SAS® Forecast Studio 1.2

User's Guide

The Power to Know®

The correct bibliographic citation for this manual is as follows: SAS Institute Inc. 2006. *SAS® Forecast Studio 1.2: User's Guide*. Cary, NC: SAS Institute Inc.

SAS® Forecast Studio 1.2: User's Guide

SAS Institute Inc., SAS Campus Drive, Cary, North Carolina 27513.

1st printing, January 2006

SAS Publishing provides a complete selection of books and electronic products to help customers use SAS software to its fullest potential. For more information about our e-books, e-learning products, CDs, and hard-copy books, visit the SAS Publishing Web site at **support.sas.com/pubs** or call 1-800-727-3228.

Contents

Chapter 1
Introduction to SAS Forecast Studio

Chapter Contents

Chapter 1
Introduction to SAS Forecast Studio

Overview of SAS Forecast Studio

SAS Forecast Studio is the client piece of the SAS Forecast Server. SAS Forecast Server has three components:

- the Analytics Platform
- SAS Forecast Server Mid-Tier
- SAS Forecast Studio

For information about each of these components and how to install them, see the *SAS Forecast Studio Administrator's Guide*.

SAS Forecast Studio is a forecasting application designed to speed the forecasting process through automation. The software provides for the automatic selection of time series models for use in forecasting timestamped data.

SAS Forecast Studio provides a Java-based graphical interface to the forecasting and time series analysis procedures contained in SAS High-Performance Forecasting and SAS/ETS software. For more information about these procedures and about the models underlying these procedures, see the *SAS High-Performance Forecasting User's Guide* and the *SAS/ETS User's Guide*.

Using this application, you can do the following tasks:

- Generate forecasts automatically using candidate model selection lists that are generated by the system. You can also generate forecasts using a candidate model selection list that you selected.
- Create your own forecasting models.
- Perform top-down, bottom- up, and middle-out hierarchical forecasting.
- Visually analyze and diagnose time series data.
- Override forecasts as needed.
- Export projects as SAS code for processing in a batch environment.

For more information about the version of SAS Forecast Studio that you are running, select **Help** → **About SAS Forecast Studio**.

Using Help

About Help

To open Help, select **Help** → **SAS Forecast Studio Help** or click  in the toolbar. You can also view Help for a specific window by clicking **Help**. Help opens in a window with three panes:

- On the left side of the window is the Navigation pane. It contains three navigational tabs: the table of contents tab, the search tab, and the index tab.
- On the right side of the window is the Topic pane. It displays the selected Help topic or the default Help topic.
- The third pane is the toolbar, which is located below the Help window title bar. Using the icons in this toolbar, you can specify the options for printing the current Help topic.

Finding a Help Topic

In the Navigation pane, click one of the following tabs:

- To browse through the table of contents, select the **Table of Contents**  tab. The Table of Contents is an expandable list of important topics. Double-click the topic name and select the subtopic that you want to view.
- To see a list of index entries, select the **Index**  tab. Then type a word in the **Find** box and press ENTER or scroll through the list. Topics are often indexed under more than one entry. Double-click the entry that you want to view.
- To locate every occurrence of a word or phrase that may be contained in a Help file, follow these steps:

  1. Select the **Search**  tab.
  2. In the **Find** text box, type the word that you want to find and press ENTER.
  3. From the results, select the topic that you want to view.

Printing Help Topics

To print the current topic in the **Topic** pane, click  in the toolbar.

To specify the page setup before printing a document, click  .

Chapter 2

Preparing an Input Data Set for SAS Forecast Studio

Chapter Contents

Chapter 2

Preparing an Input Data Set for SAS Forecast Studio

Overview of Time Series Data

In order for SAS Forecast Studio to generate a forecast, the time series data must meet the following requirements.

- The data set contains one variable for each time series.
- The data set contains exactly one observation for each time period.
- The data set contains a time ID variable that identifies the time period for each observation.
- The data is sorted by the time ID variable so that the observations are in order according to time.
- The data is equally spaced. This means that successive observations are a fixed time interval apart, and the data can be described by a single interval, such as hourly, daily, monthly.

SAS Forecast Studio creates time series data sets from your input data. For information about the requirements for the input data set, see "Required Format for the Input Data Set" on page 8.

SAS Forecast Studio creates the time series data through the following process:

1. The data is aggregated if the input is one of the following:
   - timestamped data that is recorded at no particular frequency (also called transactional data)
   - data recorded at a smaller time interval than needed for forecasting
2. Any gaps in the data are filled in. Gaps appear when there is not an observation for each time period or when the data is not equally spaced. The added observations have the required values of the time ID variable and the value that you specified for missing values. For more information, see "Missing Values in the Input Data Set" on page 12.
3. The data is sored by the BY variables and the time ID variable.

When you create a project, you select the input data set to use and assign variables to the time ID variable, BY variables, and dependent variables roles. SAS Forecast Studio uses this information to create the time series data.

Required Format for the Input Data Set

SAS Forecast Studio uses timestamped data and requires a date or datetime variable in the data set to generate forecasts. SAS Forecast Studio generates forecasts from time series data that consists of unique and equally spaced data over time. If the data is not equally spaced with regard to time, SAS Forecast Studio uses the date or datetime variable to accumulate the data into a time series before forecasting. The input data set must be a single SAS data set. You can have the following variables in the input data set.

- Dependent variables are the variables to model and forecast.
- Independent variables are the explanatory or input variables that are used to model and forecast the dependent variable.
- The time ID variable contains the date or datetime value of each observation.
- BY variables enable you to group observations into a hierarchy.
- Indicator variables are variables for holiday promotions.
- Reporting variables are not used for analysis but for reporting only.

Note: The names of the variables cannot match any of the reserved variable names that are used in the output data set. For more information, see "Reserved Variable Names" on page 9.

If you want to forecast the data hierarchically, then the lowest level of aggregation must be provided.

For more information about variable roles, see "Assign Variables to Roles" on page 34.

Here are two examples of the input data set for SAS Forecast Studio.

- This input data set contains monthly sales revenue and price information for the past 12 months. The variable Holiday indicates if there is any holiday during the month.

Date	Revenue	Avg. Price	Holiday
JAN2003	18817	26.3	0
FEB2003	52573	25.3	0
...	...	...	...
DEC2003	44205	20.3	1

- This input data set contains monthly retail sales information for different regions and product categories over the past 12 months. You can use the Region and Product variables to create a hierarchy for the sales forecasts.

Date	Sales	Avg. Price	Holiday	Region	Product
JAN2003	355	25.3	0	Region 1	Product 1
FEB2003	398	25.3	0	Region 1	Product 1
...	...	...	...	...	...
JAN2003	555	19.8	0	Region 1	Product 2
FEB2003	390	25.3	0	Region 1	Product 2
...	...	...	...	...	...
JAN2003	301	27.1	0	Region 2	Product 1
FEB2003	350	25.3	0	Region 2	Product 1
...	...	...	...	...	...
JAN2003	314	27.2	0	Region 2	Product 2
FEB2003	388	25.3	0	Region 2	Product 2
...	...	...	...	...	...
DEC2003	518	20.3	1	Region 2	Product 2

Note: In SAS Forecast Studio, the projects that contain hierarchies are limited to one dependent variable. If you want to forecast additional variables from the same hierarchy, then you need to create a separate project for each of the variables.

Reserved Variable Names

The variable names in your input data cannot start with an underscore and cannot match any of the variable names in the output data sets that SAS Forecast Studio creates. The following table lists alphabetically the variables that are created by the output data sets. For more information about the output data sets that are created, see the *SAS High-Performance Forecasting User's Guide*.

If your input data set contains one of these variables and you try to assign this variable to a role in the New Project Wizard, then an error message appears.

Table 2.1. Reserved Variable Names

Variable Name	Description
_*VariableName*	Any variable name that begins with an underscore
ACTUAL	Dependent series value
COMP	Name of the component
COMPONENT	Model component
CROSS	Cross variable name
DSVAR	Data set variable mapping
EST	Parameter estimate
FACTOR	Model factor
LABEL	Parameter or statistic label
LAG	Lag of input
LOWER	Lower confidence limit
MODE	Mode of decomposition
MODEL	Name of model
MODELVAR	Model variable mapping
NAME	Variable name

Table 2.1. (continued)

Variable Name	Description
PARM	Parameter name
PREDICT	Component forecast
PVALUE	Parameter estimate *p*-value
SEASON	Seasonal index
SELECT	Name of selection list
SHIFT	Shift
STAT	Statistic name
STATUS	Indicates success/failure in estimating parameter
STD	Prediction standard error
STDERR	Parameter estimate standard error
TIME	Time ID
TIMEID	Time ID values
TVALUE	Parameter estimate *t*-value
TRANSFORM	Transformation applied
UPPER	Upper confidence limit
AADJRSE	Amemiya's adjusted R-Square
ACF	Autocorrelations
ACF2STD	Indicates ACF beyond two standard errors
ACFLPROB	Autocorrelation log probabilities
ACFNORM	Normalized autocorrelations
ACFPROB	Autocorrelation probabilities
ACFSTD	Autocorrelation standard errors
ACOV	Autocovariances
ADJRSQ	Adjusted R-Square
AIC	Akaike information criterion
APC	Amemiya's prediction criterion
AVG	Average value
CC	Cycle component
CCF	Cross-correlations
CCF2STD	Indicates cross-correlations beyond two standard errors
CCFNORM	Normalized cross-correlations
CCFLPROB	Cross-correlation log probabilities
CCFPROB	Cross-correlation probablities
CCFSTD	Cross-correlation standard errors
CCOV	Cross-covariances
CSS	Corrected sum of squares
ERROR	Prediction errors
IACF	Inverse autocorrelations
IACF2STD	Indicates inverse autocorrelations beyond two standard errors
IACFNORM	Normalized inverse autocorrelations
IACFLPROB	Inverse autocorrelation log probabilities
IACFPROB	Inverse autocorrelation probabilities
IACFSTD	Inverse autocorrelation standard errors
IC	Irregular component
LAG	Time lag

