forecast pro ®

forecastpro.com

BUSINESS FORECAST SYSTEMS

Better Decisions

Demand

Forecast Accuracy

Unlimited V12

Forecast Pro Unlimited Reference

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

Software Version 12
Manual Last Revised: August 25, 2021

Contents

Introduction to Forecast Pro

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 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.

What You Need to Run Forecast Pro

A computer running a Windows operating system.

A minimum of 512 MB of RAM (random access memory). 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.

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 BFS product development (support@forecastpro.com or call 1-617-484-5050).

Forecast Pro Resources

Business Forecast Systems, Inc. provides the following resources to help you get the most out of Forecast Pro:

The *Forecast Pro Tutorials Manual* provides a great way to introduce yourself to the product and its capabilities. The tutorials use sample input files and Forecast Pro projects that are installed with Forecast Pro. This manual is accessible via the Start page, the File page and the Help Ribbon menu.

The *Quick Start* video provides a quick video overview of how to use Forecast Pro. This video is accessible via the Start page and the File page.

The *Forecast Pro Reference Manual* documents all functionality included in the software. This manual is accessible via the Start page, the File page and the Help Ribbon menu.

The *Forecast Pro Statistical Reference Manual* describes the statistical techniques, statistics and strategies that are implemented in Forecast Pro. This manual is accessible via the Help Ribbon menu.

The *Resources Menu* on www.forecastpro.com provides an array of Forecast Pro resources including free on-demand forecasting webinars, access to the Forecast Pro blog and a calendar of upcoming Forecast Pro events.

Our *Technical Support Team* (support@forecastpro.com or call 1-617-484-5050) strives to provide timely responses to issues or questions you may encounter while using Forecast Pro.

BFS provides on-site and web-based training services. Please contact our support team if you are interested in training.

Installing Forecast Pro

Forecast Pro can be installed as a native 64-bit or native 32-bit application. Most users will want to install the 64-bit version to optimize performance. We recommend installing the 32-bit version only if you are running a 32-bit operating system.

Both the 64-bit and 32-bit installation downloads are compressed MSI files that can can be extracted using an unzipping application such as PkZip or WinZip. The screenshots below were created using Windows File Explorer.

1. Download the ZIP file by clicking the appropriate link. Choose a location to Save the file by choosing the Save As option and/or making a note of the location where it is being saved.
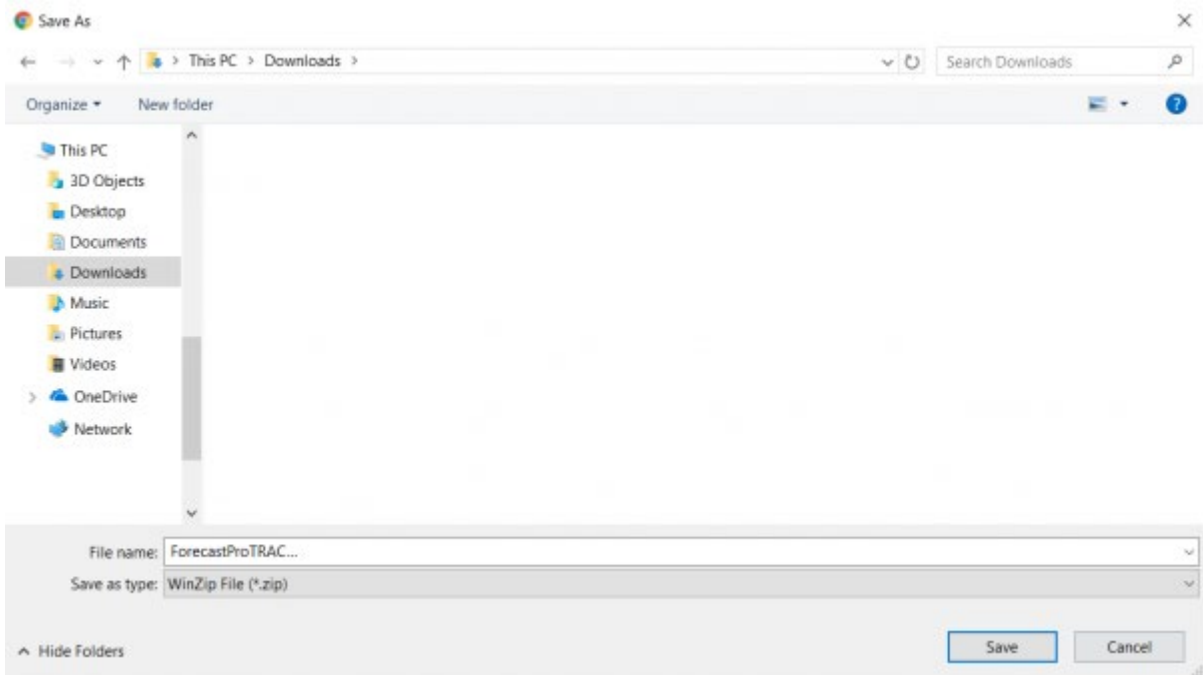
2. After the download is complete, go to the folder that contains the downloaded file and extract the files from the zip archive by selecting the Extract All option. Please do not run the installation program until it has been unzipped.
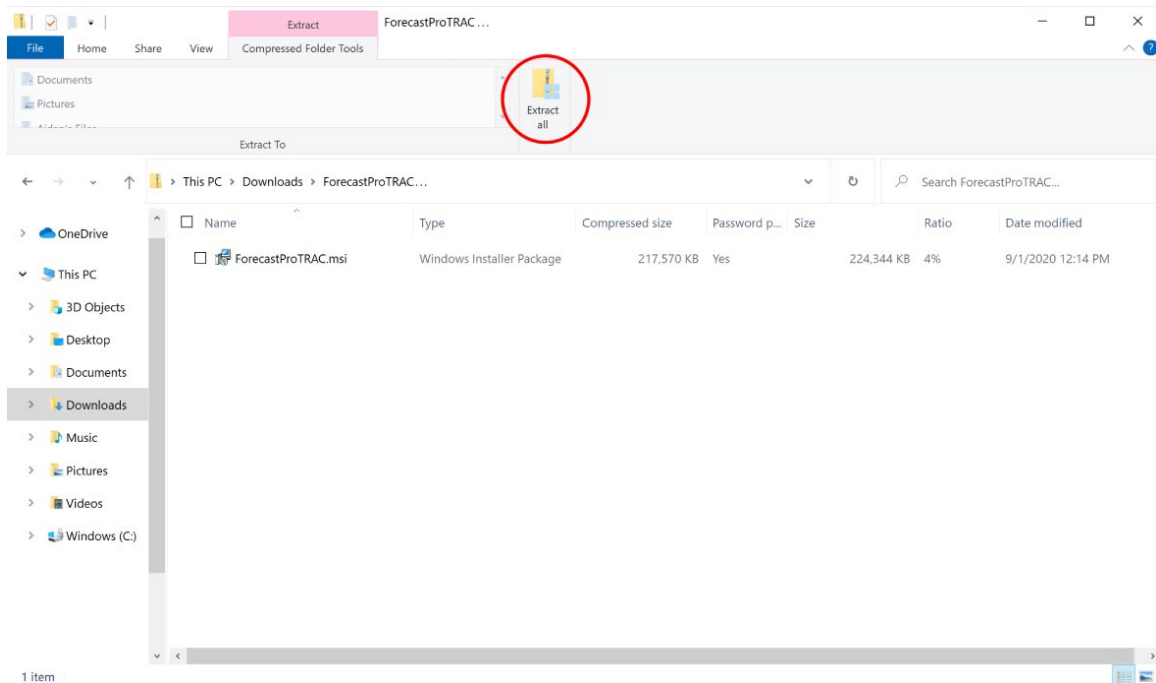
3. After selecting *Extract All*, you will be prompted for a password. Enter the password that was provided to you via email and click *Next*. Make a note of the name and location of the folder that is created.

4. Extracting the files will create an unzipped folder. Navigate to the unzipped folder to view the extracted files. Please make sure that the folder you are in looks like the one in the screenshot below and not look like the one above, and then run the MSI installation program the unzipped folder.

5. Follow the on-screen instructions of the installation program. When prompted for a Serial Number or CD-Key, enter the numbers that were provided to you via email.

6. You will be prompted to install SQLite drivers during installation. This is a required step to run Forecast Pro, so do not skip this step.

7. If you have any questions or difficulties, please do not hesitate to reach out via email to support@forecastpro.com or by phone at +1 617 484 5050 (US) or +44 (0) 20 8132 6333 (UK).

Running the installation program will 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.

The installation program will also install files to run the Forecast Pro Excel Add-in. The Forecast Pro Excel Add-in is activated the first time you start Forecast Pro. If you are planning to use the Forecast Pro Excel Add-in but are not planning on using the interactive version of Forecast Pro, you will still need to start the interactive version of Forecast Pro prior to using the Forecast Pro Add-in for the first time.
Use the Add or Remove Programs utility in the Windows Control Panel to remove Forecast Pro from your computer.

A Quick Overview of Forecasting

What is Statistical Forecasting?

Everybody forecasts, whether they know it or not. Businesses need to forecast future events 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. The latest version of Forecast Pro also includes extreme gradient boosting, a machine learning methodology that has recently emerged as a very effective for demand forecasting.

All 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.

Six 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 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 may not perform well when the data are unstable.

Machine learning is available as a univariate model and may be included in the expert selection algorithm, if desired. Like exponential smoothing and Box-Jenkins, machine learning univariate models leverage information inherent in the historic data (seasonality, trend) to create features. Using XGBoost, an extreme gradient boosting algorithm, these features are used to create an ensemble of decision trees to generate forecasts. Forecast Pro's automatic machine learning model leverages our AI-driven *expert selection* algorithm to select features and an extreme gradient boosting tree model specification appropriate for your data.

Forecast Pro's *expert selection* automatically chooses the appropriate univariate forecasting technique for each item forecasted. By default, expert selection will determine which of the first five univariate models is most appropriate for the data, optimize the model and create the forecasts. If desired, you may include machine learning in expert selection as well. Alternatively, you can dictate that a specific method be used and customize your models. Forecast Pro provides extensive diagnostics and statistical tests to help you make informed decisions.

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

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 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.

The *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, and in those instances the forecasts may not be as accurate as those from univariate methods.

Custom Machine learning models may also include event schedules and explanatory variables as features in a boosted decision tree ensemble model. As with event models and dynamic regression, you must provide historic values for events and explanatory variables.

Like Forecast Pro's automatic machine learning models, our custom machine learning methodology uses XGBoost, an extreme gradient boosting algorithm, to create an ensemble of decision trees. Forecast Pro's machine learning algorithm will automatically select the most relevant features from both the user provided ones and the Forecast Pro generated features. You can automatically train custom machine learning models using the same AI-driven logic that powers Forecast Pro's automatic machine learning univariate models, or you can specify a specific boosted decision tree structure.

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, so 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.

Setting Up Your Historic Data

Forecast Pro requires historic data to be input in a specific 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. Finally, you can map history from one or more items to your short data history item to create a longer history to use for forecasting.

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. This is controlled by several settings on the Formats tab of the Options (see below).

The encoding of missing values depends upon the data format. However, one problem is common to all 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 **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 code's 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 **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. To change treatment of missing data, you must reread the data 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 **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.

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. The year and period corresponding to the first available data point. If the first available observation for a monthly series was April 2016, then the starting year would be 2016 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. If someone has an xlsx file open in Excel, you cannot read it into Forecast Pro. 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. Reading data from a view is slower than reading from a table.

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 examples 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*.

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 that cell contains data or not. Thus, in the example spreadsheet, the first value for PG1-204-2 is for January 2017 in column I, even though the column C defines the starting year as 2016, 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*.

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. In the example spreadsheet, the first value for PG1-204-2 is for January 2017 in row 9, even though the header defines the starting period as 11 and the starting year as 2016.

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*.

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 to define product groupings within a hierarchy, 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.
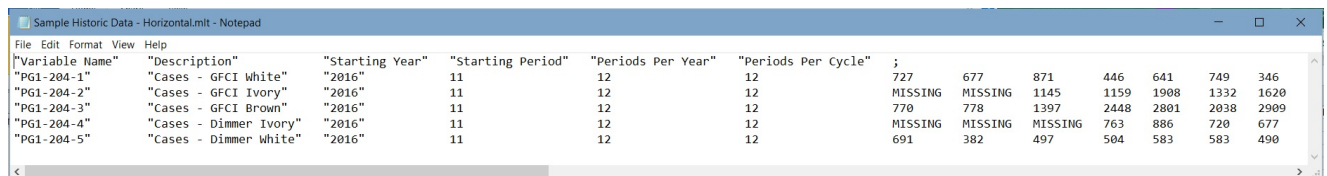
This file is stored as an editable text file with the extension MLT. You must create this file outside of Forecast Pro, making sure that you format it appropriately.

The MLT file format is straightforward. The file begins with record headings followed by a semicolon. The record headings are required for the attribute (e.g. 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 *Sample Historic Data - Horizontal.mlt*.

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 2017, even though the header defines the starting period as 11 and the starting year as 2016. Aligning 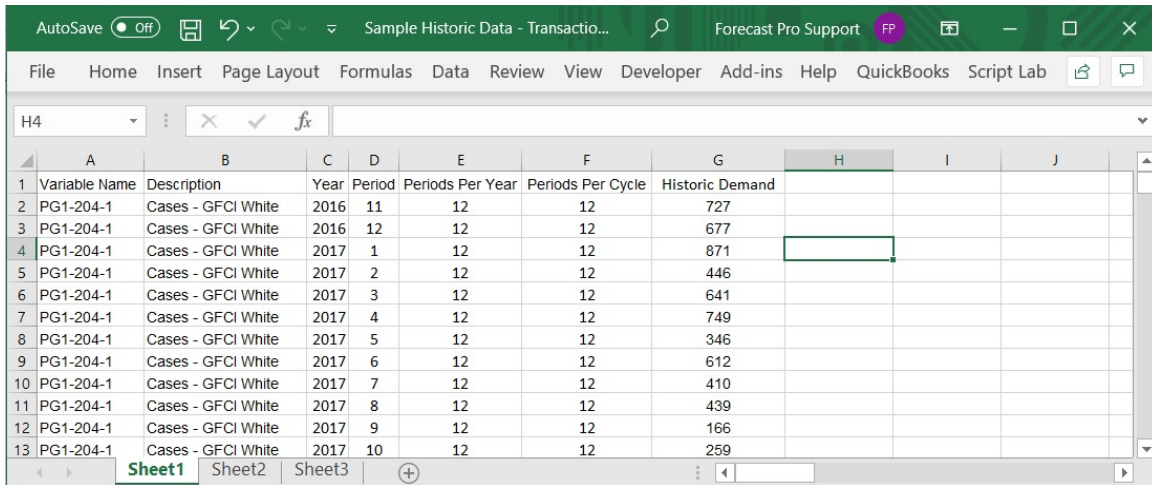the data file this way makes it easier for somebody to view and modify the file. Alternatively, we could have omitted the "MISSING" keywords and altered the starting period and starting year to correspond to each individual time series' starting date. Forecast Pro will have no problem reading a file 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 to define product groupings within a hierarchy, 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 ensure 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 and install the driver from the database manufacturer or a third party.

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*.

ItemId0	Description	Hist_Year	Hist_Period	PPY	PPC	Hist_Value
COR-12-11	CASE - 12 count Corn Muffins	2016	6	12	12	1944
COR-12-11	CASE - 12 count Corn Muffins	2016	7	12	12	3888
COR-12-11	CASE - 12 count Corn Muffins	2016	8	12	12	5832
COR-12-11	CASE - 12 count Corn Muffins	2016	9	12	12	8748
COR-12-11	CASE - 12 count Corn Muffins	2016	10	12	12	16524
COR-12-11	CASE - 12 count Corn Muffins	2016	11	12	12	16524
COR-12-11	CASE - 12 count Corn Muffins	2016	12	12	12	23328
COR-12-11	CASE - 12 count Corn Muffins	2017	1	12	12	17496
COR-12-11	CASE - 12 count Corn Muffins	2017	2	12	12	2916
COR-12-11	CASE - 12 count Corn Muffins	2017	3	12	12	6804
COR-12-11	CASE - 12 count Corn Muffins	2017	4	12	12	17496
COR-12-11	CASE - 12 count Corn Muffins	2017	5	12	12	1944
COR-12-11	CASE - 12 count Corn Muffins	2017	6	12	12	9720
COR-12-11	CASE - 12 count Corn Muffins	2017	7	12	12	3888
COR-12-11	CASE - 12 count Corn Muffins	2017	8	12	12	5832

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 to define product groupings within a hierarchy, 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.

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.

ItemId0	ItemId1	ItemId2	ItemId3	Description	Hist_Year	Hist_Period	Ppy	Ppc	Hist_Value
Total	Cakes	Food-King	BU-20-02	CASE - 20 oz. Butter Pound Cake	2016	6	12	12	1944
Total	Cakes	Food-King	BU-20-02	CASE - 20 oz. Butter Pound Cake	2016	7	12	12	3888
Total	Cakes	Food-King	BU-20-02	CASE - 20 oz. Butter Pound Cake	2016	8	12	12	5832
Total	Cakes	Food-King	BU-20-02	CASE - 20 oz. Butter Pound Cake	2016	9	12	12	8748
Total	Cakes	Food-King	BU-20-02	CASE - 20 oz. Butter Pound Cake	2016	10	12	12	16524
Total	Cakes	Food-King	BU-20-02	CASE - 20 oz. Butter Pound Cake	2016	11	12	12	16524
Total	Cakes	Food-King	BU-20-02	CASE - 20 oz. Butter Pound Cake	2016	12	12	12	23328
Total	Cakes	Food-King	BU-20-02	CASE - 20 oz. Butter Pound Cake	2017	1	12	12	17496
Total	Cakes	Food-King	BU-20-02	CASE - 20 oz. Butter Pound Cake	2017	2	12	12	2916
Total	Cakes	Food-King	BU-20-02	CASE - 20 oz. Butter Pound Cake	2017	3	12	12	6804
Total	Cakes	Food-King	BU-20-02	CASE - 20 oz. Butter Pound Cake	2017	4	12	12	17496
Total	Cakes	Food-King	BU-20-02	CASE - 20 oz. Butter Pound Cake	2017	5	12	12	1944

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.

Setting Up Your Optional 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:

Explanatory Variables. Explanatory variables files contain variables that you wish to use in conjunction with dynamic regression models.

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.

Helpers. 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.

Explanatory Variable Files

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 can use the "automatically extend" option to instruct Forecast Pro to forecast explanatory variables (using expert selection) 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.

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).

The sample file shown below, "*Cereal- Item Level Explanatory Variables.xlsx*", is provided with the software.

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 Variables - Transactions.xlsx*", is provided with the software and shown below.

Notice that the attribute fields are in the first four 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 Allocation Files

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*.

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*.

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.

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.*

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 as

well as number overrides. Formulas must begin with an equal sign (if using Excel, use an apostrophe before the equals sign to prevent Excel from executing the formula).

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.

ItemId0	ItemId1	ItemId2	ItemId3	FC_Period	FC_Year	Override_Row	Override_Formula	Override_Created	Override_Comment
Total	Cakes	Grocery-Land	BU-20-02	8	2020	2	=STAT*.8	19-May-19	G-Land to discontinue cakes at 40% of store locations
Total	Cakes	Grocery-Land	BU-20-02	9	2020	2	=STAT*.7	19-May-19	G-Land to discontinue cakes at 40% of store locations
Total	Cakes	Grocery-Land	BU-20-02	10	2020	2	=STAT*.6	19-May-19	G-Land to discontinue cakes at 40% of store locations
Total	Cakes	Grocery-Land	BU-20-02	11	2020	2	=STAT*.6	19-May-19	G-Land to discontinue cakes at 40% of store locations
Total	Cakes	Grocery-Land	BU-20-02	12	2020	2	=STAT*.6	19-May-19	G-Land to discontinue cakes at 40% of store locations
Total	Cakes	Grocery-Land	BU-20-02	1	2021	2	=STAT*.6		G-Land to discontinue cakes at 40% of store locations
Total	Cakes	Grocery-Land	BU-20-02	2	2021	2	=STAT*.6		G-Land to discontinue cakes at 40% of store locations
Total	Cakes	Grocery-Land	BU-20-02	3	2021	2	=STAT*.6		G-Land to discontinue cakes at 40% of store locations
Total	Cakes	Grocery-Land	BU-20-02	4	2021	2	=STAT*.6		G-Land to discontinue cakes at 40% of store locations
Total	Cakes	Grocery-Land	BU-20-02	5	2021	2	=STAT*.6		G-Land to discontinue cakes at 40% of store locations

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. 2019-July, or 2019-52.

