simplified forms of exponential smoothing models. This will speed up forecast generation considerably.

Specifying the models to use rather than using the expert selection will substantially improve performance. This will require some work to determine the best type of model to use for your data. The best way to determine the appropriate model for your data set is to experiment using the forecast evaluation procedures (see methodology chapter in the *Forecast Pro Statistical Reference Manual* for details).

End-User License Agreement

This End-User License Agreement ("EULA") is a legal agreement between Licensee (either an individual or a single entity) and Business Forecast Systems, Inc. ("BFS") for the BFS software that accompanies this EULA.

1. DEFINITIONS

A. "BFS" means Business Forecast Systems, Inc., 68 Leonard Street, Belmont, MA 02478 USA, the author and owner of the copyright of this computer program product.

B. "Computer" means the single computer, workstation, terminal or other digital electronic device on which You use this program.

C. "Software" means the set of computer software programs in this package, regardless of the form in which You subsequently use it.

D. "Licensee" means an individual or single entity exercising rights under and complying with all the terms of this EULA. In the remainder of this EULA, the Licensee is referred to as "You."

E. "Authorized User" means the Licensee in the case where the Licensee is an individual. In the case where the Licensee is an entity, "Authorized User" means (i) an employee of the Licensee, but only while such employee is acting within the scope of such employee's employment by Licensee; and (ii) the primary user of the Computer on which the Software is installed.

If You do not agree with the terms of this EULA, BFS does not grant any license to the Software and You should not download, install or use the Software. In such event that You do not agree with the terms of this EULA and have already installed or downloaded the Software, You must remove the Software from your Computer and return to BFS the complete software package, including disks and printed materials and destroy all copies (or, if You downloaded the Software, notify BFS in writing that You have removed the Software from your Computer and destroyed all copies). If You do not comply with these return and notification requirements within thirty days (30) of your receipt of the Software, You will not be entitled to a refund of any license fees You paid.

2. OWNERSHIP

BFS retains title, ownership and all rights and interests in and to the Software, documentation and all other materials supplied by BFS. BFS does not sell the Software or any copies thereof, but only grants limited licenses to use the Software in accordance with the terms of this EULA.

3. GRANT OF LICENSE

Upon acceptance of this EULA, BFS grants You a non-transferable, non-exclusive, limited right to use this copy of the Software. You may use the Software and the associated documentation only in connection with one Computer. You may physically move the Software from one Computer to another, provided that the Software is used on only one Computer at a time. You may not move the Software electronically from one Computer to another over a network. This license is granted for use by an individual Authorized User, and, in no event shall there be more than one Authorized User on the Computer where the Software is installed; e.g., there shall be no common sharing of a "Forecast Pro computer." Each Authorized User must have a separate license to use the Software on his or her Computer.

4. LICENSE RESTRICTIONS

You acknowledge that the Software and its structure, organization and source code constitute valuable trade secrets of BFS. You shall not, nor shall permit, assist or encourage any third party to:

- copy, modify, adapt, alter, translate, reverse engineer, decompile, disassemble, or attempt to derive the source code for the Software.
- merge the Software with other software and/or create derivative works from the Software.
- sublicense, distribute, publish, sell, lease, rent, loan, or otherwise transfer or allow access to the Software to any third party.
- use the Software in the operation of a service bureau.
- defeat, disable or circumvent any security measures BFS may use to prevent unauthorized use or copying of the Software.

5. TRANSFER OR REPRODUCTION

You may not transfer or reproduce the Software with the exception of making a maximum of two (2) copies of the Software to be used for backup purposes only. When using a Computer with a hard disk, one copy of the Software may be installed on the hard disk provided that the copy is removed if the Software is used on another Computer.