Table 2.1. (continued)

Variable Name	Description
LAG*h*	Correlation or cross-correlation statistics for lag *h*
LOWER	Lower confidence limits
MAE	Mean absolute error
MAPE	Mean absolute percent error
MAXERR	Maximum error
MAXIMUM	Maximum value
MAXPE	Maximum percent error
ME	Mean error
MEAN	Mean value
MEDIAN	Median value
MINERR	Minimum error
MINIMUM	Minimum value
MINPE	Minimum percent error
MPE	Mean percent error
MSE	Mean square error
N	Number of non-missing observations or Number of variance products
NAME	Variable name
NMISS	Number of missing observations
NOBS	Number of observations
ORIGINAL	Original series index
PACF	Partial autocorrelations
PACF2STD	Indicates PACF beyond two standard errors
PACFLPROB	Partial autocorrelation log probabilities
PACFNORM	Partial normalized autocorrelations
PACFPROB	Partial autocorrelation probabilities
PACFSTD	Partial autocorrelations standard errors
PCSA	Percent change seasonal adjusted component
PERIOD*t*	Decomposition component value or trend statistic for time period *t*
PREDICT	Predicted values
RANGE	Maxinum value
RMSE	Root mean square error
RSQUARE	R-Square
RWRSQ	Random walk R-Square
SA	Seasonal adjusted component
SBC	Schwarz Bayesian information criterion
SC	Seasonal component
SCSTD	Seasonal component standard errors
SIC	Seasonal-irregular component
SEASON*s*	Season statistic value for season *s*
SSE	Sum of squares error
STD	Prediction standard errors
STDDEV	Standard deviation
SUM	Summation value
TC	Trend component
TCC	Trend-cycle component

Table 2.1. (continued)

Variable Name	Description
TCS	Trend-cycle-seasonal component
UMSE	Unbiased mean square error
URMSE	Unbiased root mean square error
UPPER	Upper confidence limits
USS	Uncorrected sum of squares
WN	White noise test statistics
WNLPROB	White noise test log probabilities
WNPROB	White noise test probabilities

Missing Values in the Input Data Set

For missing values in other variables, including the dependent, independent, and external forecast variables, you can specify how to interpret missing values (regardless of the variable roles) when you create a project. For more information, see "Perform Additional Steps" on page 41.

Additional Information

Often your data is not in the appropriate format for SAS Forecast Studio. To avoid misleading or incorrect analyses from your time series data, you should conduct data preprocessing.

- For general information about working with time series data, see the *SAS/ETS User's Guide*.
- For more information about creating time series data from transactional data, see "The TIMESERIES Procedure" and "The EXPAND Procedure" documentation in the *SAS/ETS User's Guide*.
- For more information about creating SAS data sets from Excel files, see the IMPORT Procedure documentation in the *Base SAS Procedures Guide*.
- For more information about transposing data for statistical analysis, see "The TRANSPOSE Procedure" documentation in the *Base SAS Procedures Guide*.

Chapter 3
Using the Workspace

Chapter Contents

Chapter 3
Using the Workspace

Overview of the Workspace

The SAS Forecast Studio workspace consists of the following components:

Toolbar
: enables you to create a new project, open existing projects, or open an archived project. You can also create forecasts from the toolbar.

Hierarchy tab
: displays the project data as a tree when you choose to create a forecast hierarchy. You can expand and collapse the nodes of the tree.

Table tab
: displays the project data in a sortable table. You can select the fit criterion that you want to use to sort the data.

Forecasting View
: displays the reconciled forecast for the model that you selected.

Series Analysis View
: displays the time series plot for the selected series on the **Hierarchy** tab or **Table** tab.

Model Analysis View
: displays the forecast plots for the selected model. Using this view, you can view, edit, copy, and create models. You can also change the model that is currently selected.

Series Properties pane
: displays the properties for the currently selected series.

Status bar
: displays information about the current procedure, the username, and the metadata server for the current project.

About the Toolbar

The toolbar enables you to manage your project, generate a forecast, and open Help. The following table explains each of the icons in this toolbar.

Table 3.1. Buttons in the Toolbar

Button	Description
	enables you to create a new project
	enables you to manage existing projects
	enables you to archive a project
	enables you to reforecast the data for all series
	enables you to reforecast the data for the selected series
	opens the model selection list
	enables you to run a report
	opens SAS Forecast Studio Help

About the Hierarchy Tab

The **Hierarchy** tab displays a hierarchical view of the forecasts. This hierarchy is defined by the variables that you assign to the BY Variables role when you create the project. Once the hierarchy is created, it cannot be changed. To create another hierarchical representation of the data, you must create another project.

Note: You must create separate projects for different hierarchical representations of the data.

You can expand or collapse each node in the tree. To display the forecast for a node in the Forecasting View, select the node in the hierarchy.

Note: When you select a node on the **Hierarchy** tab, the corresponding model is selected on the **Table** tab.

If you have added a manual override for a series or added a note to a series, this is indicated with an icon. If you have used a filter to create an exception rule, the series that are flagged as exceptions appear in red. For more information, see "Managing Filters" on page 51.

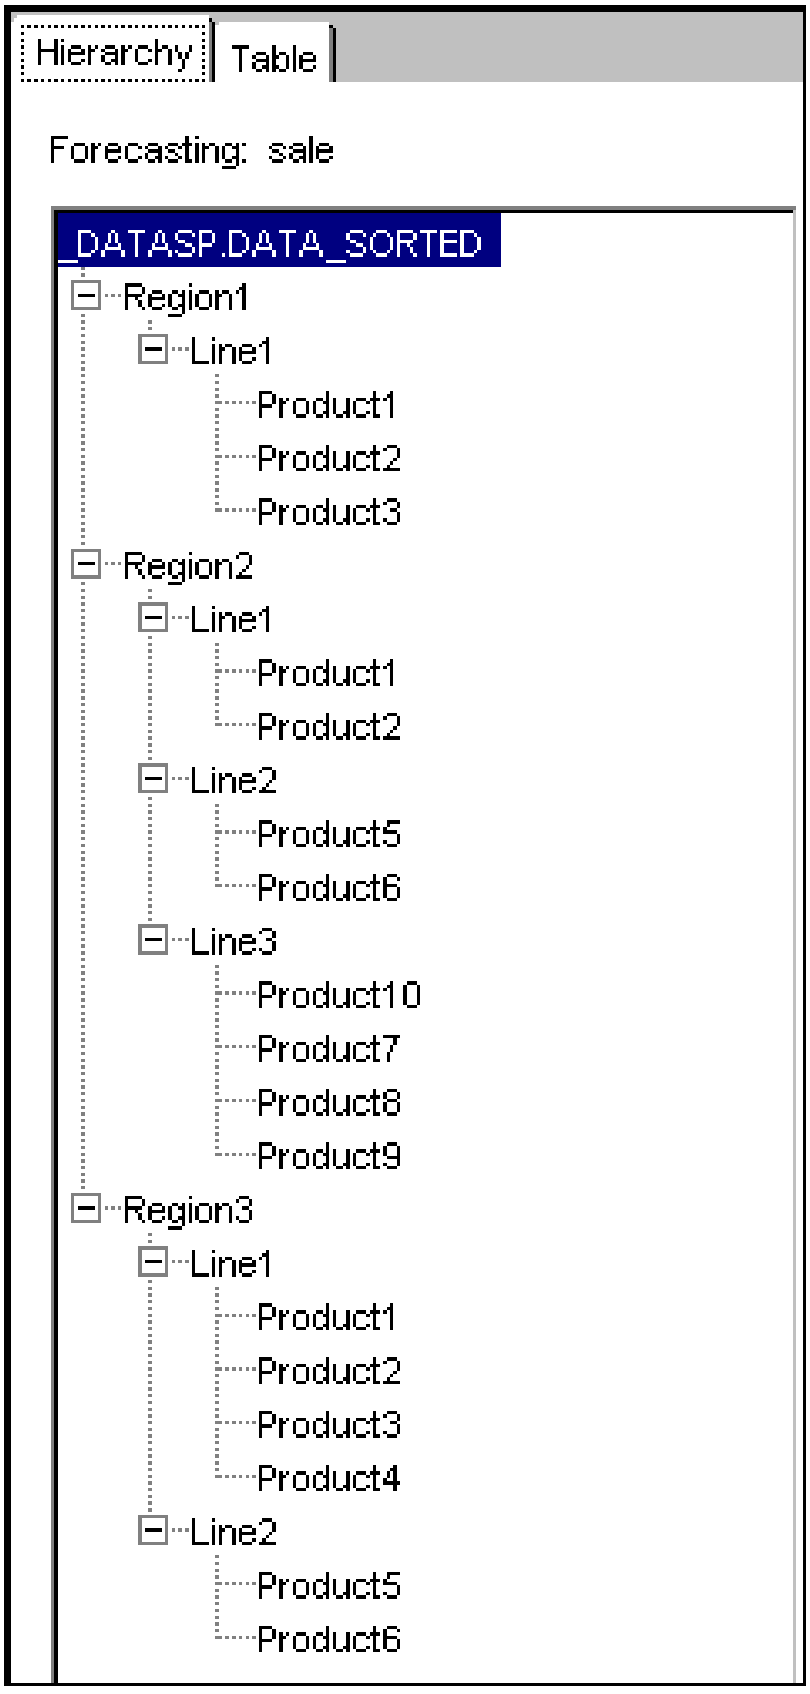

Figure 3.1. Hierarchy Tab

For more information about how to customize the hierarchy, see "Customizing the Hierarchy Tab" on page 29.