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. In addition to creating helper variables interactively for event models and forecasting by analogy, Forecast Pro allows you to import helper variables from other Forecast Pro projects and from Excel spreadsheets.

You can create event schedules 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. You may also import helper variables from other Forecast Pro projects 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 the **Create Helper** icon on the **Forecasting** tab) or you can create them in Excel or in a database and import them into your project using the Data Manager. You may also import helper variables from other Forecast Pro projects 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.

    a. The Description field defines the name of the helper variable.

    b. The fields Starting year, Starting period, Periods per Year and Periods per Cycle should all match the historic data.

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.*

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 field you'll want to examine the *Modifiers* table in the file *123 Bakery - ODBC.mdb*.

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.

The Forecast Pro Interface

This chapter presents a command overview and a description of the user interface. Details for the functionality of the views and dialog boxes available from the Ribbon menu are provided in subsequent chapters.

Operations Overview

Generating forecasts in Forecast Pro consists of executing the following steps:

1. Prepare the input data. 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 Setting Up Your Historic Data and Setting Up Your Optional Data chapters.

2. Specify the default options. This is accomplished using the Options dialog box, which may be opened with the **Options** icon on the **Home** tab. Options set by the user will apply to the current project file 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. An overview of all option tabs and settings is in Forecast Pro Options.

3. Define the historic data and other inputs. This is accomplished by specifying the appropriate entries in the Data Manager. Consult Using the Data Manager for details.

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 Artificial Intelligence (AI) driven 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 the project file so that you can update it the next forecast period rather than creating a new one from scratch. Consult Saving Your Work for details.

In addition to the above required steps, there are many optional operations that you can perform. Including:

1. View the forecasts and other information graphically

2. View on screen reports.

3. View forecasts and other information numerically in the override grid and enter any desired changes and comments.

4. Customize how the forecasts are generated by specifying the forecasting models and reconciliation techniques to use.

Forecast Pro Interface Overview

The Startup Page

When you open Forecast Pro, the program will display a startup page like the one below.
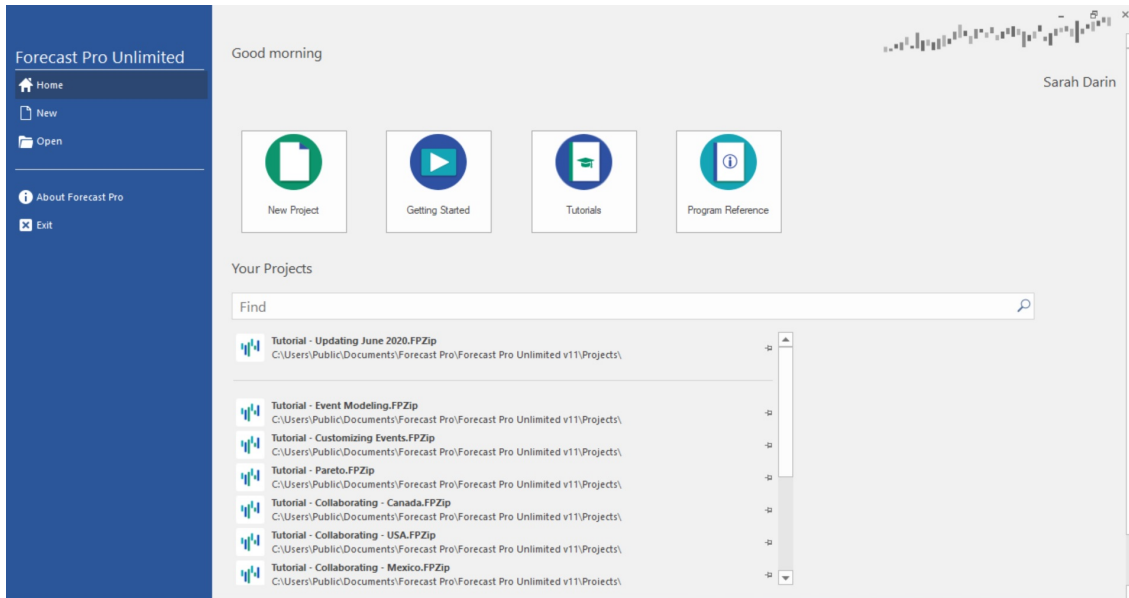
The startup page allows you to open an existing project, start a new project, view a Getting Started video, open the Tutorials manual or open the Program Reference manual.

Once you start a new project or open an existing project, the main Forecast Pro interface appears.
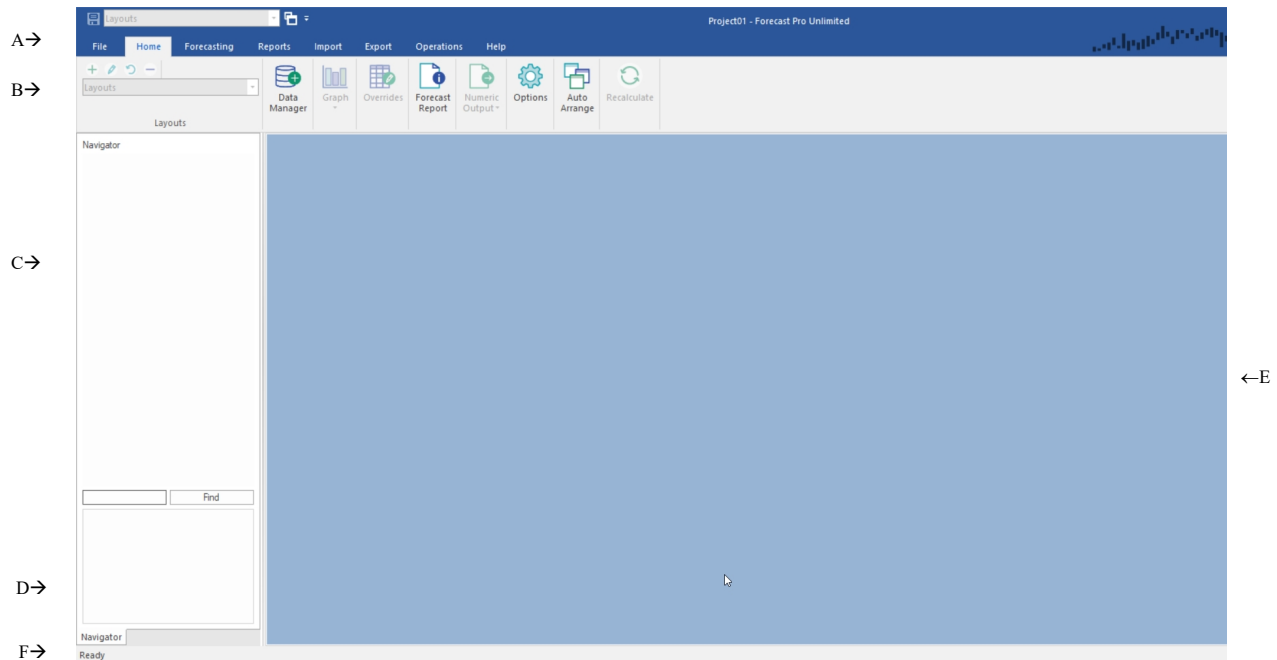
A Quick Access Toolbar **B** Ribbon Menu **C** Navigator **D** Hotlist **E** View Window **F** Status Bar

The Forecast Pro interface consists of a Quick Access Toolbar (A), a ribbon menu (B), a Navigator (C, currently unavailable), a Hot List area (D), a view window (E) and a status bar (F).

Quick Access Toolbar

Located in the upper left-hand corner of the display, the Quick Access Toolbar contains frequently used icons. The Quick Access Toolbar initially contains the Save and Auto Arrange icons as well as a Layout selection drop-down and a Customize drop-down. The Customize drop-down allows you to specify the icons you wish to display in the Quick Access Toolbar.

Ribbon Menu

The ribbon menu provides access to almost all Forecast Pro functionality. The functionality is organized into seven ribbon tabs: **File, Home, Forecasting, Reports, Import, Export, Operations** and **Help**. Clicking on one of these tabs will display the ribbon for the selected tab. Each ribbon is a collection of icons and drop-downs, often organized into groups. The most commonly used procedures are available on the **Home** tab (selected above). The seven tabs are described in *Forecast Pro Ribbon Tabs* chapter.

The Navigator

The Navigator is the primary way to select an item to view in the Forecast Report, Graph, and Override windows. It is also how you select the item you want to customize using the **Forecasting** tab. After the data have been read, the Navigator displays the available time series in a tree structure. The tree structure is defined by the hierarchy in your historic data files. 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 either *Manual forecast* or *Manual override* mode is selected on the Performance tab of the Options dialog box) indicates that an item's forecast needs to be updated with the Recalculate icon on the **Home** tab. A dark blue icon (applicable only if *Manual override* or *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 icon on the **Home** tab.

The Navigator context menu provides extensive functionality for working with your data and the Hot List. Please see *The Navigator Context Menu* for details.

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. A Hot List can also be defined based on the items currently displayed in a report. You may use filters in a report to get the desired subset of items, then use the report's context menu or hamburger menu to add all the items currently displayed in the report to the Hot list.

When you select an item on the Hot List, Forecast Pro will immediately select the item on the Navigator and update the affected views.

Right click in the Hot List area to open the Hot List's context menu. The Hot List's context menu allows you to sort the Hot List alphabetically, change how items are displayed 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.

Please consult the *Advanced Navigation* lesson in *Forecast Pro Unlimited Tutorials* for more details.

View Window

Forecast Pro allows you to open multiple views in the View Window. These views may be resized, maximized or minimized and arranged as you wish using standard Windows functionality. Left click on the border of an open window and drag the mouse to resize and left click on a window's title bar and drag the mouse to move the window.

You may also tile the open views using **Auto Arrange** icon on the **Home** tab and Quick Access Toolbar. This can be helpful when you open a new view. The drop-down on the **Auto-Arrange** icon allows you to tile windows vertically instead.

Each Forecast Pro view has a hamburger menu on the upper left-hand corner. This menu can be used to access key functionality, such as opening the settings dialog for the view or adding filters to the report.

When you add filters to a report, Forecast Pro includes a filter drop-down on each column. If you click on the filter drop-down, you will bring up a dialog box that allows you to sort and filter values.
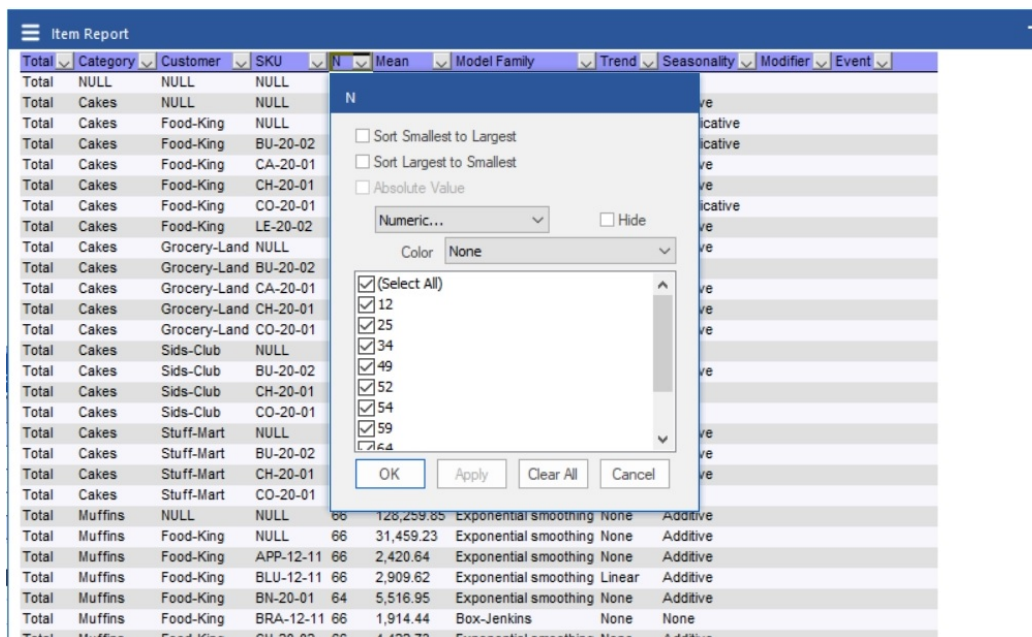
From this dialog box you may sort the report by the values (or absolute values) in the column, filter the report based on specific rules using the Numeric drop-down or apply a heat map using the Color drop-down.

On the drop-down currently selected Numeric, you may select *Numeric, Top 10, Above Average* or *Below Average*. The Above Average and Below Average selections are self-explanatory. Selecting *Numeric* brings up a dialog box that allows you to enter logic for how the report should be filtered.
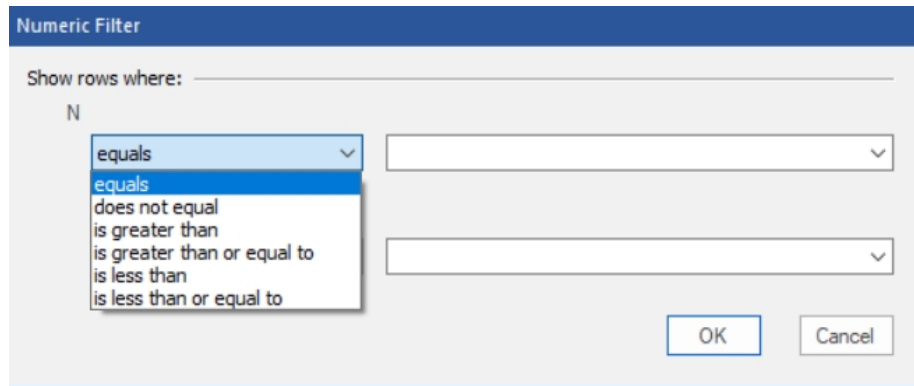
Selecting *Top 10* brings up a dialog box that allows you to show the top or bottom items, based on number of items, percent or volume. While the default is to show the top 10 items, you can change the 10 to whatever value you wish to see.
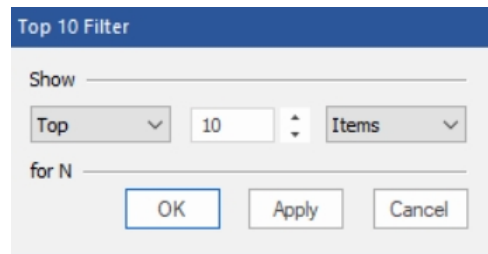
Finally, you may add heat maps to numeric columns by selecting from the Color drop-down. The heat map colors a cell based on where the cell value falls in the range of values for the column. The default selection, *None*, is to not use a heat map. The drop-down provides different color schemes to select from. If you choose a color scheme, a heat map with the selected color scheme will be applied to the column specified.

Each Forecast Pro view also has a context menu. Right click anywhere on an open view to bring up the context menu. The context menu has the same selections as the hamburger menu which is accessible in the top left corner of each view.

There are eight view windows available in Forecast Pro. The following are context specific:

Forecast Reports

The Graph Views

The Override Forecasts View

The Diagnostics View

Context specific views display information for the current Navigator selection and update automatically as you move about the Navigator.

The remaining four views are global and include:

The Item Report View

The Override Report View

Outlier Reports

The Numeric Output View

The Status Bar

The Status Bar, on the bottom of the Forecast Pro window, displays key information about the status of the open project.

On the right-hand side you will see the beginning and ending periods for your historic data. You may edit this data range on the **Operations** tab.

The left-hand side of the status bar displays dynamic information about the current project status. When using the override grid, the status bar will display the total and average for selected override grid cells. These are also displayed if you select numeric values in one of the report views.

Forecast Pro Ribbon Tabs

File Tab

Selecting the **File** tab will take you to the following screen.
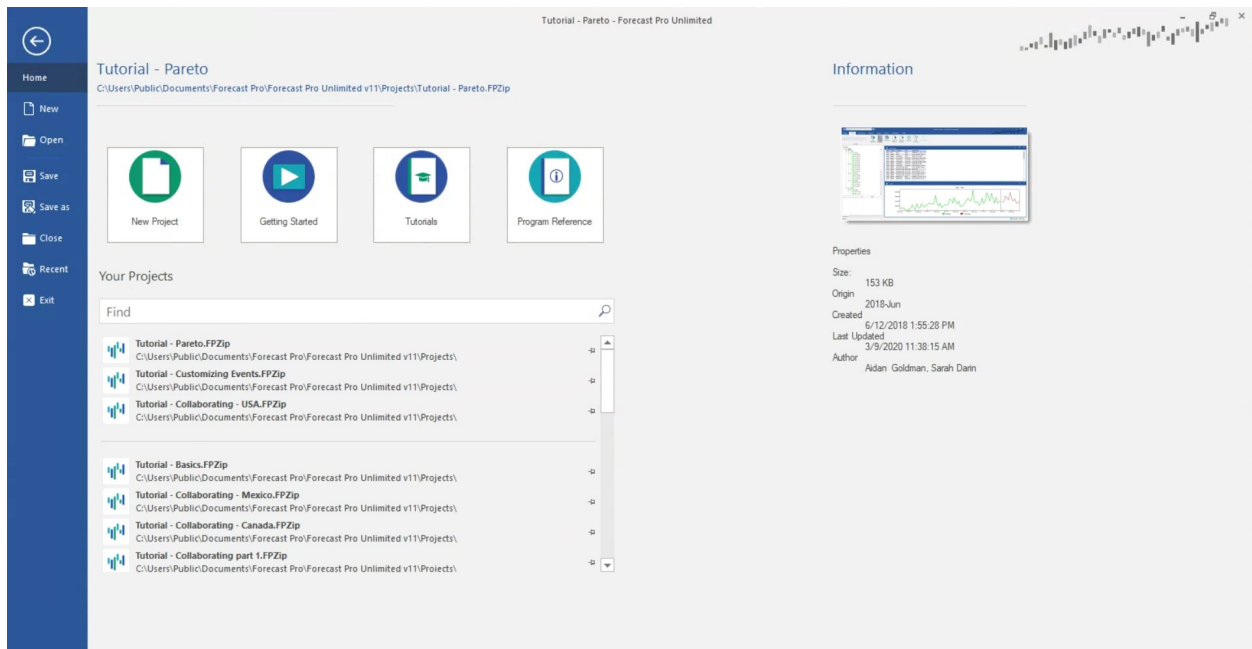
When you are on the File tab, you do not have a view of your project or the rest of the Forecast Pro Interface. On the right side, there is an Information pane displaying key details about the current project. Along the top, you will see the project's filename and the folder the project file is saved in.

On the top left, common file operations, including New, Open, Save, Save As, Close and Exit, are listed.

In the middle section there are 4 large icons: **New Project** (opens a new project), **Getting Started** (links to a short video), **Tutorials** (opens the Tutorials) and **User Guide** (opens the User's Guide). Beneath these icons you will see a pane that displays recent projects and a list of all projects in your default project directory (set in Options).

Home Tab

Forecast Pro's most frequently used commands are available on the **Home** tab.

Layouts: On the left-hand side of the **Home** tab is the Layouts group. The **Layouts** drop-down is used to select a saved layout to display. Above the drop-down are four small icons for **adding**, **managing**, **updating** and **removing** layouts. Please see details in the "Working with Layouts" section of the *Advanced Navigation* lesson in *Forecast Pro Unlimited Tutorials* or Working with Layouts in this reference.

Data Manager: To the right of the Layouts group is the **Data Manager** icon. The Data Manager is used to read data into the project. Once data is read in, the Navigator displays all items to be forecast, and the Hot List becomes available for use. Graphs will also become available after data has been read in. Consult Using the Data Manager for details.

Graphs: Click on the **Graph** icon to open a graph view for the item selected in the Navigator. Up to four additional graphs may be added to the view window by selecting *Add Graph* on the Graph icon's drop-down. Finally, you may select *Diagnostics* on the Graph drop-down to open the Diagnostics view. Details for how to work with the Graph and Diagnostic views are in The Graph Views and The Diagnostics View.

Overrides: The Override icon opens the Override view which allows you to adjust forecasts for the selected item. Consult The Override Forecasts View for details.

Forecast Report: The **Forecast Report** icon opens the Forecast Report. The Forecast Report view displays the forecasting model, within-sample statistics, the historic data and forecasts for the selected item. If you are using a holdout sample, the Forecast report view will also include an out-of-sample evaluation section. Details for using the Forecast Report view are in The Forecast Report View.

Numeric Output: The **Numeric Output** icon provides a quick way to save your Numeric Output in a single click. 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. If you will be importing the forecasts into another application, you will likely do so with the Numeric Output file.

The drop-down on the **Numeric Output** icon allows you to open the Numeric Output view (select *Show*) and the settings associated with the Numeric Output file. The Numeric Output view displays the contents and format of the currently specified Numeric Output file, which can both be adjusted using the Numeric Output settings. You do not need to display the view in Forecast Pro to adjust the settings. Consult The Numeric Output View and Numeric Output Files for more details.