6. LIMITED WARRANTY

THE SOFTWARE, THE DOCUMENTATION AND ANY OTHER MATERIALS THAT MAY BY BE PROVIDED BY BFS ARE PROVIDED "AS IS" WITHOUT ANY WARRANTY OF ANY KIND. BFS HEREBY EXPRESSLY DISCLAIMS ALL WARRANTIES WITH REGARD TO THE SOFTWARE, THE DOCUMENTATION AND ANY OTHER MATERIALS, WHETHER EXPRESS, IMPLIED, STATUTORY OR OTHERWISE, INCLUDING WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND ANY WARRANTIES ARISING OUT OF COURSE OF DEALING OR COURSE OF PERFORMANCE. BFS EXPLICITLY DOES NOT WARRANT THAT THE SOFTWARE WILL MEET ALL OF YOUR REQUIREMENTS OR THAT THE OPERATION OF THE SOFTWARE WILL BE ERROR-FREE. YOU ACKNOWLEDGE THAT YOU HAVE RELIED ON NO WARRANTIES AND NO WARRANTIES ARE MADE HEREIN BY BFS.

7. LIMITATION OF LIABILITY

BFS SHALL NOT BE LIABLE FOR ANY LOSS OR DAMAGE THAT MAY BE RELATED TO THE FURNISHING, PERFORMANCE, OR USE OF THE SOFTWARE, INCLUDING WITHOUT LIMITATION ANY INDIRECT, SPECIAL, INCIDENTAL, ECONOMIC, LOST PROFIT, OR CONSEQUENTIAL LOSS OR DAMAGE OF ANY KIND WHATSOEVER EVEN IF BFS HAS BEEN ADVISED OF THE POSSIBILITY OF THOSE DAMAGES. IN NO EVENT SHALL BFS’S LIABILITY UNDER THIS EULA EXCEED THE AMOUNT PAID FOR THE SOFTWARE LICENSE, IF ANY.

8. TERM AND TERMINATION

This EULA is effective until terminated. Your rights under this EULA will terminate automatically without notice from BFS if You fail to comply with any provision of this EULA. Upon termination of this EULA, You must destroy all copies, full or partial, of the Software. You agree that upon termination of this EULA for any reason, BFS may take actions so that Software no longer operates.

9. GENERAL TERMS AND CONDITIONS

A. You acknowledge that You have read this EULA understand it and agree to be bound by its terms and conditions. This EULA supersedes all prior agreements, oral or written communications, proposals, representations and warranties and prevails over any conflicting or additional terms of any quote, order, acknowledgment, or other communication between the parties relating to its subject matter during the term of this EULA. No modification of this EULA will be binding, unless made in writing and signed by an authorized representative of each party.

B. In the event that any portion of this EULA is held unenforceable, the unenforceable portion shall be construed in accordance with applicable law as nearly as possible to reflect the original intentions of the parties, and the remainder of this EULA shall remain in full force and effect.

C. This EULA will be governed by and construed and enforced according to the laws of the Commonwealth of Massachusetts, USA excluding its conflict of law provisions.

Suggested Reading

The *Forecast Pro Statistical Reference Manual* is a pdf document that is copied to the program directory when you install Forecast Pro. It is also accessible via the Forecast Pro help menu and is the primary suggested reference for all of the forecasting techniques, statistics and algorithms found in Forecast Pro.

The *Forecast Pro Statistical Reference Manual* includes an extensive bibliography. A few recommended textbooks include:

J. S. Armstrong [2001] Principles of Forecasting: A Handbook for Researchers and Practitioners, Norwell MA: Kluwer Academic Publishers.

K. Kahn [2006] New Product Forecasting: An Applied Approach, Armonk NY: M.E. Sharpe.

K. Ord and R. Fildes [2013] Principles of Business Forecasting, Cengage.

S. Makridakis, S. C. Wheelwright and R.J. Hyndman [1998] Forecasting Methods and Applications, Third Edition, New York: Wiley.

P. Newbold and T. Bos [1990] Introductory Business Forecasting, Cincinnati: South-Western.

S. Makridakis, S. C. Wheelwright and R.J. Hyndman [1998] Forecasting Methods and Applications, Third Edition, New York: Wiley.

Index