About the Table Tab

The **Table** tab enables you to use a fit criterion to find out which are the best and the worst fitting models. For series, you can see the value of the fit statistic for the current model.

You can filter the contents of the table by using the **Filter** drop-down list. Depending on your hierarchy and aggregation statistic, the following options are available:

- **No Filter** - lists all of the series. No filter is applied. This is the default when you generate a forecast.
- **Notes** - lists only the series that have notes.
- **Overrides** - lists only the series with manual overrides.
- **All Exceptions** - lists only the series that are identified as exceptions. If the filter for the exception rule is created from the New Project Wizard, then all exceptions are displayed by default.
- an option for each of the BY variables - lists the series for that variable.
- any user-defined filters.

To refresh the contents in the table, click **Refresh**.

Note: When you click a series on the **Table** tab, the corresponding node is selected on the **Hierarchy** tab.

An asterisk (*) in the table indicates aggregation. If there is an asterisk in a hierarchy level (BY variable) column, then all of the series at that level have been aggregated to the next level of the hierarchy.

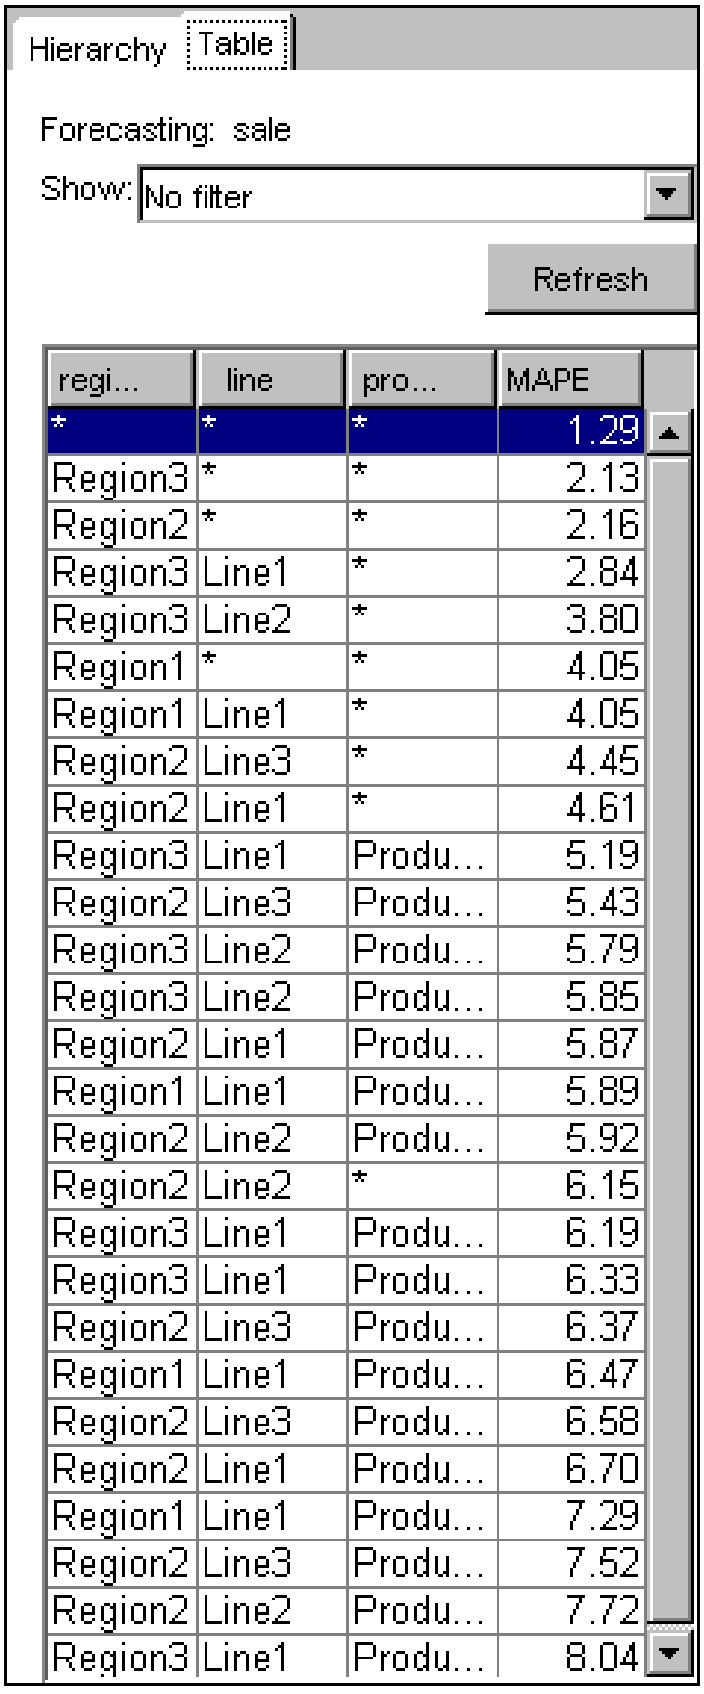

regi...	line	pro...	MAPE
*	*	*	1.29
Region3	*	*	2.13
Region2	*	*	2.16
Region3	Line1	*	2.84
Region3	Line2	*	3.80
Region1	*	*	4.05
Region1	Line1	*	4.05
Region2	Line3	*	4.45
Region2	Line1	*	4.61
Region3	Line1	Produ...	5.19
Region2	Line3	Produ...	5.43
Region3	Line2	Produ...	5.79
Region3	Line2	Produ...	5.85
Region2	Line1	Produ...	5.87
Region1	Line1	Produ...	5.89
Region2	Line2	Produ...	5.92
Region2	Line2	*	6.15
Region3	Line1	Produ...	6.19
Region3	Line1	Produ...	6.33
Region2	Line3	Produ...	6.37
Region1	Line1	Produ...	6.47
Region2	Line3	Produ...	6.58
Region2	Line1	Produ...	6.70
Region1	Line1	Produ...	7.29
Region2	Line3	Produ...	7.52
Region2	Line2	Produ...	7.72
Region3	Line1	Produ...	8.04

Figure 3.2. Table Tab

For more information about how to customize the table, see "Customizing the Table Tab" on page 29.

About the Forecasting View

Opening the Forecasting View

By default, the Forecasting View is open when you generate a forecast. To hide this view, select **View** → **Forecasting View**.

You can have multiple views open during a single session. To open the current view in a separate window, select **Window** → **Detach**. To close a window, click  .

Note: Detached windows are read-only. You cannot add manual overrides or notes from these windows.

Components of the Forecasting View

In the Forecasting View, you can view the forecast for the selected series in the **Hierarchy** tab or in the **Table** tab. The Forecasting View has three sections: a time plot, the Data Table, and a notes section. Above the time plot, SAS Forecast Studio displays the statistic of fit for the selected series and lists whether the series has broken any exception rules and the condition that led to the rule being broken.

About the Time Plot

When SAS Forecast Studio generates the forecasts for a project, only the time plot is displayed. By default, the time plot displays the historical values, the fitted values, and the reconciled forecast for the selected model.

About the Data Table

To view the Data Table, expand the **Data Table** section. The Data Table is shown by default.

By default, the Data Table also displays the following information:

- historical data
- statistical forecast
- reconciliation adjustment
- reconciled forecast
- manual overrides
- final forecast

Historical values are displayed in regular font. Future values are displayed in bold. Cells that you can edit have a white background. For more information, see "Add Overrides" on page 74.

About the Notes Section

In the Notes box, you can add notes for the statistical forecast or manual overrides. The Notes section is hidden by default. To view the notes for the selected series, expand the **Notes** section. For more information, see "Adding Notes" on page 76.