Options: The **Options** icon opens the Options dialog box. This dialog box controls project-level settings. Typically, you will want to set your options when you are starting a project. Please consult Forecast Pro Options for details.

**Auto Arrange:** Forecast Pro allows you to open multiple views in the View Window. These views may be resized and arranged as you wish. The **Auto Arrange** icon (with drop-down) tiles the open views

horizontally. This can be helpful when you open a new view. Given that auto arrange is often used after opening a new view, the **Auto Arrange** icon is available on the Quick Access Toolbar by default.

Recalculate: This option is only available when *Manual override mode* or *Manual forecast mode* is selected on the Performance tab of the Options dialog box (see The Performance Tab). Selecting recalculate will reconcile all forecasts and overrides in the project.

Forecasting Tab

By default, Forecast Pro will use expert selection, an Artificial Intelligence (AI) driven engine, to automatically select a forecasting model for each item on the Navigator. You may specify if you want expert selection to include machine learning models, Box-Jenkins (ARIMA) models, seasonal simplification and/or NA-CL models on The Performance Tab of the Option dialog box. Alternatively, you may select alternative models and/or forecasting options on the **Forecasting** tab.

Available models, modeling options and **Forecasting** tab functionality are documented in detail in Customizing Your Forecasts.

Data Tools Group

Diagnostics

Create Helper

Outliers

Start Dates

Remove Modifier

Reports Tab

The Reports tab provide access to all Forecast Pro Views. Details for each of these views are in the Forecast Pro Views chapter.

Graphs: Click on the **Graphs** icon to open and close a Graph view for the item selected in the Navigator. Up to four additional graphs may be added to the view window by selecting *Add Graph* on the **Graph** icon drop-down. Finally, you may select *Diagnostics* on the **Graph** icon drop-down to open the Diagnostics view.

Overrides: The **Override** icon opens and closes the Override view, which allows you to adjust forecasts for the item selected in the Navigator.

Forecast Report: The **Forecast Report** icon opens and closes the Forecast Report. The [Forecast Report](#) view displays the forecasting model, within-sample statistics, the historic data and forecasts for the selected item.

Reports Group: The Reports group includes the following:

> **Item Report**: The **Item Report** icon opens and closes the [Item Report](#). Item Reports may 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. The drop-down on the **Item Report** icon allows you to open the Settings dialog box associated with the Item Report. The functionality of the Item Report is essentially the same as the Numeric Output view, but the Item Report allows you to change the view without impacting the structure of the Numeric Output File.

> **Outliers**: The **Outliers** icon opens and closes the [Outlier Report](#), a global report that list all outliers identified in the project.

> **Override Report:** The **Override Report** icon opens and closes the [Override Report view](#) that shows all overrides in the project.

> **Numeric Output**: The **Numeric Output** icon opens your Numeric Output view. Numeric Output Reports 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. The drop-down on the **Numeric Output** icon allows you to open the Settings dialog box associated with the Numeric Output. Note that these setting are typically configured for exporting a file with a defined structure, so changing these settings may disrupt your forecasting process. Consult the [Numeric Output View](#) and [Numeric Output Files](#) for more details.

Import Tab

The **Import** Tab includes the functionality for reading in project data as well as several other types of information, possibly from other saved Forecast Pro projects, into your current project.

Data Manager: Clicking the **Data Manager** icon opens the Data Manager dialog. The Data Manager is used to read data into the project. Once data is read in, the Navigator displays all items to be forecast, and the Hot List becomes available for use. Graphs will also become available after data has been read in. Consult [Using the Data Manager](#) for details. Note that the **Data Manager** icon is also available on the **Home** tab.

Project Settings: The **Project Settings** icon 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.

Hot List: The **Hot List** icon 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 the **Hot List** icon on the **Export** tab or using the *Save Hot List* 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. See [The Hot List](#) for more details.

Layouts: The **Layouts** icon is used to import layouts from a saved project into the current project. Please see details in the <u>Working with Layouts</u> chapter of this reference.

Export Tab

Reports Group: The Reports Group includes functionality to export the Numeric Output or any of the other Forecast Pro reports for all or a subset of the items in your project.

On the left-hand side of the group, you can select which items you want to export. You may select *Current Display* (includes the items shown if you opened the report view, including any active filters), *Hot List* (items currently displayed in the Hot List) or *All*.

Clicking any of the Reports group icons (Numeric Output, Forecast Report, Item Report, Outliers and Override Report) will prompt you to save the report for the specified items.

Note that the **Forecast Report** icon will save a Formatted Forecast Report, including a tab for each exported item, in the saved Excel file.

Other Group: The icons in the **Other** Group prompt you to save the **Hot List, Helpers** or **Modifiers** in the saved current project. These files may then be read into another Forecast Pro project with the **Import** tab.

Operations Tab

Date Range Group: The Date Range controls are used to set the span of data used when estimating forecast models, define a holdout sample and reset dates to the default of using all available data without a holdout sample.

> **Year and Period:** The **Year** and **Period** spinners are used to define the fit set. All time series methods begin by fitting 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. By default, Forecast Pro 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 is applied to all forecasted items, unless you have used the **Start Dates** icon on the **Forecasting** tab to specific alternative start dates for individual items.

> **Holdout Periods**: This spinner is used to specify the length of a holdout sample. When using a holdout sample, Forecast Pro withholds the specified number of periods at the end of the data from the model fitting. If you withhold 2 or more points, Forecast Pro automatically generates out-of-sample evaluation statistics for any models that you build. These statistics are available in the Forecast Report. If you do not want to perform out-of-sample testing, set the holdout sample to zero (the default).

Read Data: The **Read Data** icon 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.

Forecast: Clicking the **Forecast** icon generates statistical forecasts. This option is only available when historic data has been defined in the Data Manager and the data has been read in. Changes to the date range used may require you to re-forecast the data.

Recalculate: This option is only available after changing the date range or when *Manual override mode* or *Manual forecast mode* is selected on the Performance tab of the Options dialog box (see The Performance Tab). Selecting the **Recalculate** icon will reconcile all forecasts and overrides.

Password: This icon 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.

Help Tab

Forecast Pro includes extensive online help.

Tutorials: The **Tutorials** icon opens Forecast Pro'*s* Tutorial website. The website includes online video tutorials as well as written tutorials. These tutorials are designed to give you an overview of Forecast Pro's functionality. The tutorials are organized into two sections: **General Operations** and **Forecasting Operations**. The General Operations tutorial is designed to teach you how to use the Forecast Pro interface, while the Forecasting Operations tutorial focuses on customizing your models and forecasts.

Program Reference: The **Program Reference** icon opens a pdf version of the *Forecast Pro Unlimited Reference.*

Statistical Reference: The **Statistical Reference** icon opens a pdf version of the *Forecast Pro Statistical Reference Manual*.

Check for Updates: This icon 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.

About Forecast Pro: This icon displays the Forecast Pro version number, copyright notice and your license information (Serial Number and CD Key).

The Navigator Context Menu

The Navigator's context menu provides a convenient way to build a Hot List, apply forecast modifiers and control the display of the Navigator tree. 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* refers 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 options under *Forecasting* are all used to apply forecasting models and 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. All forecasting-related functionality on the Navigator's context menu is also available on the **Forecasting** tab. Please consult Customizing Your Forecasts for details on available forecasting options and supported forecast modifiers.

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. You may also access the Create Helper Variable dialog box from the **Forecasting** tab. Please consult Helper Variables in *Customizing Your Forecasts* for details on creating helper variables.

If you select a helper variable on the Navigator, you can use the *Helper Variable>Delete* option to delete it from your project.

Working with Layouts

Layouts are an excellent way to standardize your forecasting process. Layouts allow you to save and return to setting and window view configurations. If you set up your forecasting process with layouts, you can easily access the key information you need each time you update your project. Please consult the Working with Layouts section of the *Advanced Navigation* lesson and the Touring the Five Reports section of the *Exception Reporting* lesson in *Forecast Pro Unlimited Tutorials* for more details on using layouts.

To make a layout, you first create a project view that you find useful. For example, you may want to have a graph view and an item report filtered to products with less than 12 periods of history. If you are using monthly data, this view will show you all items with less than a year of history. You can view each item in the graph by double clicking on each item in the item report list. Next, you will want to save that view as a layout.

Layouts are managed from the **Home** tab.
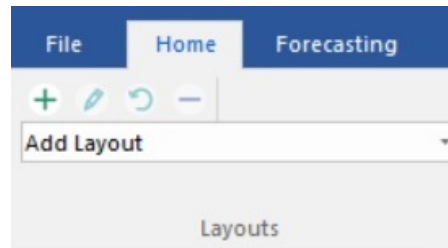
Adding Layouts

The ➕ icon opens the Add Layout dialog.

To add a layout, simply enter a layout name, a description (if desired) and click OK. If you click on the Layouts drop-down list on the Home tab or Quick Access Toolbar, you will now see your saved layout included. If you select your layout from the drop-down menu, Forecast Pro will display the view and setting configurations as saved in the selected layout.

Managing Layouts

The Layout settings can be configured by clicking the Manage Layout icon (✎) to open the Layout Manager dialog box.
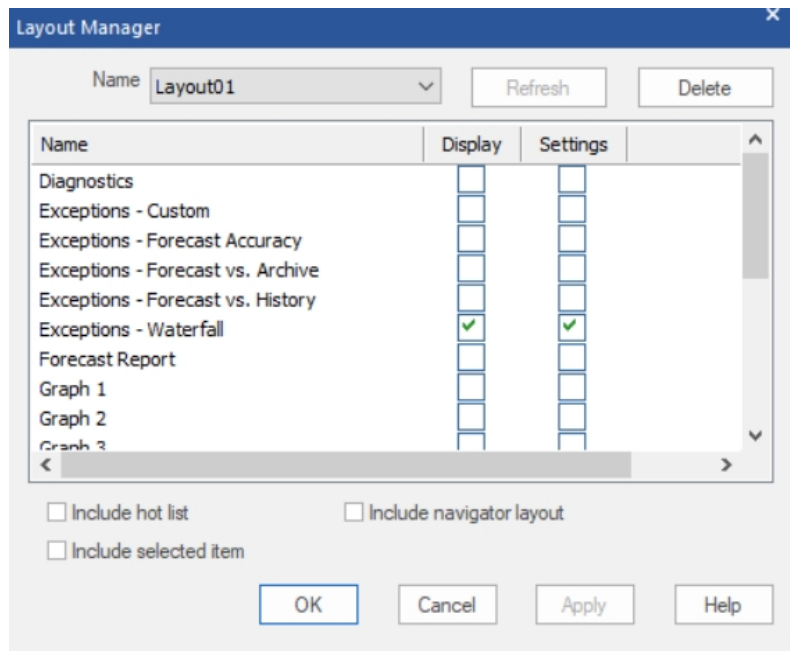
Name: This drop-down shows four Layout commands, *New*, *Rename*, *Save As* and *Delete*, as well as the layouts that have been saved to the current project. The drop-down displays the currently selected layout.

New: Opens the Add Layout dialog box described above.

Rename: Opens the Rename Layout dialog box where you can enter a new name.

Save As: Opens a Save As dialog box that will take the settings for the selected layout, along with any changes you have made in the Layout Manager dialog and create a new layout with the specified name.

Delete: Removes the selected layout from the current project.

Update: The Update button updates the currently selected layout with specified changes.

Delete: The Delete button removes the selected layout from the current project.

Display: The Display column allows you to specify, for each Forecast Pro view, if you want the selected layout to save display and/or settings configurations for the selected view. Display settings include view window status (open, closed etc.) and view location, while settings include anything that is specified in the view's Setting dialog box. Display is checked for all views that should be shown in the layout. If a view has Settings checked but not Display, Forecast Pro will load the settings as specified in the layout but not open the view window. This may be useful if you want to export reports with saved settings, but you do not want to view the Report in Forecast Pro.

Include hot list: Check this if you want the layout to save hot list items, so that when the layout is selected the hot list displays those saved items.

Include selected item: Check this if you want the layout to associate the item selected on the Navigator with the layout.

Include navigator layout: Check this if you want the layout to save the Navigator layout (e.g. which nodes are open) and update the Navigator when you select the layout.

Layouts may be removed and refreshed directly from the **Home** tab using the remove layout icon (━). Clicking on the **Remove** icon opens the Remove Layout dialog box. Simply select the layout you want to remove and click OK to remove a layout from the Layouts drop-down and the current project.

Layouts may also be updated directly from the **Home** tab with the **Update** icon (↺). Click the **Update** icon to launch the Update Layout dialog box. Selecting a layout in the dialog box and clicking OK will replace the selected layout with the current Forecast Pro view and settings.

Using the Data Manager

The Data Manager provides a centralized place to define all input data for Forecast Pro and keep it up to date.

Row Types

You invoke the Data Manager by clicking the **Data Manager** icon on the **Home** tab.
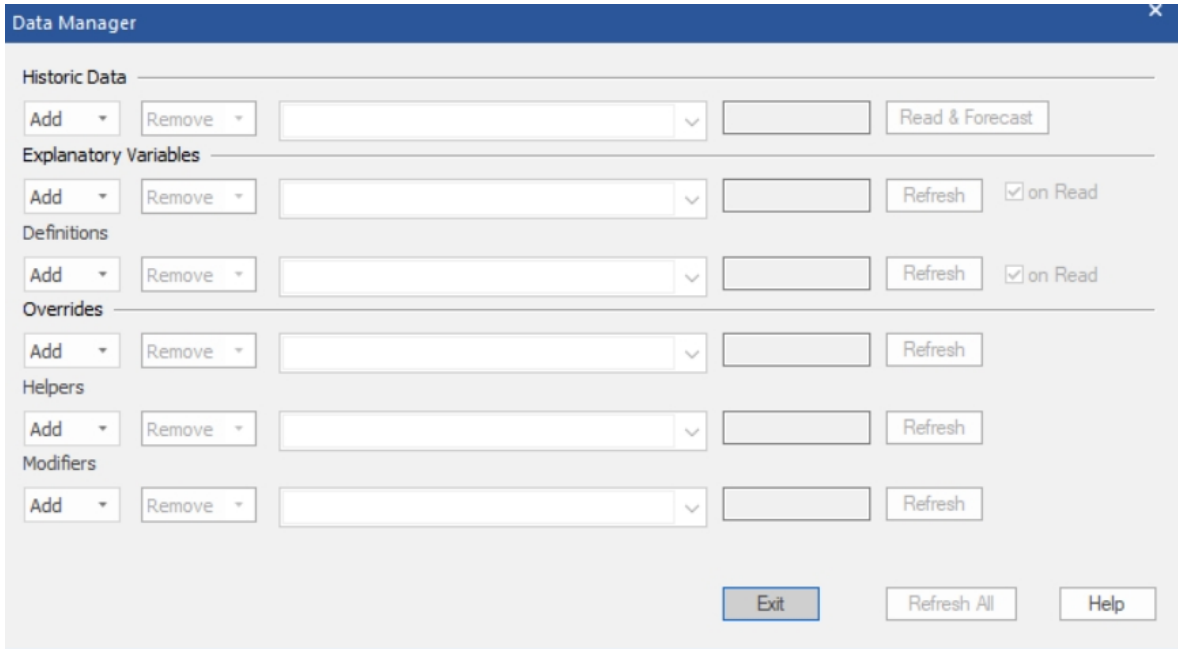
The Data Manger consists of six rows each of which allows you to read in a different type of data.

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:

Explanatory Variables. Variables that you may wish to include in a dynamic regression model may be read in with an explanatory variable file. Explanatory Variables provides details.

Definitions. Definition files contain values to use in conjunction with custom forecast allocations. Consult Custom Allocation Files for details.

Overrides. In addition to entering forecast overrides interactively using the Forecast Overrides view, Forecast Pro allows you to import them. Consult Override Files for details.

Helpers. Helper variables are used in conjunction with event models, forecast by analogy models and weighting transformations. Helper Files provides details.

Modifiers. In addition to specifying forecast modifiers interactively on the Navigator, Forecast Pro allows you to import them. Modifier Files provides details.

Row Controls

Each row has controls associated with it that allow 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. Note that you can only read from one ODBC connection at a time.
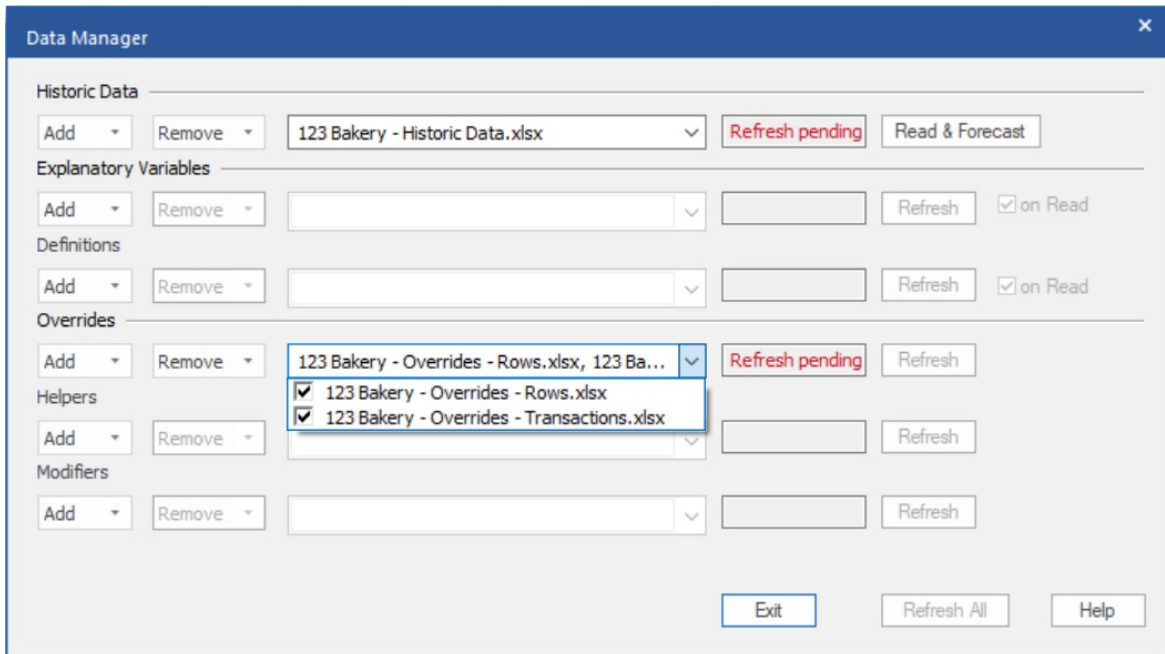
Selected data sources drop-down. After you add one or more data sources to a given row, the filenames, table names, query names, or project names will appear on the selected data sources drop-down display. If you open this drop-down, you will see a list of the selected data sources along with a checkbox 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 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. Five 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 (or clicking Refresh All). This is because overrides, modifiers, helper variables and mapped histories 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 except the Historic Data. Note that it will refresh all rows that have defined data sources even if their status is up to date.

Forecast Pro Options

The **Options** icon on the **Home** tab opens the Options dialog box which is used to change various Forecast Pro settings.

The Options dialog box contains both a **Set as Default** button and an **OK** button. Clicking the OK button will apply the defined settings to the current project. Clicking the Set as Default button will save your current settings as the default settings for all new projects.

The Options dialog box contains eight pages or tabs. We will discuss each in turn.

The Basics Tab
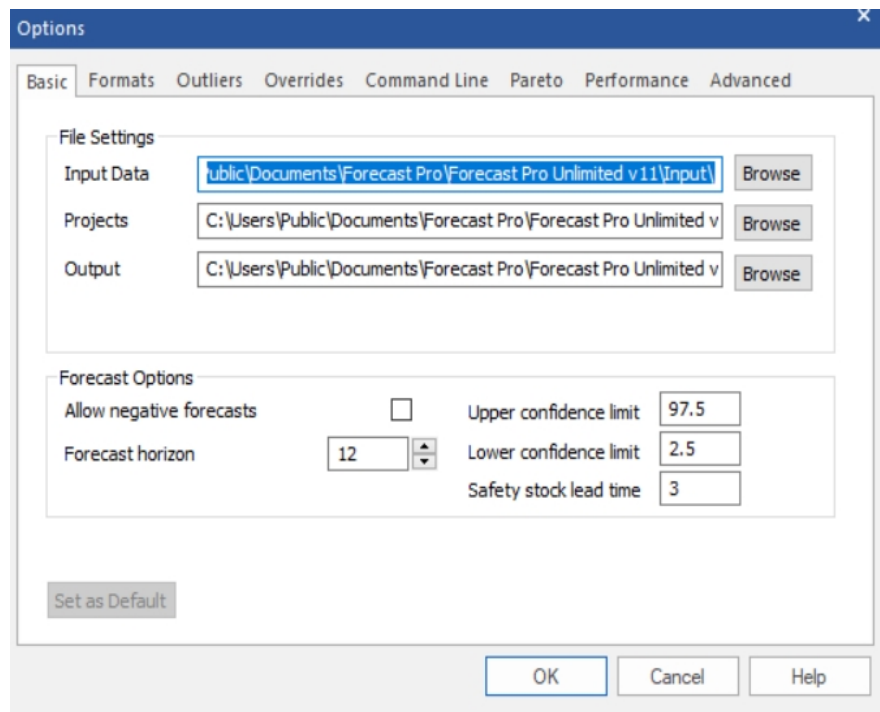
The Basic tab contains frequently accessed settings that every user should understand.

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 replace negative forecasts with zeros. 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
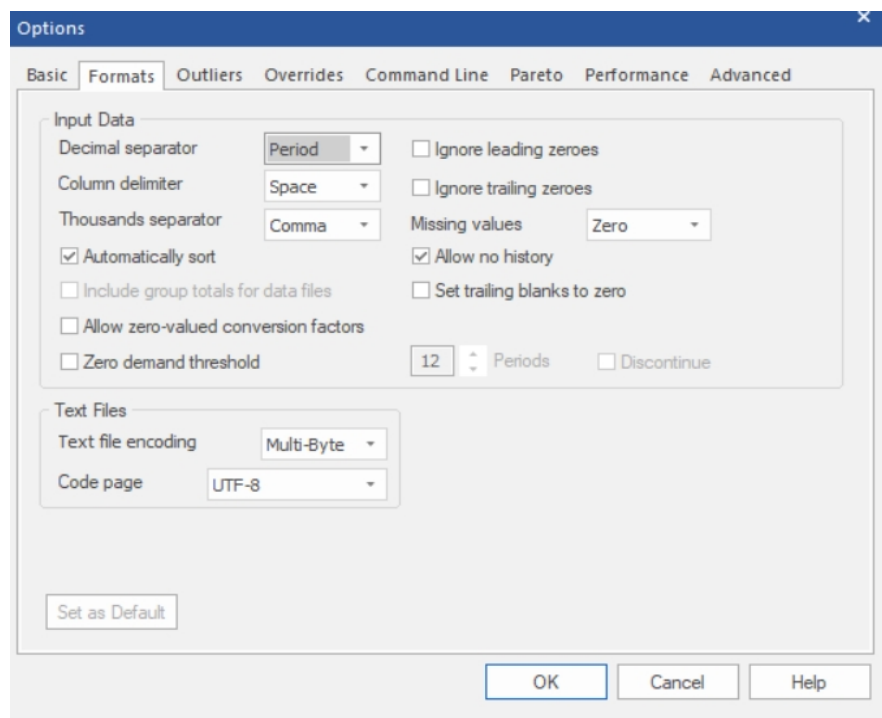
The Formats tab controls the formatting of the input and output.

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 does 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, 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 are considered 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" and not forecasted. 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. The Missing Values drop-down allows you to specify how you wish the program to treat missing values. Options include: *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 *Setting Up Your Historic Data* chapter.

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 zeros 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 chapter entitled *Setting Up Your Historic Data*.

Text Files:

Text file encoding. Specify the format to use when writing two-byte characters to text files. This option is only relevant if your data uses 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.)

The Outliers Tab
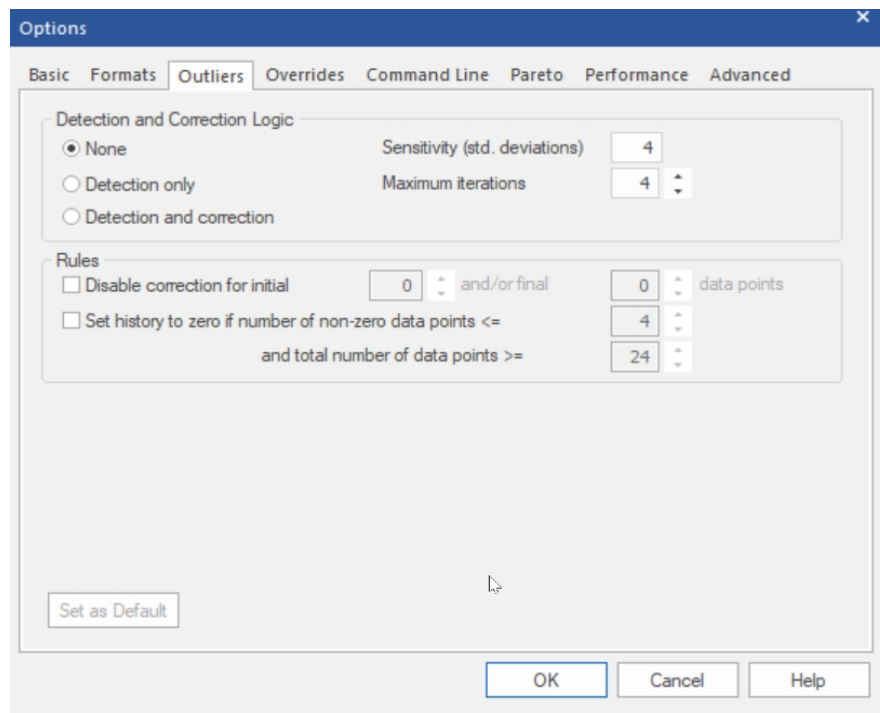
The Outliers tab is used to specify if and how Forecast Pro should detect and/or correct outliers. It also allows you to set the sensitivity settings for outlier detection and to dictate the amount of detail in the outlier report.

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 more details on outlier detection and correction.

Maximum iterations allows you to set the maximum number of iterations permitted during outlier detection for a given item. Since only one outlier is identified per iteration, this setting also defines the maximum number of outliers that can be detected for a given item. Consult the *Forecast Pro Statistical Reference Manual* for more details on how outlier detection and correction.

Rules:

Disable correction for initial X and/or final Y data points allows you to disable correction for data points at the beginning and/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/or 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. This setting only has an impact if *Outlier Detection and correction* is enabled.

The Overrides Tab
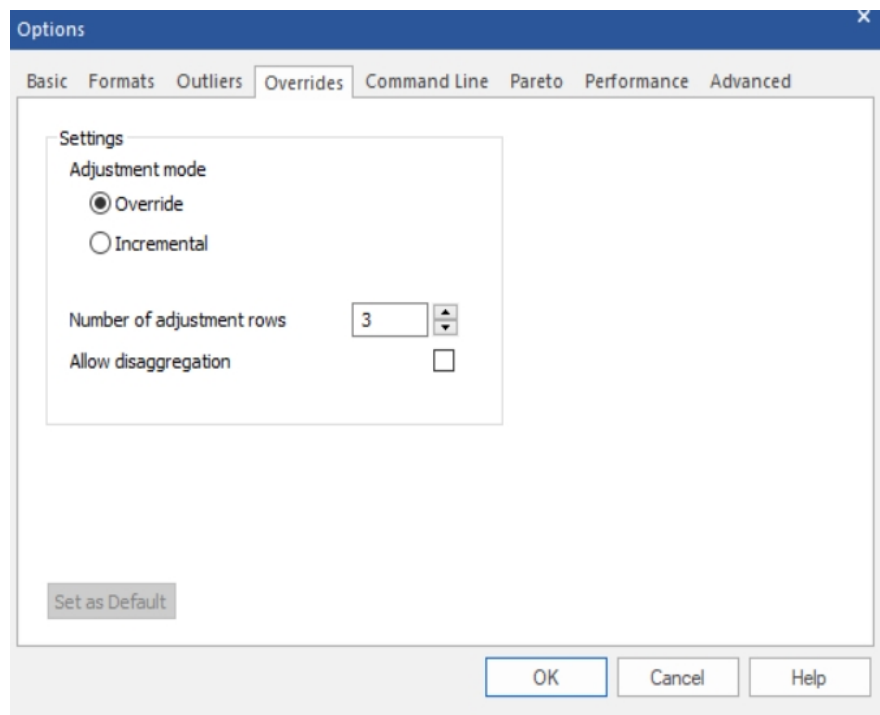
The Overrides tab controls operation and display of the Override view.

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, and 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 lower levels of the hierarchy. The disaggregation is based on a proportional allocation.

The Command Line Tab

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 Command Line Operation.

The Command Line tab allows you to dictate the options to apply if the current project is run via the command line.

Modifiers. This drop-down 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 included in the Modifiers file in the Data Manager. Thus, the Data Source will take precedence over the Project for all modifiers.

Overrides. This drop-down 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 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. You can change the default settings for new projects by pressing the Set as Default button on any tab of the Options.

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, which is why the default behavior is to only output items listed in the Hot List.

Item report, Outlier report, and Override report. If any of these options are selected, the command-line run will save the specified report(s). The format and contents (including current display vs. all items vs. Hot List only) will match the current settings in the project. Note, the contents of the reports can be altered using the default scoping (see Performance tab), export tab, or in the context and hamburger menus.

Project. If this option is selected, the command-line run will update and save a new project file.

The Pareto Tab

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

Weighting Basis:

History and *Forecasts*. These settings allow 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 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
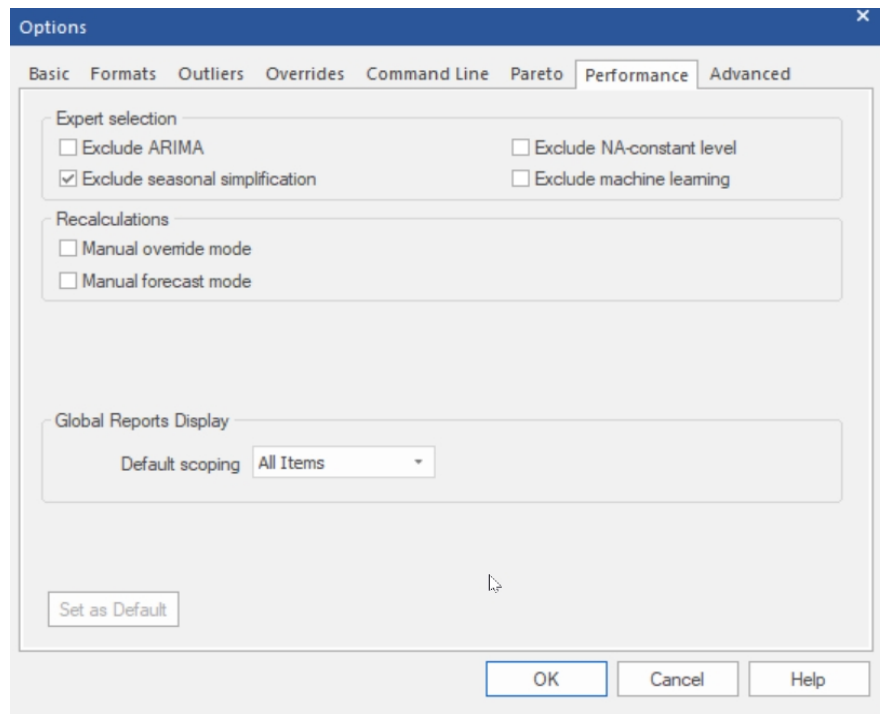
The Performance tab allows you to set expert selection options, select manual recalculation modes and set project operations. Adjusting these settings can improve processing speed and decrease time waiting for project operations and calculations.

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. Since seasonal simplification is only considered if the periods per cycle is 13 or greater, excluding seasonal simplification will speed up the forecast generation only if you have 13 or more periods per cycle.

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.

Exclude machine learning. If this switch is on, Forecast Pro's expert selection will not consider machine learning. Given that machine learning models are far more computationally intensive than other forecasting models, selecting this will speed up the forecast generation process dramatically. Please note that when this option is selected, you may still include machine learning in expert selection for select items using the \MLES modifier.

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 **Recalculate** icon (on the **Home** tab and **Operations** tabs or optionally on the Quick Access Toolbar) to perform the reconciliation once, rather than having to wait for the program to reconcile each time a new item is selected.

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 **Recalculate** icon (on the **Home** tab and **Operations** tabs or optionally on the Quick Access Toolbar). This allows you to specify multiple forecast modifiers without having to wait for the program to recalculate the forecasts after each specification.

Global Reports:

Default Scoping. When you are working with a very large project, loading reports for all end-items and group levels may take considerable time. The Default scoping drop-down allows you to specify if you want reports to show hot list items, the item selected on the navigator or all items. Selecting Hot List or Selected Item instead of All Items can speed up report loading time considerably.

The Advanced Tab
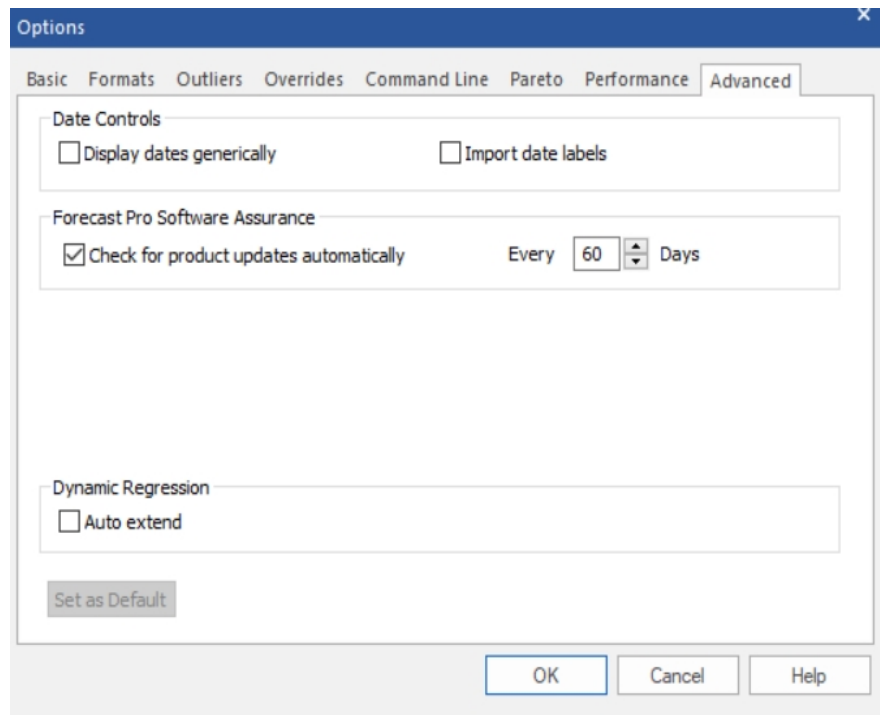
The Advanced tab includes the following:

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, "2019-01" will be used rather than "2019-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 helper variable and not use it for date labels.

Forecast Pro Software Assurance:

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 a newer version is available, clicking OK on the pop-up message should generate an email with your current version and the serial number you are using. 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 explanatory variable values and only forecast missing values. 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 Forecast Pro Views

The Forecast Report View

The Forecast Report view is a scrollable text display containing information about the forecasts and how they were generated. It provides a convenient way to view information while using Forecast Pro. The Forecast Report can be exported, but it can potentially be a very large file. For this reason, you should consider using another report designed for exporting (e.g. Numeric Output Report), or only export a subset of items in the Forecast Report.

The Forecast Report is opened and closed with the **Forecast Report** icon on the **Home** tab or the **Reports** tab.

The Forecast Report Settings dialog box allows you to customize the Forecast Report view. This dialog box is invoked by selecting **Settings** on the **Forecast Report** icon drop-down on the Reports tab, by selecting **Settings** from the Forecast Report view's context menu, or by selecting **Settings** on the hamburger menu in the upper left corner of the view window.
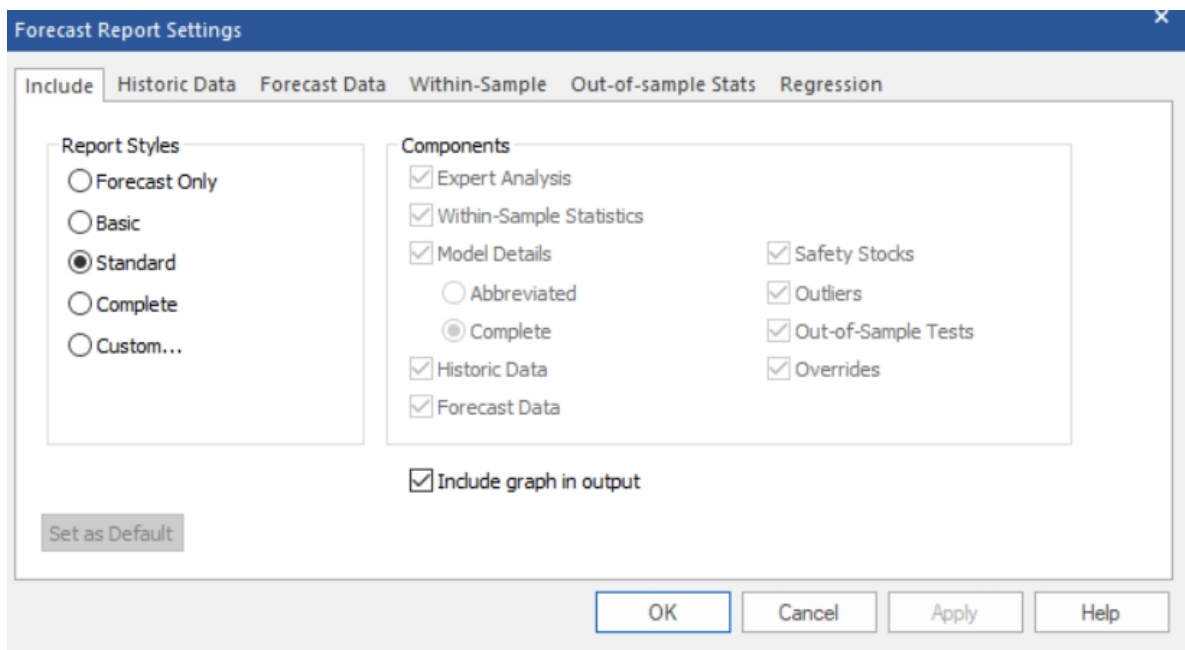
You can either select one of the four standardized report styles or define a custom report. *Include graph in output* can be used in conjunction with any style report and specifies whether to include or omit a graph when saving a Formatted Forecast Report to Excel. The settings for the saved graph may 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 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 gives you complete control over what is included in the display. The *Historic Data*, *Forecast Data, Within-Sample, Out-of-Sample* and *Regression* tabs on the Forecast Report Settings 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 may set precision for regression model details on the *Regression* tab*.* These precision settings are used for all report styles.

The Graph Views

The graph views are used to display variables and forecasts graphically. You can include up to five graph views. In addition to providing a convenient way to view the forecasts while using Forecast Pro, a 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 1 view.
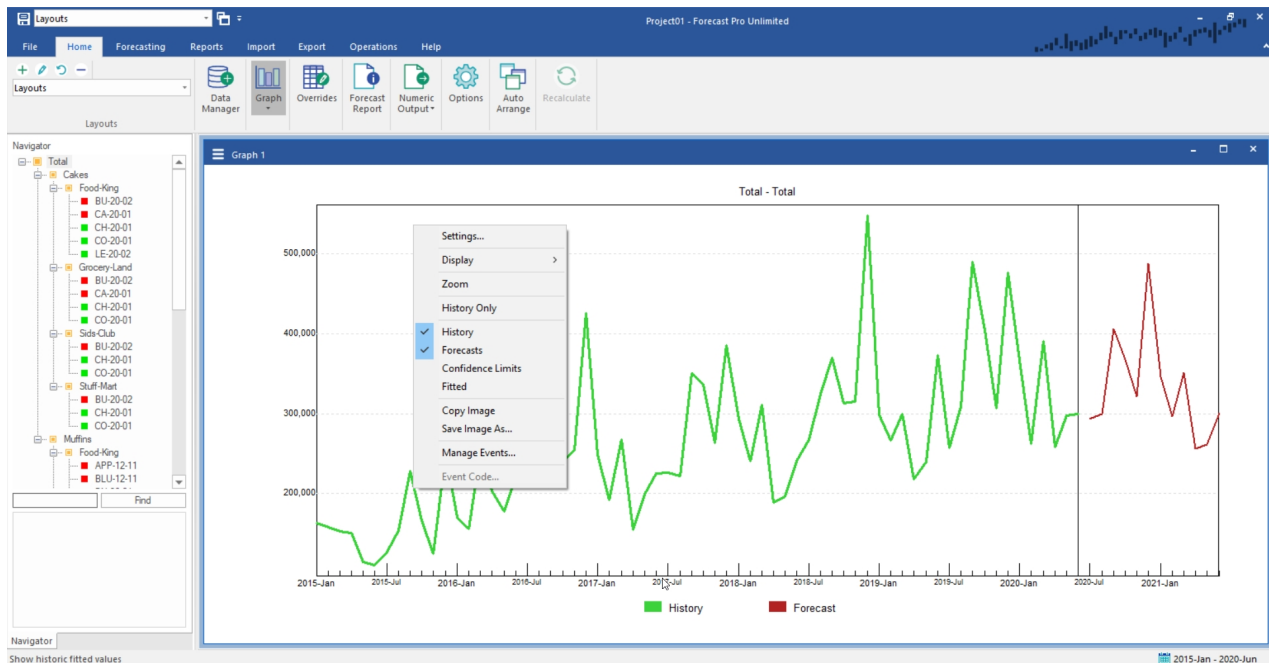
The Graph Settings dialog box allows you to customize the Graph view. This dialog box is invoked by selecting *Settings* on the view window's hamburger menu or by selecting *Settings* from the Graph view's context menu.