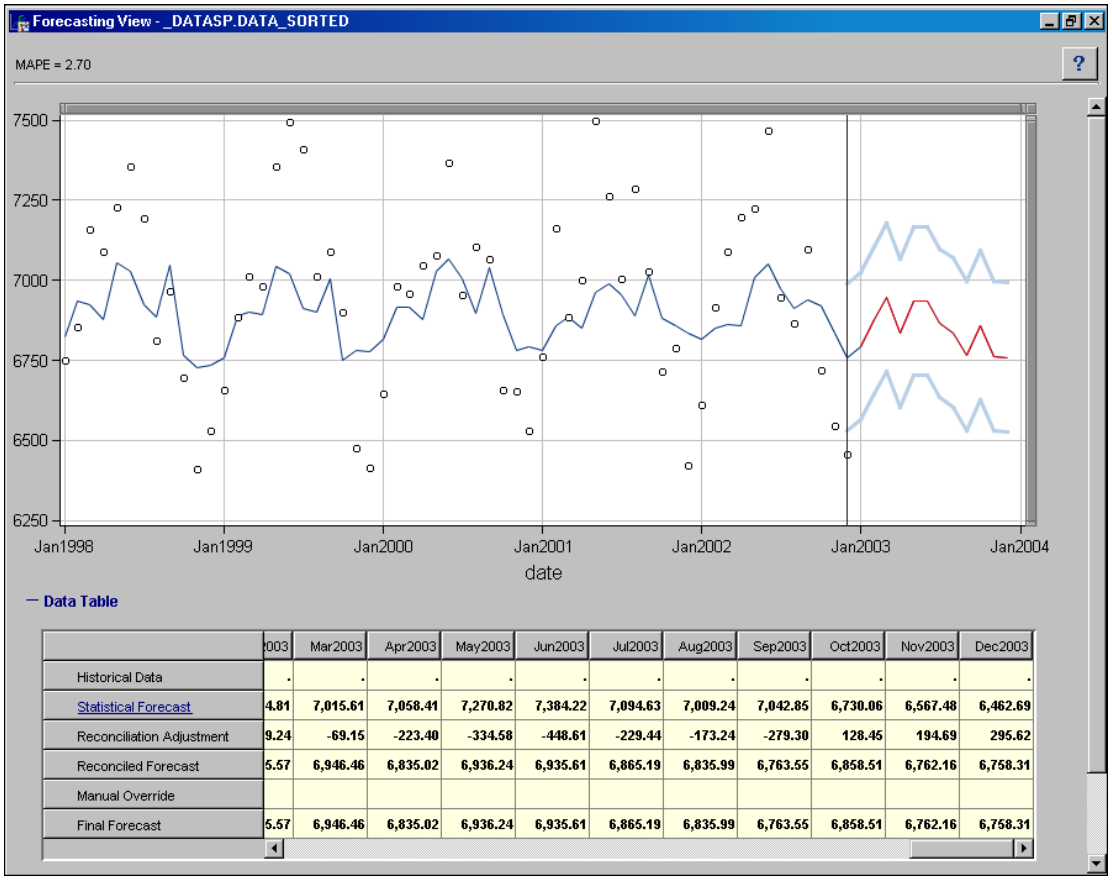

	003	Mar2003	Apr2003	May2003	Jun2003	Jul2003	Aug2003	Sep2003	Oct2003	Nov2003	Dec2003
Historical Data	.	.	.	.	.	.	.	.	.	.	.
Statistical Forecast	4.81	7,015.61	7,058.41	7,270.82	7,384.22	7,094.63	7,009.24	7,042.85	6,730.06	6,567.48	6,462.69
Reconciliation Adjustment	9.24	-69.15	-223.40	-334.58	-448.61	-229.44	-173.24	-279.30	128.45	194.69	295.62
Reconciled Forecast	5.57	6,946.46	6,835.02	6,936.24	6,935.61	6,865.19	6,835.99	6,763.55	6,858.51	6,762.16	6,758.31
Manual Override											
Final Forecast	5.57	6,946.46	6,835.02	6,936.24	6,935.61	6,865.19	6,835.99	6,763.55	6,858.51	6,762.16	6,758.31

Figure 3.3. Forecasting View

Working in the Forecasting View

In the Forecasting View, you can complete the following tasks:

- open the model selection list. For more information, see "Using the Model Selection List" on page 103.
- add overrides. For more information, see "Add Overrides" on page 74.
- add notes. For more information, see "Adding Notes" on page 76.

About the Series Analysis View

Opening the Series Analysis View

To open the Series Analysis View, select **View** → **Series Analysis View**. The Series Analysis View is not open by default.

You can have multiple windows open during a single session. To open the current view in a separate window, select **Window** → **Detach**. To close a window, click  .

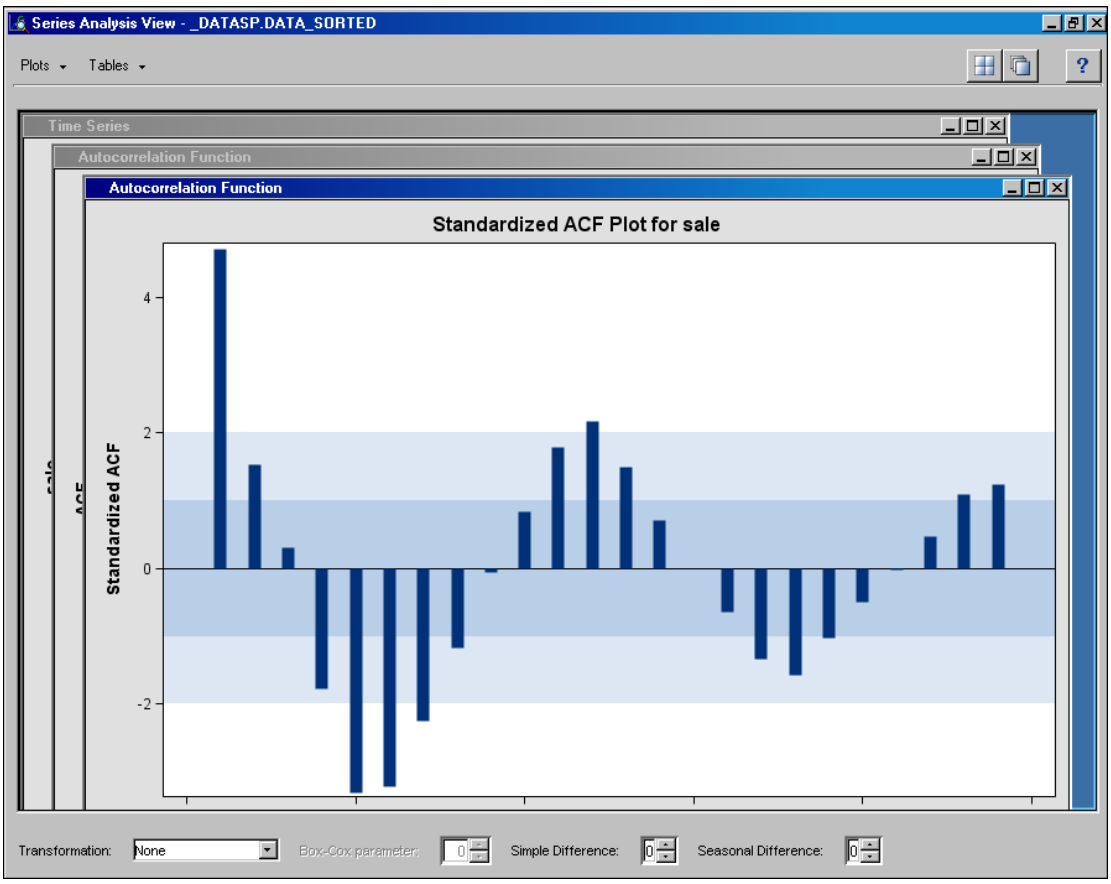

Figure 3.4. Series Analysis View

Components of the Series Analysis View

The Series Analysis View displays various plots and tables for the series that is selected on the **Hierarchy** tab or on the **Table** tab.

You can choose the time series plots and tables that you want to display. You can also transform a time series from this window.

If you have several plots open simultaneously, you might want to cascade or tile the plots in the window. To organize the plots in the Series Analysis View, select **Plots** and then **Cascade** or **Tile**.

Working in the Series Analysis View

From the Series Analysis View, you can complete the following tasks:

- view time series plots. For more information, see "View the Time Series Plots" on page 23.
- view time series tables. For more information, see "View the Time Series Tables" on page 24.
- transform a time series. For more information, see "Transform a Time Series" on page 25.
- edit the selected model. For more information, see "Edit a Model" on page 99.

View the Time Series Plots

From the Series Analysis View, you can select the plots to display. If you right-click a plot, you can select various options from the drop-down menu. For example, you can zoom in or out on a plot.

To open a time series plot, click **Plots** and select one of the following plots:

- **Autocorrelation Function** - displays plots of the autocorrelation function and standardized autocorrelation function.
- **Correlation Graphics Panel** - displays plots of the autocorrelation function, partial autocorrelation function, inverse autocorrelation function, and white noise probability tests in a single frame.
- **Cross-series Correlations** - displays correlations between the dependent variable and a second variable. From the Select Series window, you choose the second variable from a list of independent and reporting variables.
- **Cross-series Plot** - displays a two-variable plot of the dependent variable and a second variable. From the Select Series window, you choose the second variable from a list of independent and reporting variables.
- **Cycle Component** - displays a plot of the cycle component estimated by decomposing the time series.
- **Inverse Autocorrelation Function** - displays plots of the inverse autocorrelation function and standardized inverse autocorrelation function.
- **Irregular Component** - displays a plot of the irregular component estimated by decomposing the time series.
- **Partial Autocorrelation Function** - displays plots of the partial autocorrelation function and standardized partial autocorrelation function.
- **Percent-Change Adjusted Series** - displays a plot of seasonally adjusted time series that are expressed as percent change.

- **Seasonal-Irregular Component** - displays a plot of the seasonal and irregular components that are estimated by decomposing the time series.
- **Seasonal Component** - displays a plot of the seasonal component estimated by decomposing the time series.
- **Seasonal Cycles** - displays a plot of the seasonal cycles estimated by decomposing the time series.
- **Decomposition Graphics Panel** - displays trend-cycle, seasonal-irregular, irregular, and seasonally adjusted plots in a single frame.
- **Seasonally Adjusted Series** - displays a plot of the seasonally adjusted time series.
- **Series Histogram** - displays the time series in a frequency histogram.
- **Time Series** - plots the current time series.
- **Trend-Cycle-Seasonal Component** - displays a plot of the trend, cycle, and seasonal components estimated by decomposing the time series.
- **Trend-Cycle Component** - displays a plot of the trend and cycle components estimated by decomposing the time series.
- **Trend Component** - displays a plot of the trend component estimated by decomposing the time series.
- **White Noise Probability** - displays white noise probability plots and plots of the log of the white noise probability for the time series.

About the Select Series Window

From the Select Series window, you select the independent or reporting variable to plot against the current series. This window appears when you select the **Cross-series Plots** or **Cross-series Correlations** option in the Series Analysis View. For more information, "View the Time Series Plots" on page 23.