The Settings dialog box contains three pages or tabs. We will discuss each in turn.
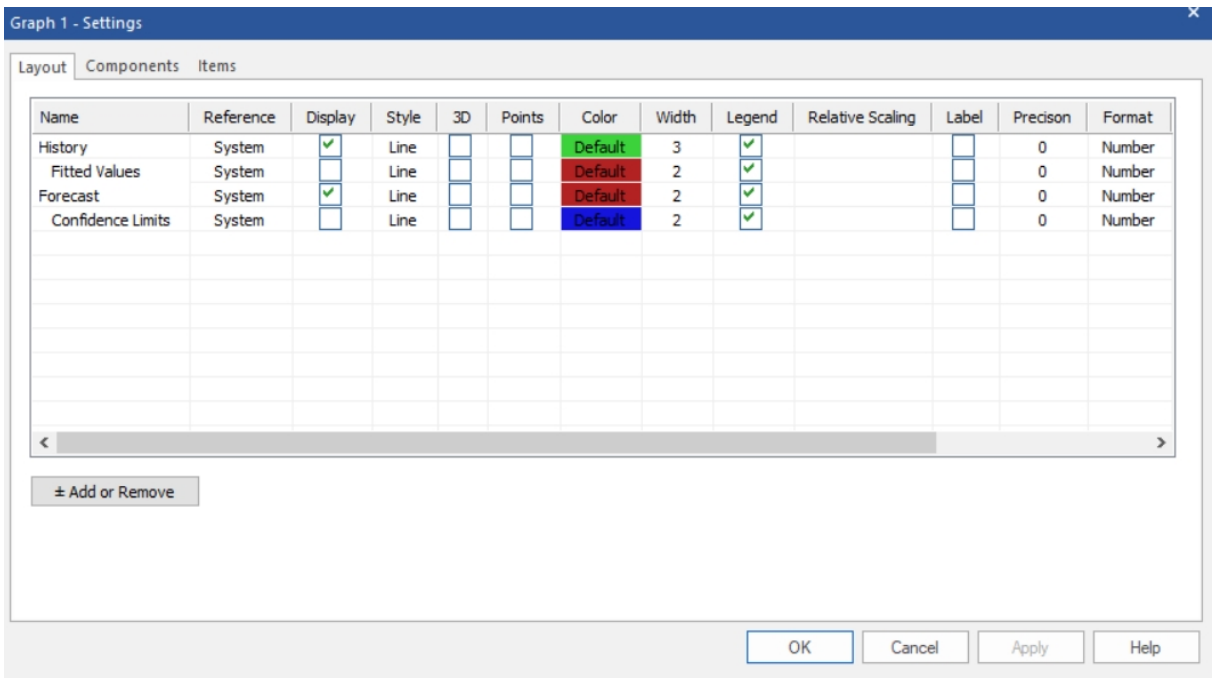
The Layout tab allows you to the select which components to display and format those components as you choose. The History, Fitted Values, Forecast and Confidence Limits components are always available in the Layout tab, and you can use the Display checkboxes to hide/unhide them from the Graph view. You can also use the context menu/hamburger menu to choose what is being displayed in the Graph view. 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.
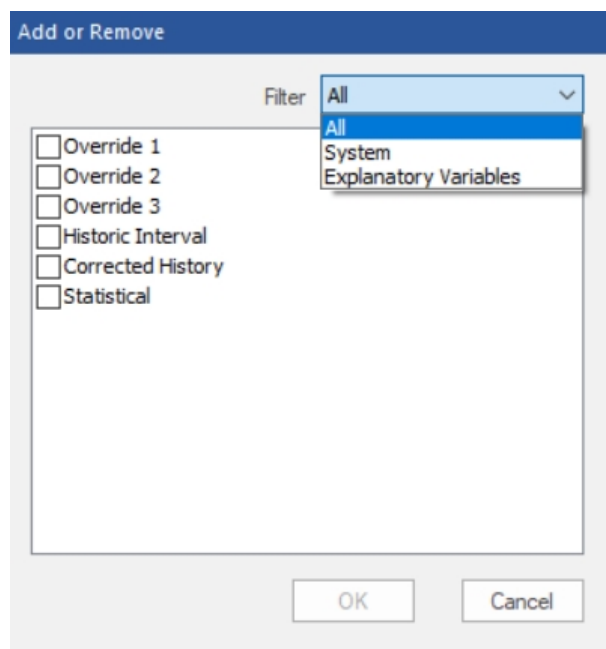
Clicking OK exits the Add or Remove dialog box and updates the list of components in the Layout tab.

You can set specific formatting options for each component.
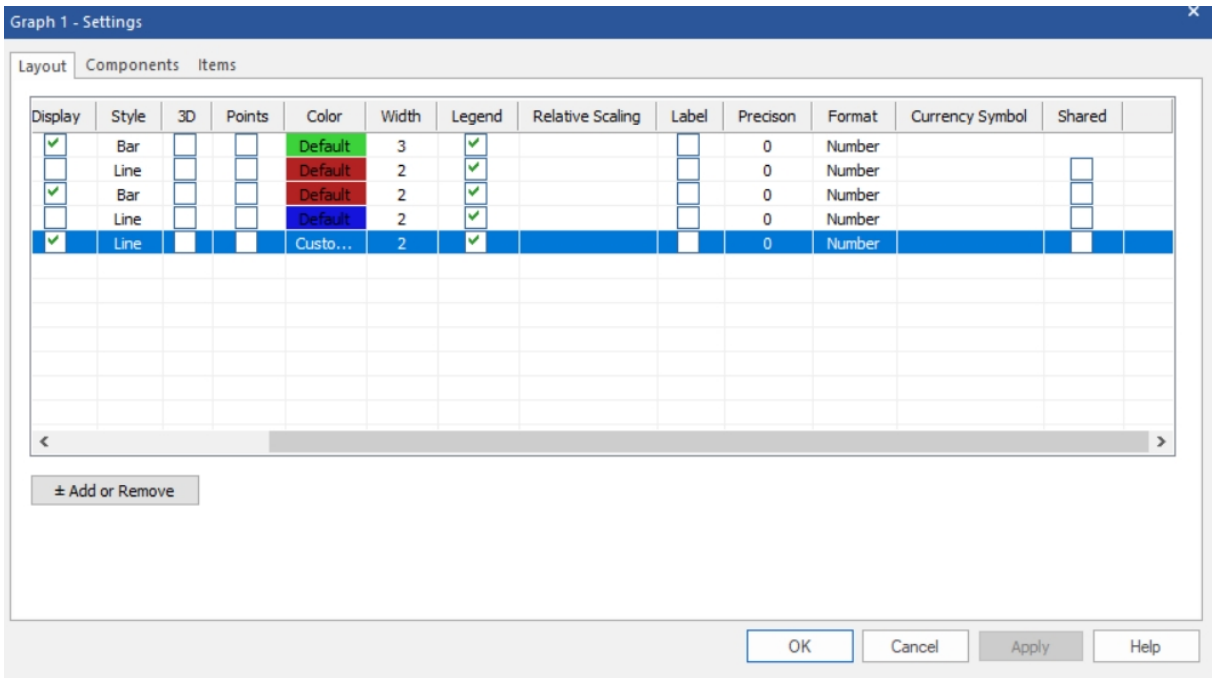
Most of these options are self-explanatory. The Style drop-down provides four options: Scatter, Line, Bar or Area. Relative Scaling is only available for Converted Forecasts, Calculated Rows, External Data and item specific Explanatory Variables. Selecting Relative Scaling will display that component 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.
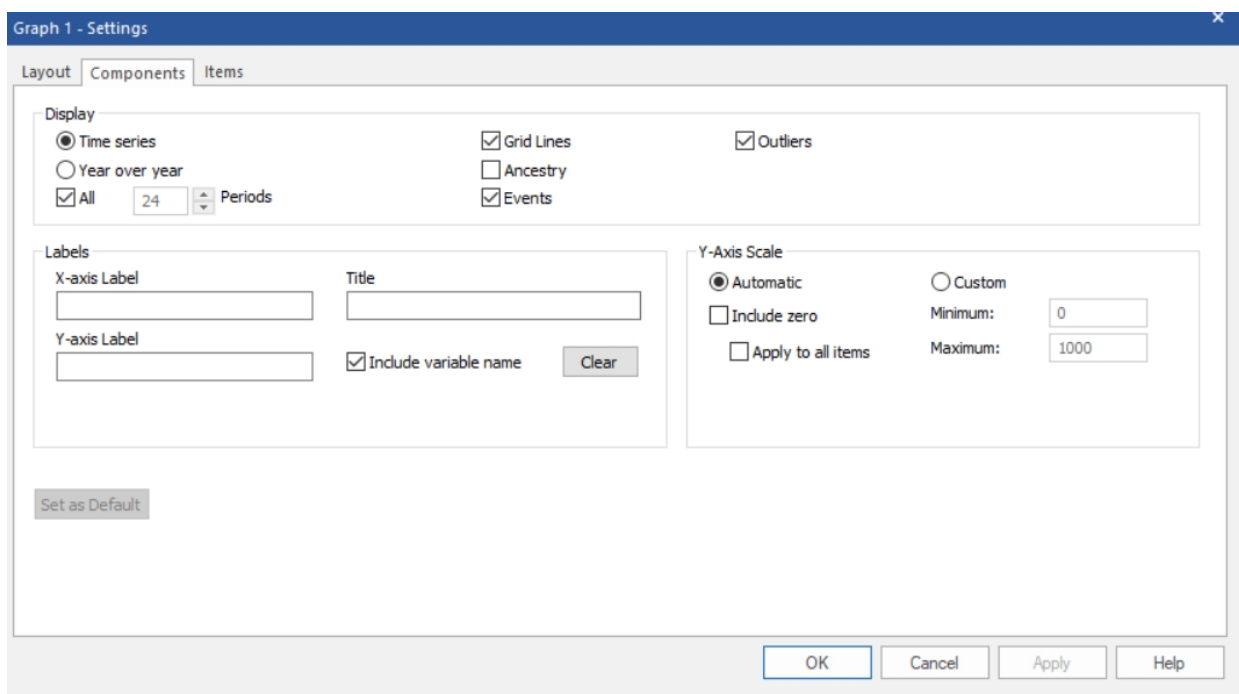
The Components tab includes three sections.

The *Display* section allows you to select the graph type.

Time series displays a graph that includes both the historic and forecast periods. The *All* option will display the entire data set in a non-scrollable 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 x-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.

Selecting *Zoom* on the Graph context menu or hamburger will activate the scrollable display. You can also switch between displaying the Graph view as a Time Series or as Year over Year using the context menu and hamburger menu.

The **Labels** section allows you to define labels for the X and Y axis and add titles to the graph. The Graph title appears within the Graph view, and the Title Bar label allows you to rename Graph 1 to a custom name. *Include variable name* will use the variable name followed by the description as the graph's subtitle. The **Clear** button erases the current labels within the Graph view (the Title Bar name will remain).

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 to all items.

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, using multiple graph views or layouts is a good way to do so.
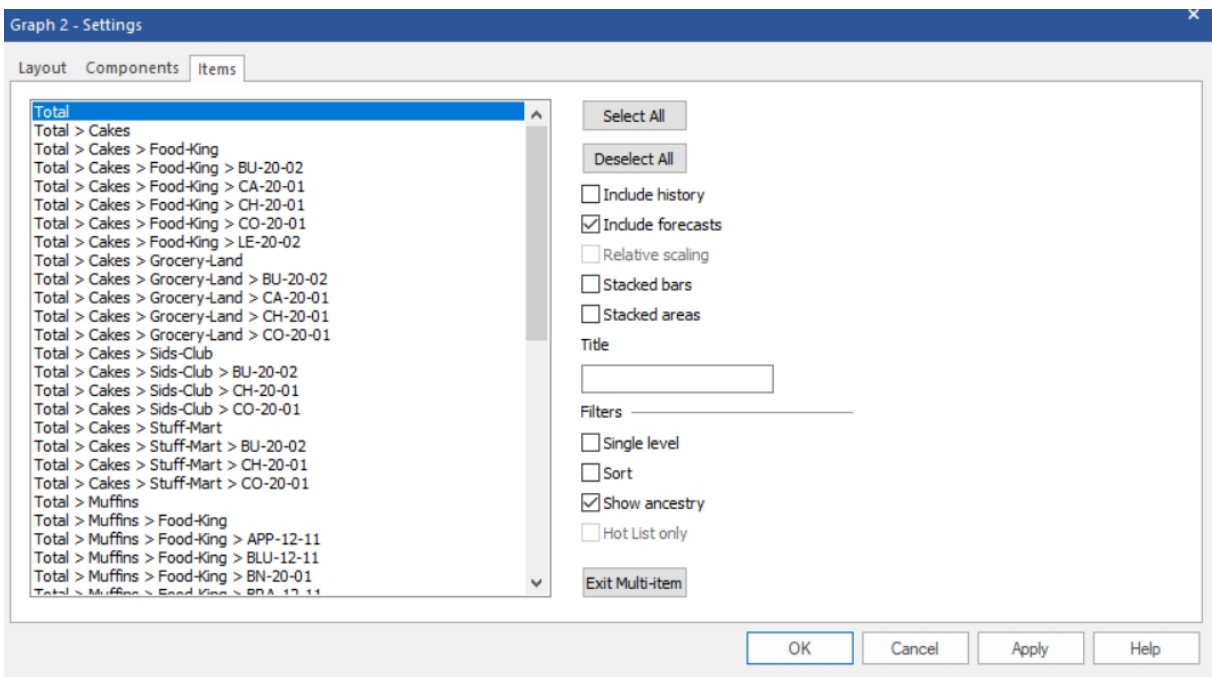
The Items tab allows you to graph multiple items and/or helper variables on a single graph. Simply select the items you want to display and then click the **OK** or **Apply** button.

Include history and *Include forecasts* indicate 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 *Stacked bars* and *Stacked area* options are useful if you want to understand the contribution of the graphed items to a total. If you use a stacked option, the height of the graph for a given period is sum of the graphed items.

You may enter a title for the graph in the *Title* text field.

The *Filter* options can be used to limit which items are displayed for selection, to sort the items or to specify if ancestry should be included for items displayed for selection.

The **Exit Multi-item** button returns the graph view back to the default, single item settings.

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, unless manual override mode has been selected on the Performance tab of the Options dialog box. Refer to *Understanding How Forecast Overrides are Allocated* for a discussion on how the reconciliation is accomplished.

To open the override view, click the **Override** icon on the **Home** tab or the **Reports** tab on the ribbon menu.

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.

If you use Excel Collaboration, the Excel Collaboration row also becomes available and can be commented and adjusted like the other override rows.

There is an Overrides tab on the 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 Overrides Tab.
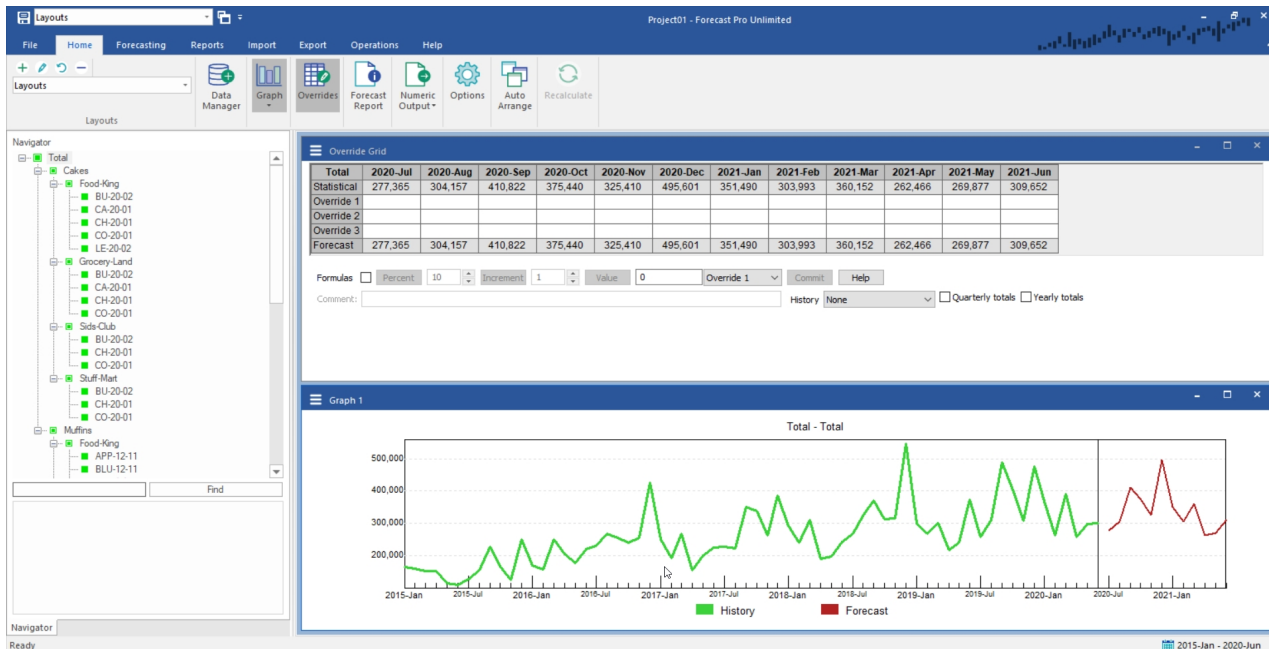
There are five rows in the example shown above.

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.

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 click the **Percent** button. If you want to adjust each selected point by a certain increment, enter the increment in the *Increment* edit box and click the **Increment** 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 click the **Value** 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 an override row cell 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 **Options** dialog box.

The Diagnostics View

The diagnostics view provides a variety of statistical displays to assist you in building and diagnosing custom models. To open the diagnostics view, click the **Diagnostics** icon in the Data Tools group on the **Forecasting** tab or select *Diagnostics* from the drop-down on the **Graph** icon on the **Home** tab or **Reports** tab.
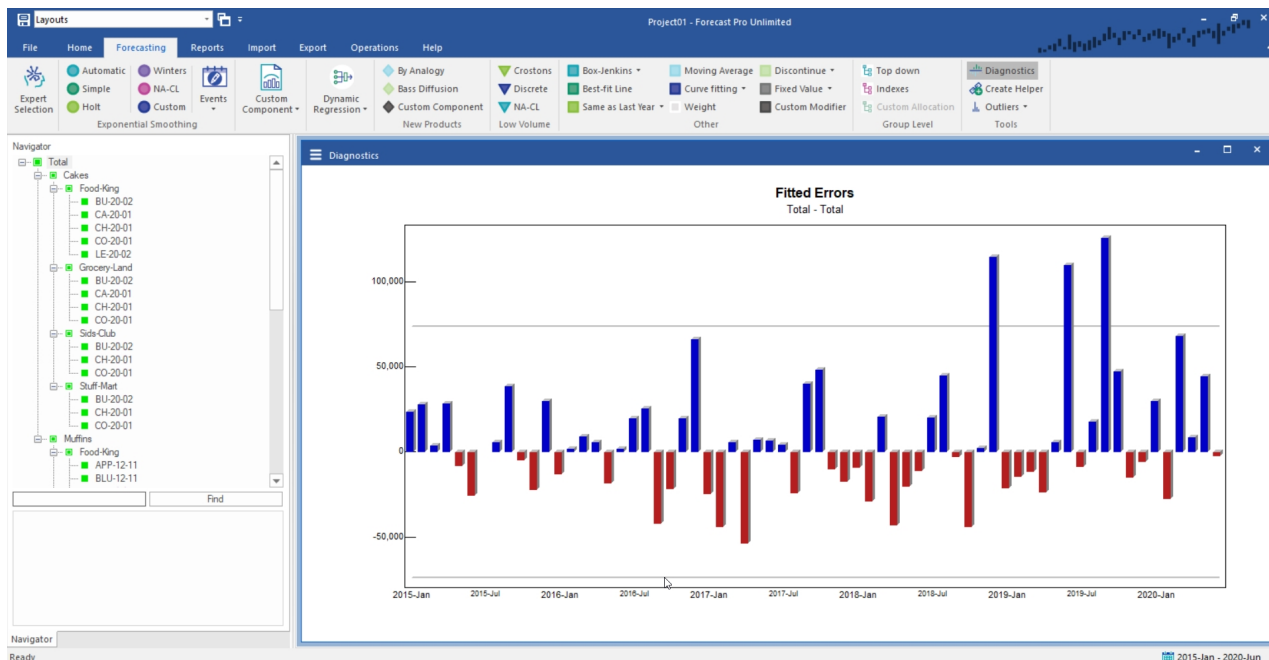
The default view is the *Fitted Errors* graph. Use this display to identify outliers or other problem areas of your data. The dashed horizontal lines mark the 2.5 and 97.5 percentiles. About one error in twenty extends beyond these limits.