View the Time Series Tables

To open a time series table in the Series Analysis View, complete the following steps:

1. Click **Tables**.
2. From the pop-up menu, select from the following options:
   - **Seasonal Statistics** - displays the seasonal statistics for the time series.
   - **Seasonal Decomposition** - displays trend-cycle, seasonal, irregular, and seasonally adjusted components. Seasonal Statistics displays n, minimum, maximum, sum, and standard deviation of the series for each season.
   - **Descriptive Statistics** - displays descriptive statistics for the time series.
   - **Accumulated Descriptive Statistics** - displays descriptive statistics for the accumulated time series. This is equivalent to Descriptive Statistics if accumulation is not being applied.

The table opens in the Series Analysis View.

After the table opens, you can right-click the table to choose from the following options:

- **Sort** - sorts the table in ascending or descending order by the selected column. (You can also do this by clicking the column headings.)
- **Display As** - displays values in numeric columns as text, graphics (a horizontal bar within the table cell), or funnel (centered horizontal bar).
- **Hide Column** - hides the selected column
- **Show All Columns** - redisplays any hidden columns
- **Copy** - copies the table to the clipboard as an image or as values you can paste into a spreadsheet.
- **Print** - prints either the visible portion or the entire table.

Transform a Time Series

From the Time Series window, you can transform a time series and the results.

To transform a time series, complete the following steps:

1. Select the transform. The available options depend on the open plots.
2. Select the decomposition.
3. Specify the simple difference.
4. Specify the seasonal difference.

All open plots show the specified transformation.

About the Model Analysis View

Opening the Model Analysis View

To open the Model Analysis View, select **View** → **Model Analysis View**. The Model Analysis View is not open by default.

You can have multiple Model Analysis View windows open during a single session. To open the current view in a separate window, select **Window** → **Detach**. To close a window, click  .

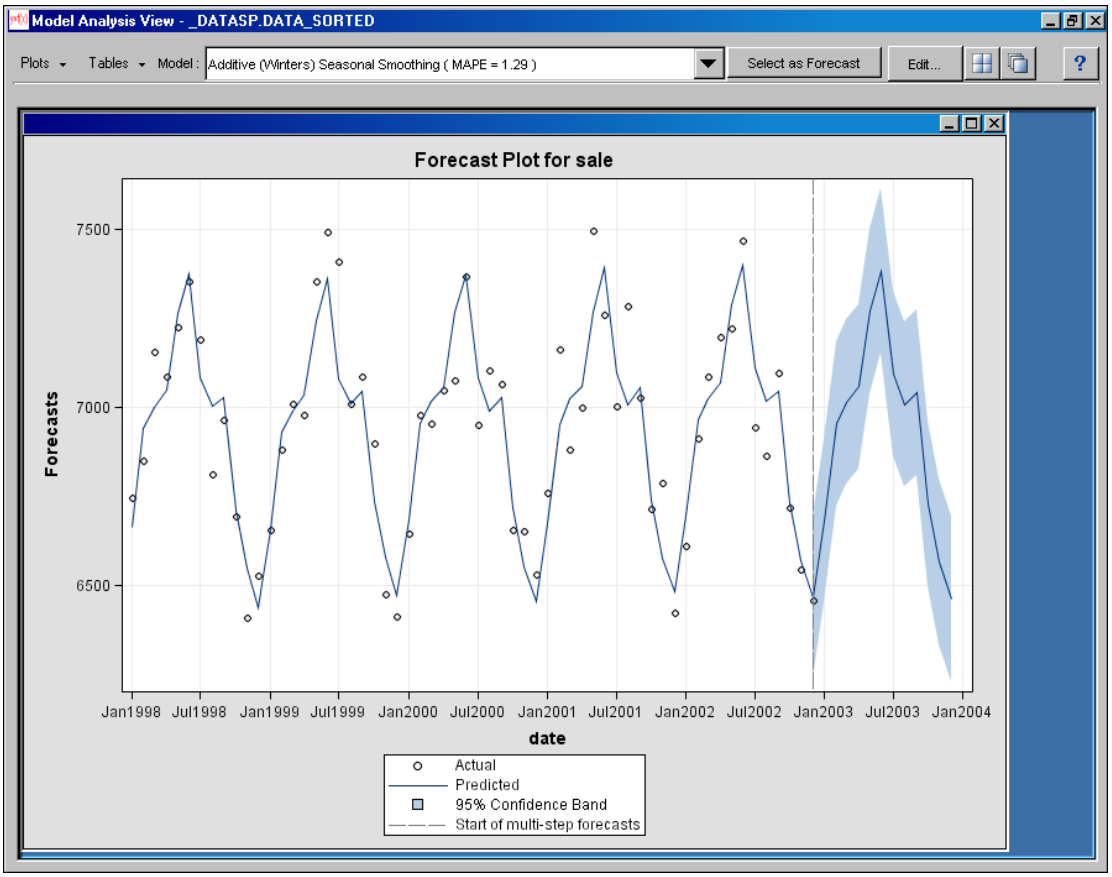

Figure 3.5. Model Analysis View

Components of the Model Analysis View

The Model Analysis View displays the model plots for the selected model. The **Model** drop-down list shows the name of the currently displayed model. This drop-down list contains the list of models that have been fitted to this series. You can choose which of the fitted models to display.

By default, the Time plot, Statistics of Fit plot, and the Model Specification appear in the window. You can open additional model plots or additional model tables. If you have several plots or tables open simultaneously, you might want to cascade or tile the plots or tables in the window. To organize the plots or tables in the Model Analysis View window, select **Cascade** or **Tile**. SAS Forecast Studio remembers your final configuration and uses that same configuration the next time that you open the Model Analysis View.

Working in the Model Analysis View

From the Model Analysis View, you can complete the following tasks:

- view model plots. For more information, see "View the Model Plots" on page 27.
- view model tables. For more information, see "View the Model Tables" on page 27.
- select a model to forecast. For more information, see "Select the Model to Forecast" on page 102.
- edit the selected model. For more information, see "Edit a Model" on page 99.
- create a new model. For more information, see "Create a New Model" on page 79.

View the Model Plots

To open a plot in the Model Analysis View, complete the following steps:

1. Click **Plots**.
2. From the pop-up menu, select from the following options:
   - **Autocorrelation Function** - displays the autocorrelation function and standardized autocorrelation function of the prediction errors.
   - **Model Components** - displays plots of each of the components that are used by the model.
   - **Forecast Only** - displays the forecast in the forecast horizon only.
   - **Forecast** - displays a plot of the model and the forecast that covers the time range of the actual data and the forecast horizon.
   - **Inverse Autocorrelation Function** - displays the inverse autocorrelation function and standardized inverse autocorrelation function of the prediction errors.
   - **Partial Autocorrelation Function** - displays the partial autocorrelation function and standardized partial autocorrelation function of the prediction errors.
   - **Model** - displays a plot of the model with predictions and confidence limits that cover the time range of the actual data.
   - **Prediction Errors** - displays a needle plot and histogram of the prediction errors.
   - **White Noise Probabilities** - displays the white noise probability plot and log white noise probability plot of the prediction errors.

View the Model Tables

To view a model table, complete the following steps:

1. Click **Tables**.

2. From the pop-up menu, select from the following options:

   - **Model Components** - displays tables for each of the components used by the model.
   - **Forecast Summary** - displays forecasted values of the dependent variable for the time range of the forecast horizon.
   - **Forecasts** - displays forecasted values, standard errors, and confidence limits for the time range of the forecast horizon.
   - **Model Parameter Estimates** - displays estimates, standard errors, and significance tests for each of the model parameters.
   - **Statistics of Fit** - displays all available statistics of fit for the forecasting model.
   - **Test of Unbiasedness** - displays all the tests for unbiasedness for the forecasting model.

The model table opens in the Model Analysis View.

After the table opens, you can right-click the table to access the following options:

- **Sort** - sorts the table in ascending or descending order by the selected column. (You can also do this by clicking the column headings.)
- **Display As** - displays the values in numeric columns as text, graphics (a horizontal bar within the table cell), or funnel (centered horizontal bar).
- **Hide Column** - hides the selected column
- **Show All Columns** - redisplays any hidden columns
- **Copy** - copies the table to the clipboard as an image or as values that you can paste into a spreadsheet.
- **Print** - prints either the visible portion or the entire table.

About the Status Bar

The status bar is located at the bottom of the main workspace. On the right side of the status bar, you see your user ID and the metadata server that you are currently using.

Figure 3.6. Right Side of the Status Bar

The left side of the status bar gives you the status of the current action in SAS Forecast Studio. For example, if you are generating forecasts, then the status bar displays the name of the procedure that is currently executing.

Customizing the Workspace

Customizing the Hierarchy Tab

To customize the Hierarchy tab, select **View** → **Customize Hierarchy**. The Customize Hierarchy dialog box opens.

To select columns to be displayed in the Hierarchy tab, select the components in the **Available items** list and click  to move them to the **Selected items** list.

You can choose the following components from the **Available items** list:

- **Notes** – They are identified in the **Hierarchy** tab by the  icon.
- **Overrides** – They are identified in the **Hierarchy** tab by the  icon.
- **Exceptions** – Series that are flagged as exceptions appear in red in the **Hierarchy** tab. The exceptions are displayed by default.

Customizing the Table Tab

To customize the Table tab, select **View** → **Customize Table**. The Customize Table dialog box opens.

You can choose the following components from the **Available columns** list:

- **Notes** – They are identified in the **Table** tab by the  icon.
- **Overrides** – They are identified in the **Table** tab by the  icon
- **Exceptions** – They are identified in the **Table** tab by . The exceptions are displayed by default.
- fit statistics other than the model selection criterion at the project level. The model selection criterion at the project level is always displayed.
- the list of series properties, such as the number of missing values, the number of zero values, the number of negative values, and the number of nonmissing values.

Chapter 4
Working with Projects

Chapter Contents

Chapter 4
Working with Projects

Create a New Project

You can have only one project open in SAS Forecast Studio at a given time. You can create a new project either when you first open SAS Forecast Studio or after you close an existing project. You create a new project by using the New Project Wizard. How you open the wizard depends on where you are in the application.

- To create a new project when you open SAS Forecast Studio, click **Create a new project** in the Welcome to SAS Forecast Studio dialog box.
- To create a new project when SAS Forecast Studio is already running, select **File** → **New Project** or click  in the toolbar.

In the New Project Wizard, you must complete the following steps:

1. Select a data set.
2. Specify how to forecast your data.
3. Assign variable to roles.
4. Configure the hierarchy (optional).
5. Enter project properties.
6. Perform additional steps.

Select a Data Set

In the New Project Wizard, you can select the data set that includes the data that you want to forecast.

To select a SAS data set, complete the following steps:

1. In the box, double-click the name of the server or directory to view the list of available libraries. Double- click the library name to view the data sets within that library. Select the data set that you want to use. The full path for the data set appears in the field below the box.

   The Information box displays the metadata for the current selection.

   Note: To return to the previous level in the directory path, click  .
2. Click **Next**.

Note: Before you can use your data in SAS Forecast Studio, it must be in the required format. For more information, see "Required Format for the Input Data Set" on page 8.

Specify How to Forecast Your Data

In the New Project Wizard, you can indicate whether you want to forecast your data in a hierarchy.

To specify how to forecast your data, complete the following steps:

1. Select from the following options:

   - **I want to forecast my data as a hierarchy** - specifies that you want SAS Forecast Studio to forecast the data hierarchically. This option is selected by default.

     When you select this option, the variable that you assign to the Dependent Variable role in the the Assign Variables to Roles step is sorted into individual series in a hierarchy using the BY variables that you assign. If you choose to create a forecast hierarchy, then only one dependent variable can be forecast.
   - **I do not want to forecast my data as a hierarchy** - enables you to forecast the data without creating a hierarchy. In the Assign Variables to Roles step, you can assign multiple variables to the Dependent Variable role.

2. Click **Next**.

Assign Variables to Roles

You can assign variables in your data set to roles in the following ways:

- When you are creating a project, use the New Project Wizard.
- After the project has been created, you can assign additional variables or change variable roles by using the Variables dialog box. To open this dialog box, select **Project** → **Variable Assignment**. The Variables dialog box opens.

For more information, see "How to Assign a Variable to a Role" on page 36.

To assign a variable to a role, complete the following steps:

1. Assign the variables to the following roles:

   - **Time ID** - specifies the time ID for the data. The time ID variable is a variable in the input data set that contains the SAS date, datetime, or time value for each observation. This variable is used to determine the frequency and ordering of the data and to extrapolate the time ID values for the forecasts. You can assign only one variable to this role, and it must be either a date variable, a datetime variable, or a numeric variable that contains date or datetime values.

     The values in the **Interval** drop-down list and the **Seasonal cycle length** box are determined by the variable that you assign to the Time ID role.

SAS Forecast Studio analyzes the variable assigned to the Time ID role to detect the time interval of the data. For example, if the Time ID variable contains data values that are spaced one month apart, then **MONTH** is the default selection in the **Interval** drop-down list, and the **Seasonal cycle length** box contains the corresponding default seasonal cycle length.

If the time interval cannot be detected from the variable that you assign, then you need to specify the interval and seasonal cycle length. Often the time interval cannot be detected with transactional data (time-stamped data that is recorded at no particular frequency). If this is the case, then SAS Forecast Studio accumulates the data into observations that correspond to the interval that you specify. For non-transactional data, you might need to specify the interval and seasonal cycle lenth if there are numerous gaps (missing values) in the data. In this case, SAS Forecast Studio supplies the missing values.

You can specify a greater time interval than that found in the input data. A smaller interval should not be used, since a smaller interval will generate a large number of observations.

You can specify a seasonal cycle length other than the default if you want to model a cycle in the data. For example, if your data contains a 13 week cycle, then type `13` in the **Seasonal cycle length** box.

- **BY Variables** - groups together observations that have the same value for the BY variable. Assigning a BY variable enables you to obtain separate analyses for groups of observations. You can assign character and numeric variables to this role.

  The order of the BY variables describes the structure of the hierarchy. If you choose to forecast the data hierarchically, then you can change the order of the hierarchy in the Configure the Hierarchy step of the New Project Wizard.

  If you have a hierarchy and get a message that the time interval of the data cannot be deleted, make sure that the BY variables are specified in the same order that the data set is sorted.

- **Dependent Variables** - specifies which variables to model and forecast. You must assign at least one numeric variable to this role. If you choose to forecast your data hierarchically, then only one dependent variable can be forecast.

- **Independent Variables** - specifies a explanatory, input, predictor, or causal factor variables. SAS Forecast Studio, by default, classifies the independent variable as stochastic. You can assign only numeric variables to this role.

2. Click **Adjustments** to add an adjustment variable. For more information, see

"Working with Adjustment Variables" on page 36.

3. Click **More Roles** to add a reporting variable. For more information, see "Add a Reporting Variable" on page 37.
4. If you are using the New Project Wizard, click **Next** to move to the next step. In the Variables dialog box, click **OK** to save your changes.

   Note: If SAS Forecast Studio cannot create the time series data, then an error message will appear. For information about these errors, see "Troubleshooting the Time ID Variable" on page 37.

   Note: If you assign a variable to a role and the variable name is reserved, then an error message will appear. For more information about variable names, see "Reserved Variable Names" on page 9.

How to Assign a Variable to a Role

To assign a variable to a role, select the variable in the **Variables to assign** box and click  for that role. To remove a variable from a role, click  .

Icons beside the variable names indicate the type of variable:

	numeric variable
	character variable
	date/time variable

Icons beside the role names indicate the type of variable that you can assign to the role:

	accepts numeric variable
	accepts character variable
	accepts date/time and numeric variables

Working with Adjustment Variables

You can add adjustment variables in the following ways:

- When you create a project, you can add an adjustment variable in the Assign Variables to Roles step in the New Project wizard.
- After you have created the project, you can manage your adjustment variables by selecting **Project** → **Variable Assignments**. The Variables dialog box opens. Click **Adjustments** to open the Adjustment Variables dialog box.

In the Adjustment Variables dialog box, you can add and remove adjustment variables from the project. You can also edit and specify the order of existing adjustment variables.

The Information box displays the description, the dependent variables being adjusted, and the operations to perform before and after forecasting for the adjustment variable that is selected in the **Adjustment Variables** list.

- To add the adjustment variable, click **New**. The New Adjustment Variable dialog box opens.
- To edit the adjustment variable, select the variable from the list and click **Edit**. The Edit *adjustment-variable-name* dialog box opens.
- To change the order of a variable in the list, select the variable from the **Variable** list and click **Move Up** or **Move Down**.
- To delete an adjustment variable, select the variable from the list and click **Delete**. You are prompted to confirm this deletion. Click **OK** to delete the variable.

Add or Edit an Adjustment Variable

To add or edit an adjustment variable, complete the following steps:

1. Optionally, specify a description for the adjustment variable.
2. Select the adjustment variable and move it to the **Selected** box. The name of this variable will appear in the **Adjustment Variables** list in the Adjustment Variables dialog box. The variable assigned as the adjustment variable cannot have any other role in the project.
3. Specify the operations to perform before and after forecasting. You can choose from add, subtract, multiply, divide, minimum, and maximum. By default, no operation is selected.
4. Select the dependent variable to adjust and move it to the **Selected** box. You must select one variable to adjust.
5. Click **OK** to return to the Adjustment Variables dialog box. For more information about this dialog box, see "Working with Adjustment Variables" on page 36.

Add a Reporting Variable

To add a reporting variable, select a variable from the **Variables to Assign** box and move it to the **Reporting Variables** box. You can only assign a numeric variable to the **Reporting Variables** role. These variables are not used to create the forecasts; however, they appear in the reports and plots.

Troubleshooting the Time ID Variable

In order to create the time series data, SAS Forecast Studio requires a specific format for the input data. For more information about how SAS Forecast Studio creates the time series data, see "Overview of Time Series Data" on page 7.

Common errors include the following:

- SAS Forecast Studio is unable to determine the time interval for the data. You must select a variable from the **Interval** drop-down list. By default, SAS Forecast Studio selects an interval when you assign a variable to the Time ID role.

- You must specify a time interval when the input data is transactional. This is when the input data is time stamped at no particular frequency.
- If the time interval is too small, then a large number of the observations will have missing values. For example, you select WEEK as the interval, but there is not enough data for most of the weeks. A larger interval, such as MONTH, should be selected.
- Assigning variables to incorrect roles is a common source of errors. **Note:** Incorrectly assigning variables to the BY variables role does not always generate errors. For example, if you forget to assign a BY variable, then that variable is not included in the hierarchy and the data is aggregated over that variable.

Configure the Hierarchy

If you are forecasting your data hierarchically, you can configure your hierarchy in the following ways:

- When you are creating a project, use the New Project wizard.
- After the project has been created, you change the hierarchy by using the Hierarchy Setup dialog box. To open this dialog box, select **Project** → **Hierarchy Setup**.

Note: This option is available only if you choose to forecast your data hierarchically.