You can right click on the *Fitted Errors* graph to bring up the Advanced Diagnostics context menu or use the hamburger menu in the top left of the view 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.
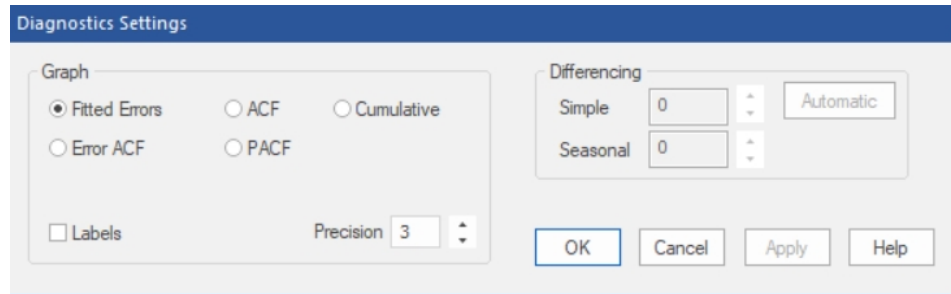
Use the **Automatic** button to display the differencing values selected by Forecast Pro's automatic Box-Jenkins algorithm. Use the partial autocorrelation functions (PACF) to identify the number of simple and seasonal AR terms. Refer to the *Forecast Pro Statistical Reference Manual* for further information 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 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 when loading archived override sets.

To open the Override Report view, click the **Override Report** icon on the **Reports** tab**.**

If you double-click an item on the Override Report the Navigator will jump directly to that item, so you can use the Override Report as the equivalent of a Hot List to navigate through the listed items. All global report views support this kind of navigation.

The content and format of the report is controlled using the Override Report Settings dialog box. This dialog box is invoked by selecting *Settings* on the **Override Report** icon drop-down, or by selecting **Settings** on the view's hamburger menu or 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 the **Archived Overrides** icon on the **Imports** tab.

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 lists items where outliers have been detected and/or corrected. To open the outlier report view, click the **Outliers** icon on the **Reports** tab or in the Tools group on the **Forecasting** tab.

The content and format of the report is controlled using the outlier report's context menu or hamburger menu. The context menu also allows you to print the currently displayed report and save it to Excel.

The Outliers tab of the **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. The Outlier Options are accessible by clicking Settings on the Outlier Report icon drop-down on the Reports tab and Forecasting tab, or on the view's context menu or hamburger menu. Consult [Outliers](#) for more details. The *Forecast Pro Statistical Reference Manual* also provides additional details on Forecast Pro's outlier detection methodology.

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 (though not necessary) to have the Numeric Output view open. To open the Numeric Output Report, click the **Numeric Output** icon on the **Reports** tab or select *Show* from the **Numeric Output** icon drop-down on the **Home** tab.

The Numeric Output Settings dialog box is used to specify the format and content of the Numeric Output Report and file. To open the Numeric Output Settings dialog box, select **Settings** from the drop-down on the **Numeric Output** icon on the **Home** tab or the **Reports** tab. You may also select Settings on the view's hamburger menu or the view's context menu to open the Numeric Output Settings dialog box.

The Numeric Output 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 and close the dialog box. Clicking the Apply button will also save your current settings for use with the current project,

but the dialog box will remain open so you can continue changing the settings. Clicking the Set as Default button will save your current settings as the default settings for all new projects.

The Numeric Output Settings dialog box contains three pages or tabs. We will discuss each tab in turn.
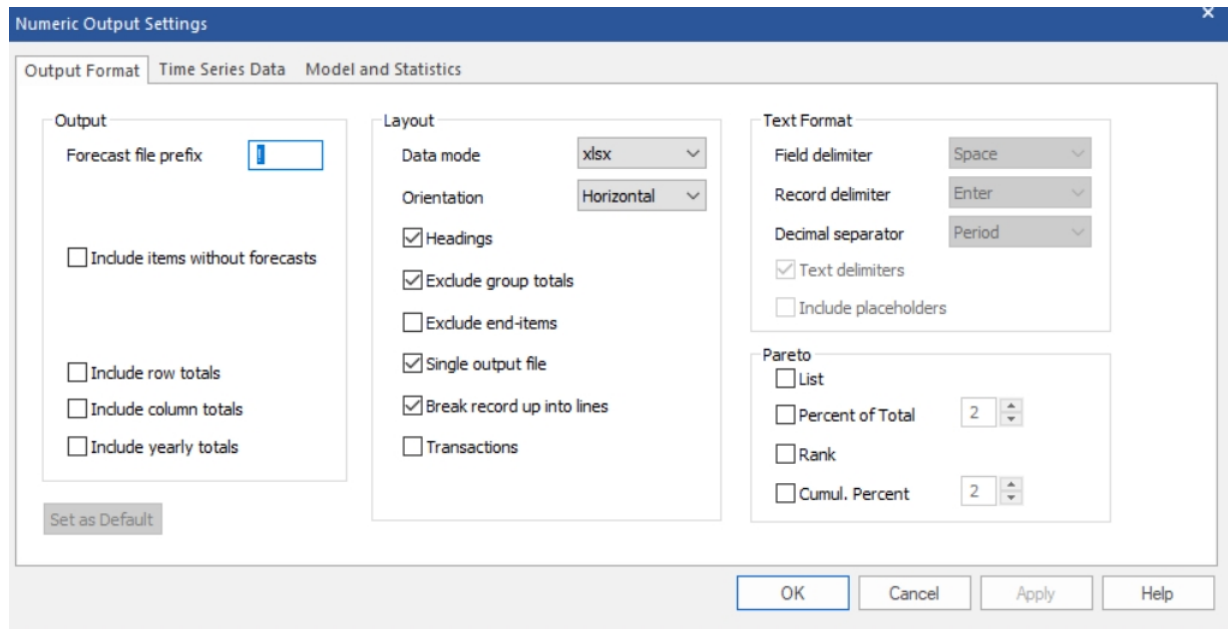
The Output Format tab is used to specify the basic layout for the Numeric Output file.

The Output Section

> *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.

> *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.

> *Include row totals.* If this option is selected, the Numeric Output file with include a column on the far right of the report, called Total, that contains the totals for all entries in each row.

> *Include column totals.* If this option is selected, the Numeric Output file with include a row on the bottom of the report, called Total, that contains the totals for all entries in each numeric column.

> *Include yearly totals.* If this option is selected, the Numeric Output file with include columns that contains the totals for each year. Each year column will be inserted to the right of the last period in each year.

The Layout Section

> The Layout section of this tab allows you to specify the default data mode, the orientation (row vs. column), the option to include group level output and the option 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.

Transactions. If selected, the data will be output using a transaction style layout (see Transactions Format). Transactional reports can only include time series data.

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* 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.

The Data Labels Section lets you specify what information you want to include for each item in the report.

Attributes. Include separate attribute fields for each level of the hierarchy. If you used attribute fields in your input data files to define groups, you will most likely want to select this option so that your output file will match your input file. Consult Using Attribute Fields to Define a Hierarchy 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. If *Include record type* is not active, then 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 lists the different types of data series you can include in the report.

History and Forecast Components. Check the boxes next to the information you want to include. Use the radio buttons under Seasonal Indexes to indicate if you want to include indexes for historic

periods and/or the forecast horizon. The spinners next to the items are used to specify the numeric precision of the selected series. If the precision is set to 0, the numbers will not include a decimal.

Custom Date range. By default, the Numeric Output Report will include all forecast periods if a forecast is selected and all history periods if history is selected. If you want to change the periods shown in the report, select *Custom date range* and use the spinners to define the beginning and ending dates.

If you are outputting a fixed width text file, you may set the column widths instead of the numeric precision by selecting *Width* in the *Adjust* radio buttons.

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, you may set the column widths instead by selecting the *Width* radio button instead on *Precision*. The spinners next to the *Within-Sample Statistics* selections and the *Safety Stock Basis* selections will apply to column width or numeric precision accordingly.

Most of the options are self-explanatory. However, we have noted some details below.

Model Properties. The Model Properties drop-down contains 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 Item Report View

The Item Report contains essentially the same information as the Numeric Output. 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. The purpose of the Item Report is to provide a way to look at custom views that may use different settings than those required for the Numeric Output. The Numeric Output settings are typically defined for how

the data needs to be exported within a forecasting process, while the Item Report is intended to be used interactively.

To open the Item Report view, click the **Item Report** icon on the **Reports** tab.

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 global report views support this kind of navigation.

The content and format of the report is controlled using the Item Report Settings dialog box. This dialog box is invoked by selecting **Settings** from the **Item Report** icon's drop-down, or by selecting **Settings** on the view's hamburger menu or from the Item Report view's context menu. The Item Report Settings dialog box functionality is a subset of the Numeric Output Settings dialog box. Please consult The Numeric Output View section of this chapter for details

Saving Your Work

Forecast Pro *Forecast Projects* allow you to save your forecasting session so that you can return later and pick up where you left off or share the session with others. The *Forecast Project* saves the data, forecasts, overrides, Hot List, modifiers and all other information associated with the project as it exists when you save it. This data can also be exported to file.

Additionally, Forecast Pro can export various types of output data including the fifteen following reports, which can be reviewed and customized on the **Reports** tab and exported from the **Export** tab.

> *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, for the currently displayed items including any filters applied or just for the Hot List. If you will be importing forecasts into another application, these are typically 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.

> The *Item Report* can be 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. You can include output for either all items forecasted, for the currently displayed items including any filters applied or just for the Hot List.

> The *Outlier Report* can be saved to Excel. This report lists items where outliers have been detected and/or corrected.

> The *Override Report* can be saved to Excel. This report lists items where overrides and/or comments have been made.

Each of the above reports may be exported for all items in the project, for items on the Hot List or for only items included in the Current Display (including filters). The following sections discuss forecast projects and the output files listed above 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 will 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 the **File** tab and then select *Save* on the left bar**,** click the **Save** icon on the far left of the default Quick Access Toolbar or simply use Ctrl+S. To open a forecast project, you select the **File** tab and then select *Open* on the left bar or simply use Ctrl+O.

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 where you left off.

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.
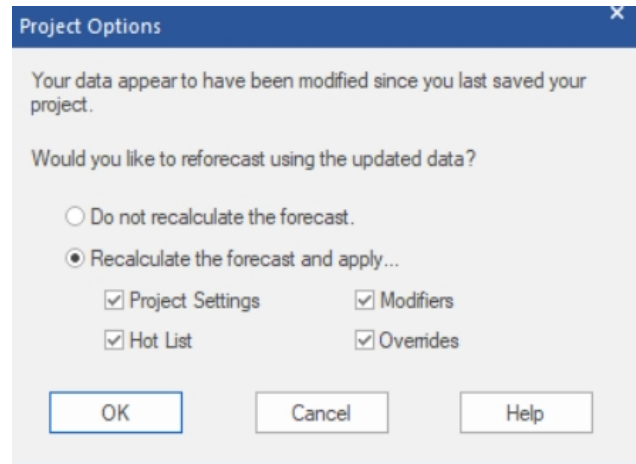
If you open a project and data files have changed, a dialog box will appear allowing you to either generate new forecasts using the updated data or restore the conditions present when the project was saved (i.e., do not recalculate the forecasts using the new data).

Project files

The project file is usually 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 Unlimited 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, select *FPProj* as the file type in the Save as and 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 all items forecasted, what is currently displayed (including filters), the currently selected item or the Hot List. If you will be importing the forecasts into another application these are the files you will likely want to use.

You can specify the format and content of Numeric Output File in Numeric Output Settings. The operation of the Numeric Output Settings dialog box is described in [The Numeric Output View](#).

You can preview the contents of the current Numeric Output File by clicking the Numeric Output icon on the **Reports** tab or selecting *Show* from the **Numeric Output** icon drop-down on the **Home** tab.

You can save a Numeric Output file on the **Export** tab. First, on the left of the tab, select if you want to export the report for All Items, the Hot List or the items Currently Displayed in the report. Then, click the **Numeric Output** icon to save the report.

You may open Numeric Output settings without previewing the report by selecting Settings on the drop-down of the **Numeric Output** icon on the **Reports** tab or the **Export** tab.

Forecast Reports

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 defined by the Graph 1 settings.

You can specify the format and content of Forecast Report using Forecast Report Settings and Graph 1 Settings. You can open the settings dialog from the hamburger or context menu on the Forecast Report and Graph 1 views. The operation of the Forecast Report Settings and Graph Settings dialog boxes are described in [The Forecast Report View](#) and [The Graph Views](#).

You save a Forecast Report from the **Export** tab. First, select if you want to export the report for All Items, the Hot List or the items currently included in the report. Then click the small **Forecast Report** icon.

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. You can use the context menu or hamburger menu to control the items included in the report (e.g. all items, the currently selected item or only items in the current hot list). This is true for all global reports.

To open an Item Report, you click the **Item Report** icon on the Reports tab. You save an Item Report from the **Export** tab. First, select if you want to export the report for All Items, the Hot List or the items Currently Displayed in the report. Then click the small **Item Report** icon.

You can specify the format and content of the report using Item Report Settings. You can open the settings dialog from the hamburger or context menu on the Item Report. The operation of the Item Report Settings dialog box is described in [The Item Report View](#).

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. You can use the context menu or hamburger menu to control the items included in the report (e.g. all items, the currently selected item or only items in the current hot list). This is true for all global reports.

To open an Outlier Report, click the **Outlier Report** icon on the **Reports** tab. You save an Outlier Report from the **Export** tab. First, select if you want to export the report for All Items, the Hot List or the items Currently Displayed in the report. Then, click the small Outlier Report icon.

The content and format of the report is controlled using the Outlier Report's context menu. Please see The Outlier Report View for details.

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 with a time stamp and the user who made the override/comment. You can use the context menu or hamburger menu to control the items included in the report (e.g. all items, the currently selected item or only items in the current hot list). This is true for all global reports.

To open an Override Report, click the **Override Report** icon on the **Reports** tab. You save an Override Report from the **Export** tab. First, select if you want to export the report for All Items, the Hot List or the items Currently Displayed in the report. Then, click the small **Override Report** icon.

You can specify the format and content of the report using Override Report Settings. You can open the settings dialog from the hamburger or context menu on the Override Report. The operation of the Override Report Settings dialog box is described in The Override Report View.

Customizing Your Forecasts

By default, Forecast Pro will use expert selection, an Artificial Intelligence (AI) driven engine to automatically select a forecasting model for each item on the Navigator. The expert selection option works extremely well and is the method of choice for most Forecast Pro users. The expert selection algorithm is described in the *Forecast Pro Statistical Reference Manual*.

Alternatively, you can select alternative models and/or forecasting options. The tools you need to customize forecast models are on the **Forecasting** tab. Forecast modifiers are added to items on the Navigator as you create custom models. These modifiers are used by Forecast Pro to document your customizations and generate models according to the customized specification. Note that the functionality available on the **Forecasting** tab can also be accessed from the Navigator's context menu.

This chapter documents the model options and tools available on the **Forecasting** tab, the associated forecast modifiers, and the operation of the model specification dialog boxes. Please note that in all cases, the described functionality is applied to the item selected in the Navigator.

Expert Selection

By default, Forecast Pro will use expert selection, an Artificial Intelligence (AI) engine, to automatically select the appropriate forecasting method. In this case, the item will not have a forecasting modifier. If an item on the Navigator does have a modifier that specifies a statistical model to be used instead of Expert Selection, when you click the **Expert Selection** icon on the **Forecasting** tab or select *Forecasting>Expert Selection* on the Navigator's context menu, the forecasting modifier(s) will be removed, and Forecast Pro will use expert selection to automatically select the appropriate forecasting method.

On The Performance Tab of the Options dialog, you can specify methodologies that expert selection should not consider. These methodologies include: ARIMA, (Box-Jenkins), seasonal simplification, the NA-CL model and machine learning.

Machine Learning

Automatic Univariate Machine Learning

Machine learning is available as an automatic univariate model, which may be included in expert selection, or as a custom model. Like exponential smoothing and Box-Jenkins, machine learning models leverage information inherent in the historic data (seasonality, trend) to create variables (features in

machine learning). While the automatic machine learning models rely on only this information, custom machine models allow you to include additional explanatory variables and/or events.

Click the **Machine Learning** icon to use an automatic univariate machine learning model.

The modifier associated with Automatic Machine Learning is:

\ML. Use automatic machine learning.

Select *Expert with ML* from the **Machine Learning** icon drop-down to include machine learning in expert selection for the selected item. This is useful when you want to consider machine learning in expert selection for only a subset of items. This is only available when you exclude machine learning from Expert Selection globally via the Performance tab of the Options.

\MLES. Include machine learning in expert selection and apply the model that is automatically selected.

Custom Machine Learning Models

To specify a custom machine learning model, Select *Manage* from the **Machine Learning** icon drop-down to open the Machine Learning dialog.

Name: Forecast Pro names each of your custom machine learning model specifications and saves them in the project's database. These named specification sets provide a convenient way to apply the same machine learning 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.

All custom machine learning models will include the same features used in the automatic univariate models (e.g. seasonality/interventions). The custom dialog allows you to include additional variables.

Events: The Events display lists all available event schedules. Check the event schedules you want machine learning to consider as features in the model. Because the event codes have no natural order (that is event codes 1 and 2 are simply two different events), Forecast Pro creates a separate feature for each event code in an event schedule. The machine learning algorithm will determine which event schedules and event codes should be used in the model.

Explanatory variables: The Explanatory variables field lists all global and item-specific explanatory variables available in the project. Check the variables that you want Forecast Pro to consider as features in your machine learning model. Do not include variables that are categorical, or unordered. If you have an unordered or categorical variable, include it as an event schedule, not an explanatory variable. Note that the machine learning algorithm will consider all checked variables, but it will only use a variable if it improves the machine learning forecasts.

Auto Extend: Check the *Auto Extend* checkbox 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.

Parameters:

Automatic: Check the *Automatic* checkbox if you want Forecast Pro to automatically select the *Maximum Tree Depth* and *Number of Trees*. The parameters are described below, in the Machine Learning Methodology.

If you uncheck the *Automatic* checkbox, the *Maximum Tree Depth* and *Number of Trees* spinners are operational. Consult Machine Learning Methodology for more details on the parameters described below.

Maximum Tree Depth: The tree depth is the length of the longest path to a leaf or forecast node. Use the spinner to select a maximum tree depth. Please note that as you increase the tree depth, the in-sample model fit will improve, but the forecasts may be less accurate. You should consider using out-of-sample statistics to choose an appropriate maximum tree depth.

Number of Trees: Use this spinner to choose the number of trees to include in the ensemble model. As you increase the number of trees, the in-sample model fit will improve, but the forecasts may be less accurate. You should consider using out of sample statistics to choose an appropriate number of trees.

\ML=name. Use a custom machine learning model with the specifications defined in *name*.

Machine Learning Methodology

Forecast Pro's machine learning methodology uses extreme gradient boosted decision trees. A decision tree applies a set of logical rules to a set of features to "decide" an appropriate forecast value. Below is an example of a decision tree.

In this case, there is only one feature (month), and tree depth is two. The tree depth is the length of the longest path to a terminal node. To generate a forecast, the model starts at the top of the tree and uses the feature values to move through the tree until a terminal ("leaf") node is reached. Here, the leaf nodes are the "Demand" nodes. For example, suppose it is June, or month = 6.
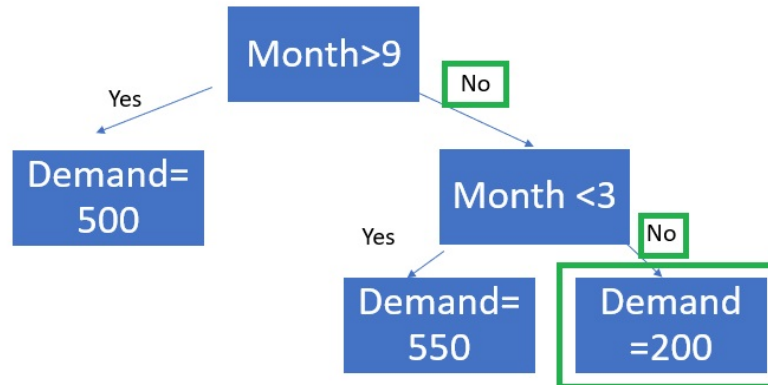
The demand forecast for June is 200.

Extreme gradient boosting is an ensemble approach. An **ensemble** model combines multiple models. Extreme gradient boosting generates multiple decision trees and then the forecast is equal to the average forecast across all trees. **Boosting** algorithms add models sequentially. In each iteration, the new tree improves on the weakness of the previous trees. Extreme gradient boosting finds the tree structure and parameters that optimize a specified objective function, typically Mean Square Error. Forecast Pro's implementation uses XGBoost, an open-source library. Please consult XGBoost documentation for more details on how extreme gradient boosted trees work.

XGBoost requires several key inputs, including features, maximum tree depth and the number of trees the ensemble model should include. These inputs are important determinants of forecast accuracy and manually selecting them can be difficult. When using automatic univariate machine learning models, Forecast Pro uses its own rule-based algorithm to decide which features to automatically build and then automatically determines the maximum tree depth and number of trees using a combination or rules and out-of-sample holdout testing.

In addition to the Forecast Pro generated features used in the automatic univariate model, *custom machine learning* models may also include event schedules and explanatory variables as features in a boosted decision tree ensemble model. You can automatically train custom machine learning models using the same AI-driven logic that powers Forecast Pro's automatic machine learning univariate models, or you can specify a specific boosted decision tree structure.

Machine Learning Forecast Report

Because ensemble models involve multiple models, it is difficult to summarize the impact of an individual feature on demand in the same way that you can with coefficients for dynamic regression or indexes for exponential smoothing. Instead of showing the size or direction of the impact of a feature on demand, Forecast Pro shows the relative importance of a feature in the ensemble model.

Model Details

User Defined \ML=MachLearn02			
Machine Learning (Automatic)			
ML(MachLearn02)			
Tree Depth Used	4		
Number of Trees	50		

Feature Importance	Gain	Coverage	Frequency
Feature and Display	41.01%	17.94%	11.63%
Price	37.06%	19.58%	12.53%
Interventions	5.82%	11.78%	19.69%
Feature	3.61%	13.06%	12.53%
Seasonal period	3.41%	10.45%	10.74%
One period lag	2.91%	6.13%	8.05%
Display	1.95%	6.98%	7.38%
Seasonal lag	1.85%	9.12%	9.40%
One year moving average	1.21%	3.17%	6.04%
_Stock Outs - Stock Out	1.17%	1.79%	2.01%
_Stock Outs -	0.00%	0.00%	0.00%

The above Model Details section of the Forecast Report is for a custom model with automatic parameter selection. Forecast Pro has used a maximum tree depth of 4 and 50 trees. The model includes 4 explanatory variables (Price, Feature and Display, Feature, Display) and one Event Schedule (_Stock_outs) with 2 event codes (Stock Out and No Event). The remaining variables were automatically generated by Forecast Pro.

There are three importance measures displayed for each feature: **Gain**, **Coverage** and **Frequency**. Gain measures the percentage of overall model improvement or accuracy that is attributable to a feature. Coverage shows the percentage of observations that are forecasted at a leaf node that are decided by a feature. Frequency is the percentage of tree splits associated with a feature. Gain is widely considered to be the most relevant feature importance measure for forecasting. For more details on feature importance computations, please consult the XGBoost documentation.

Exponential Smoothing

Automatic Exponential Smoothing

Select *Automatic* from the **Exponential Smoothing** icon drop-down or simply click the icon for Forecast Pro to automatically select an exponential smoothing model.

The modifier associated with Automatic Exponential Smoothing is:

\EXSM. Use the automatic fitting exponential smoothing model.

Simple Exponential Smoothing

Select *Simple* on the **Exponential Smooth** icon drop-down to use a simple exponential smoothing model.

The modifier associated with Simple Exponential Smoothing is:

\SIMPLE. Use the simple exponential smoothing model.

Holt Exponential Smoothing

Select *Holt* on the **Exponential Smoothing** icon drop-down to use the Holt exponential smoothing model.

The modifier associated with Holt Exponential Smoothing is:

HOLT. Use the Holt exponential smoothing model.

Winters Exponential Smoothing

Select *Winters* on the **Exponential Smoothing** icon drop-down to use the Winters exponential smoothing model.

The modifier associated with Winters Exponential Smoothing is:

WINTERS. Use the Winters exponential smoothing model.

NA-CL Exponential Smoothing

Select *NA-CL* on the **Exponential Smoothing** icon drop-down or click the **NA-CL** icon in the Low Volume group to use a non-trended Additive Seasonality Constant Level exponential smoothing model. Note that NA-CL is included in the Low Volume group because it works well for items with seasonal, low volume demand.

The modifier associated with NA-CL Exponential Smoothing is:

NA-CL. Use the non-trended Additive Seasonality Constant Level exponential smoothing model.

Custom Exponential Smoothing

Selecting *Custom* on the **Exponential Smooth** icon drop-down on the **Forecasting** tab or selecting *Forecasting>Exponential Smoothing* 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. 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 include:

\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.

\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).

\SS. Use Forecast Pro's automatic identification procedure to determine whether to use seasonal simplification and the appropriate bucket size.

\SS=n. Use seasonal simplification with n periods for each seasonal index or "bucket".

Event Models

Event 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 code (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.

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.

You invoke the Event Manager by clicking the **Events** icon on the **Forecasting** tab or selecting *Manage events* on the Graph 1 view's context menu.

The Event Manager works in conjunction with the Graph 1 view. If the Graph 1 view is not open, Forecast Pro will automatically open it when you click the **Event** icon on the **Forecasting** tab.

The Manage Events dialog box is shown above. The Date drop-down at the top of the box is used to assign event codes to the current schedule. To specify the date to which you want to assign an event code, you select a historic point on Graph 1. You can then use the drop-down to the right of the date to associate the selected date 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, not just the dates with assigned events. You may select periods in the display grid and use the drop-down to assign event codes instead of selecting periods on the graph.
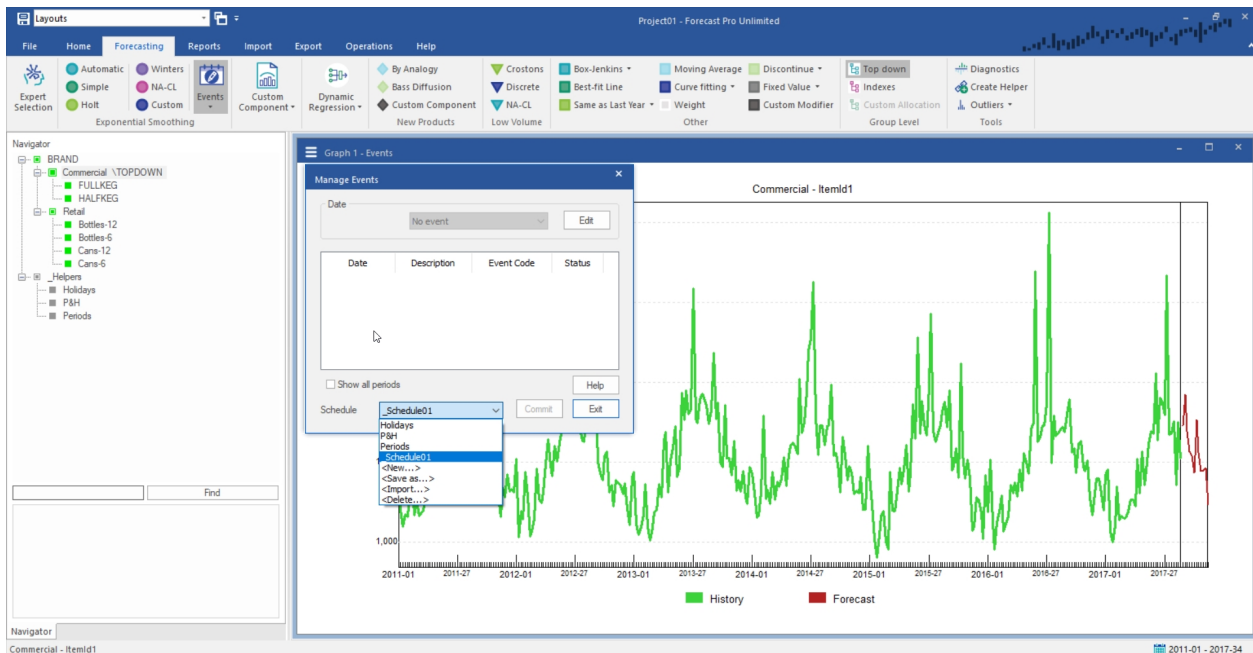
The Schedule drop-down is used for several purposes. When the drop-down is closed it displays the name of the event schedule associated with the current item. Opening the drop-down displays the currently available event schedules and allows you to select one. It also provides options to disassociate the currently selected schedule from the currently displayed item and create a new schedule (*New*), rename the current schedule (*Save as*), import events from other schedules into the current schedule so the two schedules are merged and delete the current schedule from the Forecast Pro database.

Clicking the **Edit** button on the Event Manager or using the Event Codes option in the Events icon drop-down will invoke the Manage Event Codes dialog box, which allows you to define new event codes and edit existing ones. You can select an event code using the spinners or by clicking on the event code 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. You will be warned if the event code is being actively used and given the option to save changes to the event codes. Any new event code defined 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.

Finally, the **Events** icon drop-down lists all event schedules associated with the project. Select a schedule from the drop-down to apply the modifier *EVENT=_X*, where X is the selected schedule or helper, to the selected item on the Navigator.

For examples of constructing event variables and building event models, consult the *Building Event Models* lesson in *Forecast Pro Unlimited Tutorials*. For statistical details, consult the *Forecast Pro Statistical Reference Manual*.

The following modifiers are used for Event Models:

EVENT=_X. Use an event model with helper _X.

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=.

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=.

Custom Component Model

The custom component model includes each of the components found in a custom exponential smoothing model: 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 can either let Forecast Pro estimate the final value for a given component or 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. Because custom component models are particularly useful for new products, Forecast Pro includes a **Custom Component** icon in the New Product group in addition to the primary icon on the **Forecasting** tab.

Clicking either **Custom Component** icon or selecting *Forecasting>Component Model > Manage* 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* 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.

Click OK or Apply on the Component Model Settings dialog box to apply the model shown in the name drop-down to the item selected on the Navigator. Alternatively, you may open the drop-down on the large **Custom Component** icon on the **Forecasting** tab and select a model name.

The following modifier is associated with the custom component model.

\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.).

Clicking the **Dynamic Regression** icon on the **Forecasting** tab or selecting *Forecasting>Dynamic Regression>Manage* 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 database. 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 Explanatory Variable Files 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 lists 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. In this mode, you select the explanatory variables to include, 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): If this option is selected, the program automatically determines the dynamic terms to include. The user still needs to specify the explanatory variables to include. Automatics 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. 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, open the Options dialog box by clicking the Options icon on the **Home** ribbon tab and then go to the Advanced Controls tab on the Options dialog box.

Click OK or Apply on the Dynamic Regression dialog box to apply the model shown in the name drop-down to the item selected on the Navigator. Alternatively, you may open the drop-down on the **Dynamic Regression** icon on the **Forecasting** tab and select a model name.

The following modifier is associated with the dynamic regression model.

\DR=name. Use a dynamic regression model with the specifications defined in *name*.

New Product Models

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.

Clicking the **By Analogy** icon on the **Forecasting** tab or selecting *Forecasting>New Products>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 launch, 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 June 2021, Forecast Pro would use the Forecast by analogy approach until your historic data included July 2021.

Click OK or Apply on the Forecast by Analogy Model Settings dialog box to apply the by analogy model.

The modifiers associated with Forecast by Analogy are:

\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.

\ANALOG=_X(A,B). Use the by analogy model with a specified launch total of A and a specified launch horizon of B.

\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.

\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.

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 0.3 and 0.5.

Clicking the **Bass Diffusion** icon or selecting *Forecasting>New Products>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 *Generate coefficients automatically*, you'll need to supply the model coefficients and potential market size.

The modifiers associated with the Bass diffusion model are:

\BASS. Use the Bass diffusion model with automatically generated coefficients and potential market size.

\BASS(p,q,m). Use the Bass diffusion model with user defined parameters p, q and m.

Custom Component Model

Click the **Custom Component** icon to invoke the Component Model Settings dialog box. See Custom Component Model for a detailed description of the dialog box.

The following modifier is associated with the custom component model.

\COMP=name. Use a custom component model with the specifications defined in *name*.

Low Volume Models

Intermittent Data

Click the **Croston** icon or the **Forecasting** tab or select *Forecasting>Low Volume>Croston* on the Navigator context menu to use Croston's intermittent demand model.

The following modifier is associated with the Croston's intermittent demand model:

\INTER. Use the Croston's intermittent demand 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.

Discrete Data

Click the **Discrete** icon or the **Forecasting** tab or select *Forecasting>Low Volume>Discrete* on the Navigator context menu to use simple exponential smoothing model and base confidence limits on one of the discrete distributions (Poisson or negative binomial).

The following modifier specifies a discrete model:

\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).

NA-CL Exponential Smoothing

Click the **NA-CL** icon to use a Non-trended Additive Seasonality Constant Level exponential smoothing model. This model was designed for data with a selling season.

\NA-CL is the modifier for NA-CL exponential smoothing.

Other Forecasting Options

Box-Jenkins (ARIMA)

Clicking the **Box-Jenkins** icon in the Other group instructs Forecast Pro to automatically select a Box-Jenkins (ARIMA) based on out-of-sample fit combined with some Artificial Intelligence expert system rules.

BJ is the modifier for automatic Box-Jenkins model.

Selecting *Custom* from the drop-down on the **Box-Jenkins** icon on the **Forecasting** tab or selecting *Forecasting>Box-Jenkins* on the Navigator's context menu will open the 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 custom Box-Jenkins modifiers are:

***ARIMA(p,d,q).* Use a non seasonal Box-Jenkins model with model orders p, d and q.

***ARIMA(p,d,q)*(P,D,Q).* Use a seasonal Box-Jenkins model with model orders p, d, q, P, D and Q.

***CONST*. Include a constant intercept in the Box-Jenkins model.

Same as Last Year

Click the **Same as Last Year** icon on the **Forecasting** tab to set each forecasted value to equal the value for the same period last year (either history or forecast as the case may be).

Selecting *Plus Percent* or *Plus increment* from the drop-down on the **Same as Last Year** icon will open the dialog box shown below.
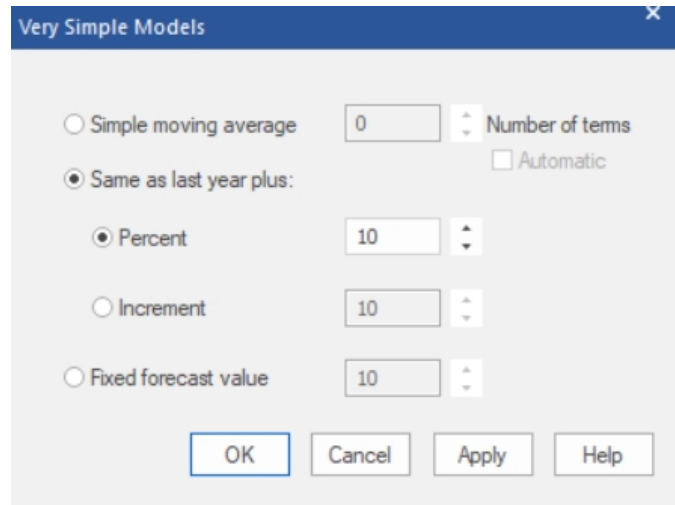
Same as last year (selected by default when opened from drop-down) sets each forecasted value to equal the value for the same period last year (either history or forecast as the case may be). The *Percent* option (selected by default when *Plus percent* is selected from the drop-down) sets each forecasted value to equal the value for the same period last year plus the specified percentage. The *Increment* option (selected by default when **Plus increment** is selected from the drop-down) sets each forecasted value to equal the value for the same period last year plus the specified increment.

Selecting *Plus delta percent* from the drop-down on the **Same as Last Year** icon will 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.

Selecting *Plus delta increment* from the drop-down on the **Same as Last Year** icon will 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.

The modifiers associated with Same as Last Year models are:

\SALY. Set each forecasted value to equal the value for the same period last year (either history or forecast as the case may be).

\SALYP=n. Set each forecasted value to equal the value for the same period last year plus the specified percentage (*n*).

\SALYI=n. Set each forecasted value to equal the value for the same period last year plus the specified increment (*n*).

\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.

\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.

Discontinue

Click on the **Discontinue** icon in the Other group on the **Forecasting** tab to discontinue the item selected on the Navigator. When an item is discontinued, Forecast Pro no longer generates a forecast for that item. This means that you cannot override the forecast in the override grid, and the item will not be

included in the Numeric Output, by default (though this can be changed in the Numeric Output Settings or in the Numeric Output context menu).

Select *Discontinue after* on the **Discontinue** icon drop-down to open a dialog box where you can specify the date after which the selected item should be discontinued.

Select *Set to zero* on the **Discontinue** icon drop-down to set an item's forecast to zero for the entire forecast horizon. When you set the forecast to zero, Forecast Pro is generating a forecast. This means that you can still override the forecast in the override grid and the item will be included in the Numeric Output.

Select *Set to zero after* on the **Discontinue** icon drop-down to open a dialog box where you can specify the date after which the forecast for the selected item should be set to zero.

The discontinue modifiers are:

\DISCONT. Treat the specified item as discontinued. Discontinued items are not forecasted but their historic demand does go into group-level aggregates.

\DISCONTAFT(Year, Period). Treat the specified item as discontinued after the specified date.

\FIXED=0. Set the item's forecast for the entire forecast horizon.

\ZEROAFT(Year, Period). Set the item's forecast to zero after the specified date.

Best-fit Line

Select *Best-fit Line* icon on the **More** icon drop-down on the **Forecasting** tab or select *Forecasting>Curve Fit>Straight Line* on the Navigator's context menu to fit a straight line to the data.

\CFIT=LINE is the modifier to fit a straight line to the data set.

Moving Average

Select *Moving Average* on the **More** icon drop-down to open the dialog box shown below.

Specify the number of terms used or check *Automatic* to have Forecast Pro automatically select the number of terms. Click OK or Apply to apply the moving average model.

The modifiers associated with moving averages are:

\SMA=n. Use an n-term simple moving average.

\SMA. Use a simple moving average where Forecast Pro automatically selects the number of terms to use.

Curve Fitting

Forecast Pro offers the following curve fitting approaches: Best-fit line, Quadratic, Exponential and Growth. Each of these methods can be selected from the *Curve fitting* drop-down on the **More** icon drop-down or from *Forecasting>Curve Fit* on the Navigator's context menu.

Select *Auto* from the *Curve fitting* drop-down on the **More** icon drop-down or select *Forecasting>Curve Fit>Auto* on the Navigator's context menu to use the curve fitting routine that best fits the historic data.

The modifiers for moving average models are:

\CFIT. Use the curve fitting routine that best fits the historic data.

\CFIT=LINE. Fit a best-fit straight line to the data set.

\CFIT=QUAD. Fit a quadratic curve to the data set.

\CFIT=EXPO. Fit an exponential curve to the data set.

\CFIT=GROW. Fit a growth curve (S-curve) to the data set.

Weight (Weighting Transformations)

A weighting transformation is commonly used to de-seasonalize your demand 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 a weighting transformation, you must create a *helper* variable containing the weights. 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.

To apply a weight transformation, select *Weights* on the **More** icon drop-down on the **Forecasting** tab or select *Forecasting>Weights>Select* on the Navigator's context menu to open the dialog box shown below.

Select the helper variable with the weights you want to apply and click OK to apply the forecast modifier.

The following modifier is used to specify the weighting transformation:

\WGT=_X. Use a weighting transformation. _X is the name of the helper variable 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 de-weighted 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.

Fixed Value

Fixed forecast value sets the forecast to the specified value.

Select *Fixed Value* on the **More** icon drop-down on the **Forecasting** tab to open the Very Simple Models dialog box with *Set forecasts to* selected.

Enter the desired forecast value, specify a date, and click OK or Apply to apply the model. By default, the date spinner is set to the first period of the forecast horizon.

The modifiers associated with fixed value forecasts are:

\FIXED=n. Set each forecasted value to *n*. This modifier is used if the "after" date on the spinner is the first forecast period.

\FIXEDAFT=(Year, Period, n). Set each forecasted value which falls after the year and period specified to *n*.