To configure a hierarchy, complete the following steps:

1. In the **Levels** box, specify the order of the variables in the hierarchy. This order determines the hierarchy levels. The variables in this box are the variables that you assigned to the BY Variables role.

   Click **Preview** to preview the hierarchy. The Hierarchy Preview window opens. For more information, see "Previewing the Hierarchy" on page 40.
2. Select the **Reconcile hierarchy** check box to reconcile the hierarchy.
3. In the **Reconciliation Method** box, select the level in the hierarchy that contains the data that you want to forecast. You can choose from the following options:

   - **Bottom Up** - uses the data at the lowest level of the hierarchy to generate the forecasts. These forecasts are then used to generate the reconciled forecasts for the higher levels in the hierarchy.

     The Bottom Up method enables you to see any patterns (such as seasonality) in the data; however, because you are at the lowest level of the hierarchy, you can also have too much noise or randomness in the data. Also, these forecasts might fail because the data at the lowest level of the hierarchy can be sporadic or too sparse.

- **Middle Out** - adds up the data from the lower levels and then uses these values to generate the forecasts for the middle level. Because some hierarchies have more than one middle level, you need to specify the level that you want to use. The forecasts at the middle level are used to generate forecasts for both the higher and lower levels.
- **Top Down** - adds up the data from the lowest levels in the forecast and then uses this value to generate the forecasts. The forecast at the highest level is used to generate the forecasts for lower levels in the hierarchy.

  The Top Down method enables you to remove the excessive noise from the data at the lower levels of the hierarchy; however, you also might lose the pattern (such as the seasonality) in the forecast.

4. In the **Disaggregation Method** box, select the disaggregation method for the hierarchy. You can choose from the following options:

   - **Equal Split of the Difference** - reconciles forecast based on forecast deviations between levels.
   - **Forecast Proportions** - reconciles forecast based on historical mean. This is the default.

5. In the table, specify the statistic for aggregating data across the levels of the hierarchy. By default, all the dependent variables are listed in the table. In the Statistic column, select the aggregation statistic to use for each variable.

   For example, your data set contains the sales for a group of products. You might want to know the total sales for a category. In this case, you would choose **Sum** as the aggregation statistic. If your data contains the price of each product, you might want to know the average price for a product line. In this case, you would choose **Average** as the aggregation statistic.

6. If you are using the New Project Wizard, click **Next** to move to the next step. In the Hierarchy Setup dialog box, click **OK** to save your changes.

The following example shows how the forecasts are generated for the hierarchy Region > Product Category > Product Line > Product, based on the reconciliation method that you choose. Aggregation forecasts for the higher levels in the hierarchy are created based on the aggregation statistic that you select. Disaggregation forecasts for the lower levels in the hierarchy are created based on the disaggregation method that you select.

Table 4.1. Example: Forecasts for the Region > Product Category > Product Line > Product Hierarchy

Hierarchy Level	Reconciliation Method			
	Bottom Up	**Middle Out - Product Category**	**Middle Out - Product Line**	**Top Down**
Region	Aggregation Forecast	Aggregation Forecast	Aggregation Forecast	Forecast
Product Category	Aggregation Forecast	Forecast	Disaggregation Forecast	Disaggregation Forecast
Product Line	Aggregation Forecast	Aggregation Forecast	Forecast	Disaggregation Forecast
Product	Forecast	Disaggregation Forecast	Disaggregation Forecast	Disaggregation Forecast

Previewing the Hierarchy

After you have specified the order of the levels in the hierarchy in the New Project Wizard or Hierarchy Setup dialog box, you might want to see how the hierarchy will be arranged. (The levels are the variables that you assigned to the BY Variables role.) The Hierarchy Preview window enables you to view the levels in the hierarchy and the order in which they will appear. This preview is of the organizational hierarchy; it is not the actual hierarchy of the data. You cannot change the order of the levels from this window. To change the order of the hierarchy, see "Configure the Hierarchy" on page 38.

Enter Project Properties

In the New Project Wizard, you can specify a name and a description for your project.

To specify the project details, complete the following steps:

1. Specify the name of the project. By default, the project name is Project*n*, where *n* is the lowest available integer value. The project name must be a valid SAS name. The project name can be 32 characters long, and it must start with a letter (A-Z). Subsequent characters can be letters or numeric digits (0-9). Both upper- and lowercase letters are valid. For more information about SAS naming conventions, see *SAS Language Reference: Concepts*.
2. Provide a brief description of the project. A description helps other users better understand the function of your project. This description is optional.
3. Select the **Allow other users to view and edit this project** check box to share this project. This option is not selected by default. When this check box is unchecked, only the person who created the project can open it.
4. Click **Next**.

Note: If a project with the same name already exists, then you are prompted to overwrite the existing project or rename the new project. If you click **Yes**, then the existing project will be overwritten once you have finished creating the project. If you click **No**, then you must rename the new project before you can continue.

Perform Additional Steps

In the New Project Wizard, you can perform the following additional steps before forecasting your data:

- **Set forecasting options** - enables you to specify the forecasting options to use for the project. The Project Options dialog box opens. For more information, see "Viewing the Forecast Options for the Project" on page 69.
- **Create events** - enables you to add events for the project. The Events dialog box opens. For more information, see "Managing Events" on page 60.
- **Create filters** - enables you to add filters and exception rules to the project. The Filters dialog box opens. For more information, see "Managing Filters" on page 51.

Before you click **Finish**, specify whether to generate the forecasts when you finish or to save the SAS code to run later. By default, the forecasts are generated when you exit the wizard. You can monitor the progress of the forecasts from the status bar. The Summary window opens after SAS Forecast Studio is finished.

Viewing the Forecast Summary

Once SAS Forecast Studio has finished generating the forecasts for the project, the Summary window appears. This window gives you a read-only summary of the project run. It describes the last task that was run, the number of series forecasted, and the total time that elapsed. Click **Close** to close the window and view the project results.

Open a Project

You can open a project in the following ways:

- When you first invoke SAS Forecast Studio, you can open a project from the Welcome to SAS Forecast Studio window.
- After SAS Forecast Studio is running, you can manage your projects and open existing projects from the Projects dialog box. For more information, see "Managing Your Projects" on page 42.

  Note: You can have only one project open at a time in SAS Forecast Studio.

The Welcome to SAS Forecast Studio window opens by default when you start SAS Forecast Studio. From this window, you can create a new project or open an existing project.

You can choose from the following options:

- **Open an existing project** - enables you to select a project from the list. Projects in grey can only be opened by the person who created the project. To share this project, select the **Allow other users to view and edit this project** check box in the Properties window. For more information, see "View Project Properties" on page 47.

  To sort the projects, click either the Project or Date Updated column heading. To view more information about the project, select the project from the table. The Project Information box displays the project name, the description, the data set that was used, and who created the project.

- **Create a new project** - opens the New Project wizard. For more information, see "Create a New Project" on page 33.

After you have selected an option, click **OK**.

Managing Your Projects

You can manage your projects using the Projects dialog box. You can open this dialog box by selecting **File** → **Projects** or click  in the toolbar.

The project list displays the project name and the date that each project was last saved. To view more information about the project, select the project from the table. The Project Information box displays the project name, the description, the data set that was used, and who created the project.

From the Projects window, you can complete the following tasks.

- To create a new project, click **New**. The New Project Wizard opens. For more information, see "Create a New Project" on page 33.
- To open an existing project, select the project from the list and click **Open**. You can only have one project open at a time.

  Note: Projects in grey can only be opened by the person who created the project. To share this project, select the **Allow other users to view and edit this project** check box in the Properties window. For more information, see "View Project Properties" on page 47.

- To delete a project, select the project that you want to delete from the list and click **Delete**. You are asked to confirm your selection. If you are certain that you want to delete the project, click **OK**. You can only delete a project that you created.

Archive a Project

You can create an archive (or backup) of the current project. You might want to create a new archive before you generate new forecasts or add a model. After you have created the project, you can manage your project archives by selecting **File** → **Archives** or by clicking  in the toolbar.

To view the name and description for an archived project, select the project name from the list. The name and description information appears in the Archive Information box.

In the Archives dialog box, you can complete the following tasks:

- create an archived project
- restore an archived project
- delete an archived project

Create an Archived Project

You can create a new archive from the New Archive dialog box.

To create an archived project, complete the following steps:

1. Select **File** → **Archives** to open the Archives dialog box.
2. Click **New** in the Archive dialog box. The New Archive dialog box opens.
3. Specify a name for the archived project. The name you specify must be a valid SAS name. By default, the archive is named Archive*n* where *n* is the lowest available integer value.
4. Optionally, specify a description for the archived project.

   Note: You cannot specify a value in the **Date** box. This information is provided by the system when you create the archive.
5. Click **OK**. The archived project now appears in the list in the Archives dialog box.

Restore an Archived Project

To restore the current project to the selected archive, complete the following steps:

1. Select **File** → **Archives** to open the Archives dialog box.
2. Select the project name from the list and click **Restore**.

Note: Any changes that you have made to the current project are lost when you restore an archive, so you might want to create an archive of the current project before restoring an archive.

Delete an Archived Project

You can manage your filters by selecting **File** → **Archives**. The Archive dialog box opens.

To delete an archived project, select the project from the list in the Archives dialog box and click **Delete**. You are prompted to confirm this deletion. Click **OK** to delete the selected project.

Save a Project

To save the current project under a different name, select **File** → **Save Project As**. The Save As dialog box for your operating environment opens. Specify the filename and location for the project. After you have finished providing this information, you return to the SAS Forecast Studio workspace.