Custom Modifiers

Select *Custom* on the **More** icon drop-down on the **Forecasting** tab or select *Custom Modifier* on the Navigator's context menu to enter modifiers explicitly.
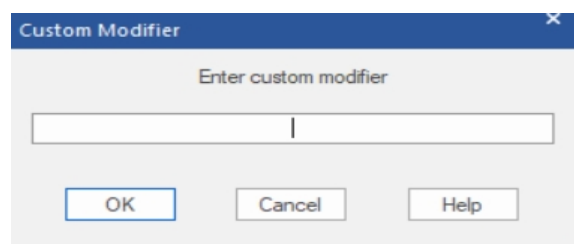
All modifiers may be entered using the Custom Modifier dialog box. There are also some modifiers that must be entered using the Custom Modifier box, including:

\LOWER=l. Set the lower confidence limit equal to the value *l*. *l* must be between 0.1 and 50.0 inclusive.

\UPPER=u. Set the upper confidence limit equal to the value *u*. *u* must be between 50.0 and 99.9 inclusive.

\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.

\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.

\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.

\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.

Group Level (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 all groups and items.

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 or leverage group data structure. These modifiers can only be used on group-level data.

Top-down

To apply top-down reconciliation, click the **Top Down** icon in the Group Level group on the **Forecasting** tab or select *Forecasting>Group Level>Top-down* from the Navigator's context menu.

This procedure begins by preparing forecasts for all groups and items. Then it proportionally adjusts the nested lower level item and group forecasts to sum to the selected group's forecast. Please consult Reconciling the Hierarchy for more details.

Top-down reconciliation uses the following modifier:

\TOPDOWN*.* Directs Forecast Pro to perform top-down reconciliation for all members of the indicated group.

Custom Allocation

To apply custom allocation, click the **Custom Allocation** icon in the Group Level group on the **Forecasting** tab or select *Forecasting > Group Level > Custom Allocation* from the Navigator's context menu.

Note that Custom Allocation is only available when you have defined the proportions in a file, table or query and read it into Forecast Pro using the Definitions row in the Data Manager, and it is only available for the group(s) one hierarchy level above the level at which the proportions are defined. Consult Definitions: Custom Allocations in *Setting up Your Optional Data* 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 are using Custom Allocation. Please see Custom Allocation for more details.

Custom Allocation uses the following modifier:

\CALL. Directs Forecast Pro to use user-defined proportionalities to allocate the group-level forecast down to the next level.

Indexes

Click the **Indexes** icon in the Group Level group on the **Forecasting** tab or select *Indexes* from the Navigator's context menu to use a Group level seasonal index for all items within the group. Please see Indexes in Working with Hierarchies for more details.

Indexes uses the following modifier:

\INDEXES. Directs Forecast Pro to calculate the seasonal indexes at the indicated group level and use them to de-seasonalize the histories for the nested items and groups. The de-seasonalized series are then forecasted and as a final step re-seasonalized. 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 re-estimated for each item of the \INDEXES group.

Data Tools

Diagnostics

Click the **Diagnostics** icon in the Data Tools group on the **Forecasting** tab or select *Diagnostics* from the drop-down on the **Graph** icon on the **Home** tab or **Reports** tab to open the Diagnostics View.

!

Consult [The Diagnostics View](#) for details on available functionality.

Create Helper

Helper variables are required for event modeling, forecast by analogy and the weighting transformations. Helper variables may also be used in Custom Component, Dynamic Regression, and Machine Learning models.

Event schedules are created interactively using the Event Manager, or 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. Event Schedules are not created with the **Create Helper** icon.

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 forecast by analogy and weighting transformation helper variables using the **Create Helper** icon 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 create a helper variable within Forecast Pro, use the Navigator to select the item which you want to base the helper variable upon, and then click on the **Create Helper** icon in the Data Tools group on the **Forecasting** tab or select *Helper variable>Create* on the Navigator's context menu to open the Create Helper Variable dialog box.

To create the helper variable, 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. You may rename the helper variable 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 category-level group. You would select the category-level group, create the helper variable and then import the indexes from the helper when you build a custom component model for the SKU. See Custom Component Models for more details.

If you wish to delete a helper variable select the helper variable on the Navigator, then right click to invoke the Navigator's context menu and then select *Helper variable>Delete.* Please note that once you delete a helper variable, it cannot be restored.

Outliers

Click on the **Outliers** icon in the Data Tools group on the **Forecasting** tab to open the Outliers Options dialog box. This is identical to the Outliers tab in the Options.

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 or 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).

Please consult The Outliers Tab in Forecast Pro Options for more details on the settings available in the Outliers tab.

You may override project level-settings for an end-item selected on the Navigator by selecting *Off*, *Detect* or *Correct* from the drop-down on the **Outliers** icon. You can also revert back to the global

default settings by selecting *Default* from this drop-down. Forecast Pro only detects outliers for end-items (i.e., non-group level data), so the selections are not available if a group level is selected on the Navigator.

Outliers are associated with the following modifiers:

(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.

\OUTLIER=OFF. Do not detect or correct outliers for the specified item.

\OUTLIER=DETECT. Detect and report outliers for the specified end-item, but do not correct for them (i.e., base the forecasts on the uncorrected values).

\OUTLIER=CORRECT. Detect and correct outliers for the specified end-item (i.e., base the forecasts on the corrected values).

Start Dates

Click on the ***Start Dates*** icon in the Data Tools group on the **Forecasting** tab or select *Forecasting>Set Start Dates* on the Navigator context menu to open the Set Start Dates dialog.
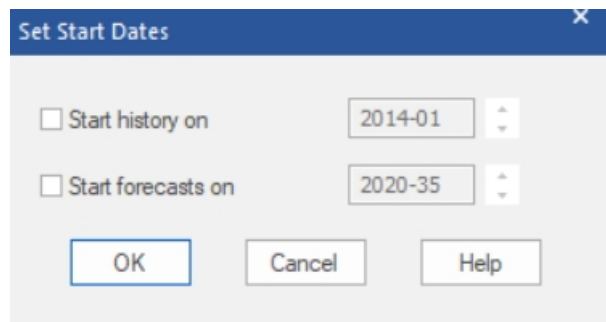
To only use history after a specified date to generate a forecast, select *Start history on* and use the spinner to set a start date.

To start forecasting an item starting with a specified date, check *Start forecasts on* and use the spinner to set the launch date.

Start Dates are associated with the following modifiers:

\START(Year, Period). Only use historic data after (and including) the specified date to estimate the specified item's forecast model.

\LAUNCH(Year, Period). Start forecasting an item on the specified date.

Remove Modifiers

Click on the **Remove Modifiers** icon in the Data Tools group on the **Forecasting** tab or select *Remove Modifier(s)* on the Navigator context menu to remove all modifiers from the selected item.

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*.

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 all groups and end-items.

Forecast Pro then recomputes the group-level forecasts by aggregating the constituent forecasts (bottom up). The original group-level forecasts are replaced, but the widths of their confidence limits are retained and re-centered on the new forecasts.

There are two alternative reconciliation approaches-- top-down and custom allocation. Both approaches are only available for group-level items.

Top-down Reconciliation

Top-down reconciliation directs Forecast Pro to perform top-down reconciliation for all members of the indicated group. First, Forecast Pro prepares forecasts for all groups and end-items. Then, it proportionally adjusts the nested lower-level end-item and group forecasts to sum to the group

forecast. If the top-down group is itself nested within larger groups, the forecasts for the larger groups are computed by the bottom-up approach.

You may select top-down reconciliation by selecting on the Navigator the group item at which you want to apply top-down and then selecting the **Top-down** icon in the Group Level group on the **Forecasting** tab. Alternatively, you may right click on the group level item on the Navigator and select **Top-down** from the context menu. In either case, you will see a **\TOPDOWN** modifier appear on the group level item which you selected.
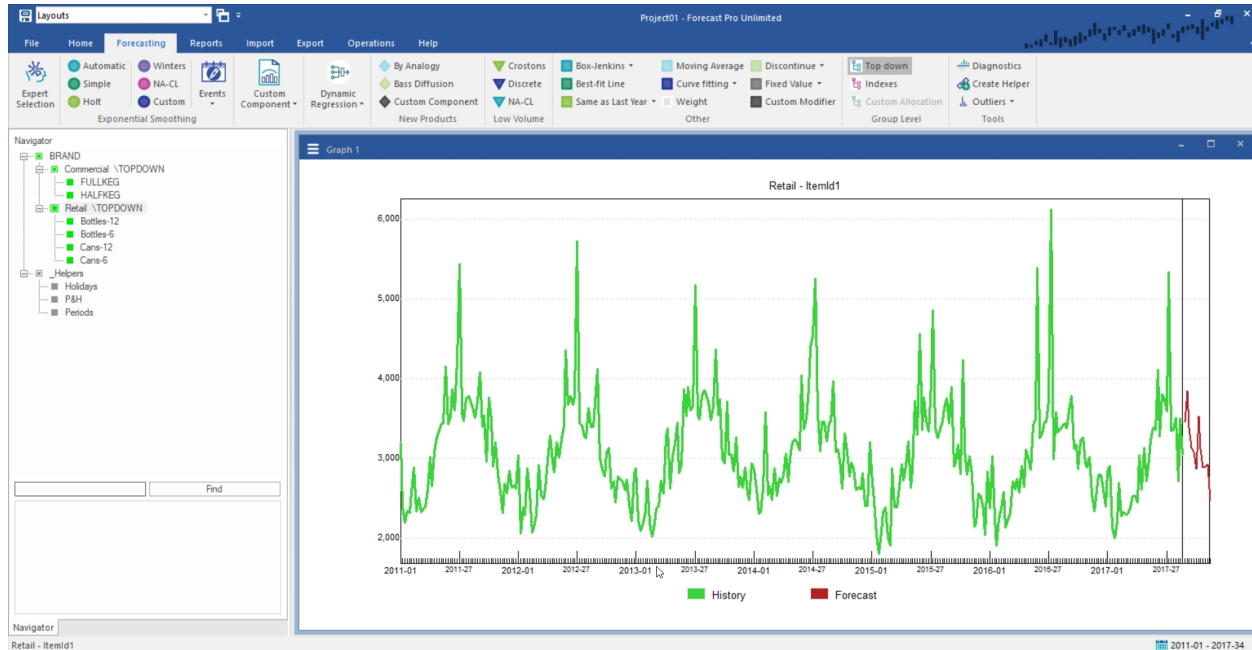
In the example above, \TOPDOWN directs Forecast Pro to compute forecasts and confidence limits for the groups Commercial and Retail at the group level. The end-item forecasts and confidence limits are multiplicatively adjusted so that the end-item forecasts sum correctly to the group level forecasts. 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.

Using Top-down reconciliation for Brand rather than Retail and Commercial would direct 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.

Custom Allocation

Custom allocation uses 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 Definitions: Custom Allocations in *Setting Up Your*

Optional Data for details. Note that Custom Allocation can only be applied to the group level directly above the level at which you defined the proportions.

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 are using custom allocation.

To apply custom allocation reconciliation, on the Navigator select a group item one level above the level at which the proportions are defined and then select the **Custom Allocation** icon in the Group Level group on the **Forecasting** tab. Alternatively, you mat right click on the group level item on the Navigator and select **Custom Allocation** from the context menu. In either case, you will see a **\CALL** modifier appear on the group level item which you selected.

Indexes

The third icon that appears in the Group Level group on the **Forecasting** tab is Indexes. Clicking the **Indexes** icon applies an **\INDEXES** modifier to the selected group level item on the Navigator. Using Indexes directs Forecast Pro to calculate multiplicative seasonal indexes at the selected group level and then use those indexes to de-seasonalize the histories for the nested end-items and groups. The de-seasonalized series are then forecasted and re-seasonalized.

The Indexes approach works well when the end-items share the same seasonality. It allows one to deal with items whose histories are too short to calculate seasonality.

Indexes can be used if the selected group level series does not include zeros or negative numbers. You cannot nest \INDEXES groups within each other. Model parameters are re-estimated for each end-item and nested group of the \INDEXES group.

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 de-seasonalizes each end-item's history using the indexes from the parent group, forecasts the resulting nonseasonal data, and then re-seasonalizes 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 [Basics Tab](#) on the Options dialog box.

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.

```
                              Total
                              S = 5
                              F = 5

           GroupA                              GroupB
           S = 2                               S = 3
           F = 2                               F = 3

    ItemA1          ItemA2            ItemB1            ItemB2
    S = 1           S = 1             S = 1             S = 2
    F = 1           F = 1             F = 1             F = 2
```

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.

```
                              Total
                              S = 5
                              F = 79
              ┌─────────────────┴─────────────────┐
           GroupA                              GroupB
           S = 2                               S = 3
           F = 76                              F = 3
       ┌─────┴─────┐                     ┌──────┴──────┐
    ItemA1       ItemA2               ItemB1         ItemB2
    S = 1        S = 1                S = 1          S = 2
    O = 75       F = 1                F = 1          F = 2
    F = 75
```

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.

```
                              Total
                              S = 5
                              F = 77
              ┌─────────────────┴─────────────────┐
           GroupA                              GroupB
           S = 2                               S = 3
           F = 2                               O = 75
                                               F = 75
       ┌─────┴─────┐                     ┌──────┴──────┐
    ItemA1       ItemA2               ItemB1         ItemB2
    S = 1        S = 1                S = 1          S = 2
    F = 1        F = 1                F = 25         F = 50
```

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.)

```
                              Total
                              S = 5
                              F = 79
                 ┌──────────────┴──────────────┐
              GroupA                          GroupB
              S = 2                           S = 3
              F = 76                          F = 3
         ┌──────┴──────┐               ┌──────┴──────┐
      ItemA1         ItemA2          ItemB1        ItemB2
      S = 1          S = 1           S = 1         S = 2
      O = 75         F = 1           F = 1         F = 2
      F = 75
```

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.

```
                                  Total
                                  S=5
                                  O = 475
                                  F = 475
                                  Locked=75
                                Unlocked=400
              ┌───────────────────────┴───────────────────────┐
           GroupA                                           GroupB
           S=2                                              S=3
           Unlocked S=1                                     Unlocked S=3
   Unlocked Allocation from Total=100         Unlocked Allocation from Total=300
           Locked=75                                        Locked=0
           F = 175                                          F = 300
           Unlocked=100                                     Unlocked=300
      ┌──────────┴──────────┐                     ┌──────────┴──────────┐
   ItemA1               ItemA2                  ItemB1               ItemB2
   S=1                  S=1                     S=1                  S=2
   O = 75               Unlocked S=1            Unlocked S=1         Unlocked S=2
   F = 75               F = 100                 F = 100              F = 200
```

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.

```
                                    Total
                                    S=5
                                  O = 475
                                  F = 475
                                 Locked=150
                                Unlocked=325

          GroupA                                        GroupB
          S=2                                           S=3
       Unlocked S=0                                  Unlocked S=3
Unlocked Allocation from Total=0          Unlocked Allocation from Total=325
       Locked=150                                    Locked=0
        F = 150                                       F = 325
       Unlocked=0                                    Unlocked=325

   ItemA1        ItemA2                       ItemB1            ItemB2
   S=1           S=1                          S=1               S=2
  O = 75        O = 75                   Unlocked S=1       Unlocked S=2
  F = 75        F = 75                     F = 108.33         F = 216.67
```

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, you will get an error message and the override is 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.

```
                            Total
                            S = 5
                           F = 302

          GroupA                              GroupB
          S = 2                               S = 3
         F = 150                             F = 152

   ItemA1        ItemA2              ItemB1            ItemB2
   S = 1         S = 1              S = 1             S = 2
  O = 75        O = 75             O = 150           F = 2
  F = 75        F = 75             F = 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.
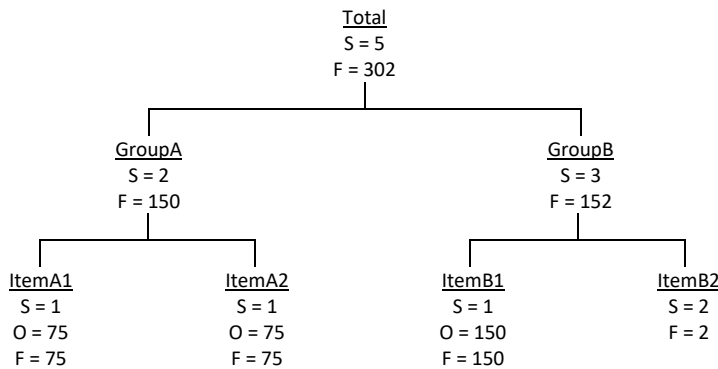
Figure 8.

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, these examples had 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 possible.

Appendices

Command Line Operation

You can run Forecast Pro Unlimited 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. Please note that this feature is not available for Forecast Pro 100.

To run Forecast Pro 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.) Run the command "ForecastProUnlimited [Project Name] /b [options…]". For example,

"C:\Program Files\Forecast Pro\Forecast Pro Unlimited v11 64-bit\ForecastProUnlimited.exe" Test.FPZip /b

runs the program located in "C:\Program Files\Forecast Pro\ Forecast Pro Unlimited v11 64-bit" for a project named Test.FPZip. The project should be in your current project directory. The "/b" parameter is used to indicate that you wish to run in command-line (batch) mode.

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.

Command line options include the following:

/BOOKMARK=[Bookmark name] causes the program to load the indicated bookmark prior to generating any output files. If the bookmark is not found, then the program continues using its current settings. Until build 5.1.1.4, this required a case-specific match; now it is case-independent. If the program cannot find the bookmark, processing continues normally to completion.

/LAYOUT=[Layout name] has the same functionality as **/BOOKMARK**.

/NOSAVE causes the program to forego re-saving the project at the end of command-line execution, regardless of the actual project setting (Settings > Options > Command Line > Project). The latter checkbox setting is respected if /NOSAVE is omitted.

/NOFORECAST the program simply loads the project, assuming it's in a forecast and reconciled state, and generates the output.

/TITLE=[file title to use for output files] causes the program to replace the project name, e.g. "Project01", with whatever is specified by the User here to build the output file names. For example /TITLE=Bookmark_1 would result in output files named accordingly, e.g., !Bookmark_1 Exceptions Report.xlsx. This setting is also applied to the ODBC output table, overriding any previous project setting. In the latter case, the User is advised to use a /TITLE value that complies with SQL naming conventions (e.g., no spaces).

/PLOG and **/ELOG** will force the program to generate project and error log files, respectively. This overrides the corresponding settings in the INI file and in the project definition. If a file title is specified using /TITLE=[…], it will be applied to the log file(s).

The following options allow the User to force the program to output the corresponding outputs/reports, even if they are unchecked within the project settings.

/OUTPUT – forces export of the Numeric Output regardless of the project setting.

/ITEM – forces export of the Item Report regardless of the project setting.

/OUTLIER – forces export of the Outlier Report regardless of the project setting.

/QUERY – forces export of any previously checked Query Report(s) regardless of the project setting.

/PWD=[project password] allows the command line execution of projects that contain a password.

/MODE=[XLSX or CSV or XLS] allows specifying a data mode other than XLSX for the Item, Exceptions, Outliers, and Overrides reports.

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.

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, as the network is often the performance bottleneck.

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 to reduce the amount of data that needs to be passed through the network. 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. 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 for all applications.

Use a Solid State hard drive, ideally with PCIe storage (virtualizing RAM with this type of hard drive is significantly faster). This will improve the read/write times which should make program performance for certain operations more efficient (e.g. opening and saving projects, as well as reading in/saving data files).

5. Use Manual Override Mode

If you make overrides to large complex hierarchies, the reconciliation can take a time. To defer this time consuming reconciliation, 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 **Recalculate** icon on the **Home** tab to perform the reconciliation. The control to turn on manual override mode is found on the Performance tab of the **Options** dialog box.

6. Use Manual Forecast Mode

By default, Forecast Pro will automatically recalculate the forecasts and re-reconcile the relevant nodes of 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 **Recalculate** icon on the **Home** tab. This allows you to specify multiple forecast modifiers without having to wait for the program to recalculate the forecasts in between each specification. The control to turn on manual forecast mode is found on the Performance tab of the **Options** dialog box.

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 machine learning" on the Performance tab of the **Options** dialog box will speed up forecast generation considerably. This option will remove machine learning 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 **Options** dialog box. If this switch is on (it is on by default), 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).

11. **Use Year over Year view of the history in the Override grid instead of Time Series view**

The Time Series view of the history in the Override grid allows you to enter formulas in Calculated Rows for historic periods. This may slow down override reconciliation, as the entire override grid must be recalculated each time an override is committed. Consider limiting the number and complexity of the formulas being used, or consider using Year over Year view of the history when committing overrides.

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., 465 Waverley Oaks Rd. Suite 318, Waltham, MA 02472 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.