Close a Project

You can have only one project open in SAS Forecast Studio at a given time. To close an existing project, select **File** → **Close Projects**. After you have closed the current project, you can create a new project or open an existing project.

To close SAS Forecast Studio, select **File** → **Exit**. If you currently have a project open, a dialog box appears that asks you to confirm your exit. Click **Yes** to exit the application. Click **No** to return to your current project.

Update the Project Data

Before reforecasting the data, you might want to see if the project data has been updated since you opened the project.

To update the project data, complete the following steps:

1. Select **File** → **Update Project Data**. The New Data window opens.
2. Select the action that you want to perform and click **OK**.

SAS Forecast Studio updates the project data and reforecasts based on the option that you selected.

About the New Data Window

The New Data window opens when you select **File** → **Update Project Data** or when you open SAS Forecast Studio and the application detects updated data.

From the New Data window, you can determine what action is performed after the new data is read. You can choose from the following options:

- **Choose the best model from current model list(s)** - chooses the best model from all fitted models. This is the default.
- **Re-estimate parameters for the current forecast model** - re-estimates the parameters for the current model and then reforecasts the data.
- **Use existing parameters for the current forecast model** - reforecasts the data using the current parameters for the current model.
- **Cancel - do not use the new data** - does not include the new data.

After you have made your selection, click **OK** to reforecast the data.

View Series Properties

You can toggle on and off the properties information for a series. To view the properties for a series, select the series in the Hierarchy or Table tab, and select **View** → **Show/Hide Series Properties**.

By default, the properties information for a series is hidden. When you select this option, the Series Properties pane appears at the bottom of the **Hierarchy** and **Table** tabs.

You can see the following information for the selected series:

- name
- start date
- end date
- number of values
- number of missing values
- mean
- standard deviation
- minimum value
- maximum value

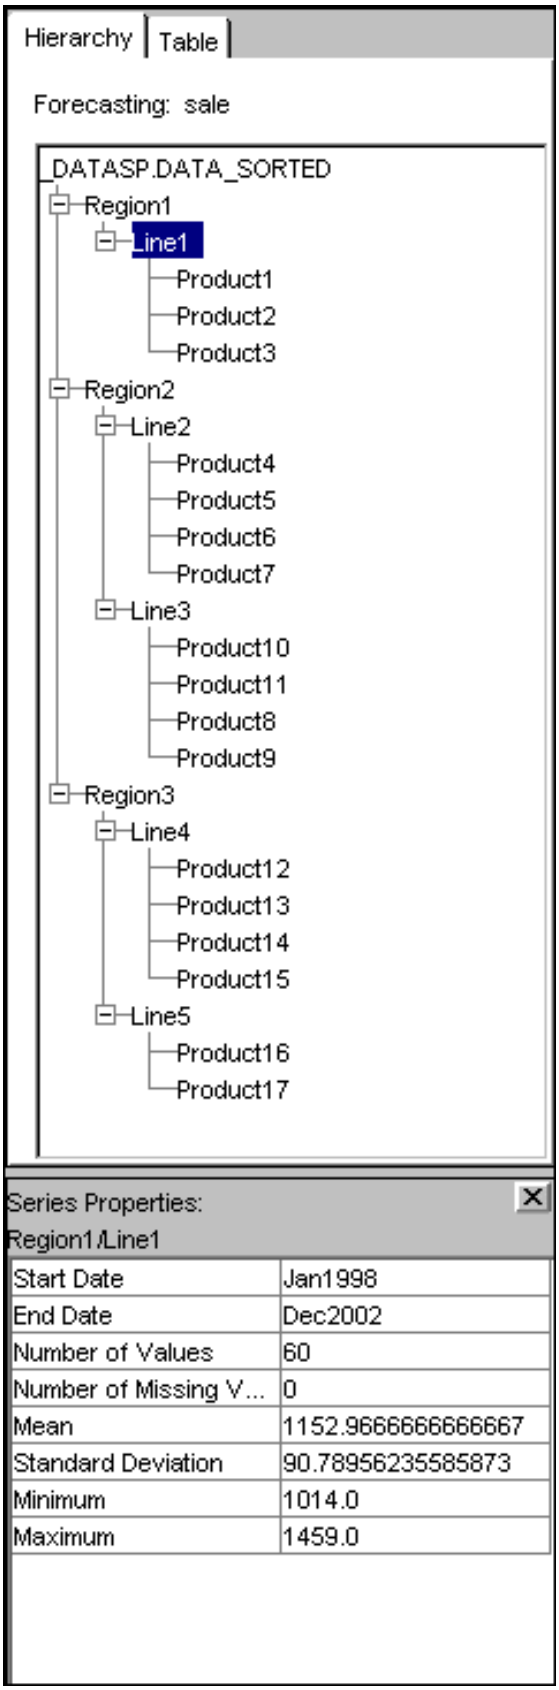

Figure 4.1. Hierarchy Tab with Series Properties Pane

View Project Properties

To view the project properties, select **File** → **Project Properties**. The Project Properties window opens.

In this window, you can view the project name, the location where the project is saved, input data set for the project, who created the project, and when it was created. These values are assigned when you created the project. You cannot change these project properties.

You can add a project description from this window. Type the description in the **Description** box. You can also select the **Allow other users to view and edit this project** check box to share this project. This option is not selected by default. When this check box is not selected, only the person who created the project can open it.

Then click **OK** to save your changes. To close the Project Properties window without saving your changes, click **Cancel**.

View the SAS Code

To view the code that SAS Forecast Studio generated, select **Tools** → **SAS Code**. The SAS Code window opens.

If you are viewing the code for the project, you can select the that code is displayed from the **Task** drop-down list. The available options include **Recreate Project**, **Reforecast Models**, **Reselect Models**, and **Refit Models**. For more information about these options, see "About the New Data Window" on page 44.

To include code for importing new data, select the **Include code to import data**. This option is not available if you select **Recreate Project** from the **Task** drop-down list.

You can expand and collapse sections of code by using the plus and minus signs next to the line numbers.

You might want to save this code to a file so that you can run the code in batch mode or in interactive SAS mode.

To save this code to a file, complete the following steps:

1. Click **Save As**. The Save As dialog box opens.
2. Select the location where you want to save this file and click **Save**.

To exit the SAS Code window, click **Close**.

View the SAS Log

To view the SAS log that SAS Forecast Studio creates when generating the forecast, select **Tools** → **SAS Log**. The SAS Log dialog box opens. You can review this log to see if there were any errors or warnings that occurred while the data was being processed.

You might want to save this information to a file so that you can refer to it later.

To save the SAS log, complete the following steps:

1. Click **Save As**. The Save As dialog box opens.
2. Select the location where you want to save this file and click **Save**.

To exit the SAS Log dialog box, click **Close**.

Note: If you need more information about a warning or error in the log, contact SAS Technical Support.

Chapter 5
Working with Filters

Chapter Contents

Chapter 5
Working with Filters

What Is a Filter?

You create filters to help manage your forecasts and to create exception rules. These exception rules help you identify forecasts that do not meet a specific criterion. After the filter is applied, series that are flagged as exceptions appear in red in the Hierarchy tab.

For an example of how to use a filter to create an exception rule, see "Example: Creating a Filter for MAPE" on page 52.

Managing Filters

SAS Forecast Studio can generate large numbers of forecasts. Filters help you focus on those forecasts that are less desirable. You can manage filters by using the Filters dialog box.

You can open the Filters dialog box in the following ways:

- When you are creating a project, you can create filters when you specify additional steps in the New Project Wizard. For more information, see "Perform Additional Steps" on page 41.
- Once you have created a project, you can manage your filters by selecting **Tools** → **Filters**.

The **Filter** list displays a list of the filters that you have created and their status. In the Filter Information box, you can view the name and description for the selected filter.

From the Filters dialog box, you can complete the following tasks:

- create a filter
- edit a filter
- delete a filter

Create or Edit a Filter

To add a new filter or edit an existing filter, complete the following steps:

1. Select **Tools** → **Filters**. The Filters dialog box opens.

2. To add a new filter, click **New**.

   To edit an existing filter, select the filter from the list and click **Edit**.

   The New Filter or Edit *Filter-name* dialog box opens.
3. Specify a name for the filter.
4. Optionally, specify a description.
5. If you want this filter to be used to flag exceptions, select the **This filter will be used to flag exceptions** check box. This option is selected by default.
6. In the first drop-down list, select the variable to evaluate.
7. In the second drop-down list, select the operation to use to evaluate the variable.
8. In the third drop-down list, select or type the value that the variable will be evaluated against.
9. Click **Add Row** if you need to add additional filter criteria.

   To remove a row, click  .

   Note: To clear all of the expressions, click **Reset**.

For an example, see "Example: Creating a Filter for MAPE" on page 52.

Delete a Filter

You can manage your filters by selecting **Tools** → **Filters**. The Filters dialog box opens.

To delete a filter, select the filter from the list in the Filters dialog box and click **Delete**.

To restore a deleted filter, click **Cancel** in the Filters dialog box.

Example: Creating a Filter for MAPE

In this example, you want to know which series has a Mean Absolute Percentage Error (MAPE) greater than 7.

To create this filter, complete the following steps:

1. Select **Tools** → **Filters**. The Filters dialog box opens.
2. Click **New**. The New Filter dialog box opens.
3. In the **Name** box, type **`MAPE`**.
4. In the **Description** box, type **`Filter to determine which series have a MAPE greater than 7`**.
5. Select the **This filter will be used to flag exceptions** check box.
6. In the first drop-down list, select **MAPE**.
7. In the second drop-down list, select **Greater than**.

8. In the third drop-down list, type **7**.

   The final window should look like the following:

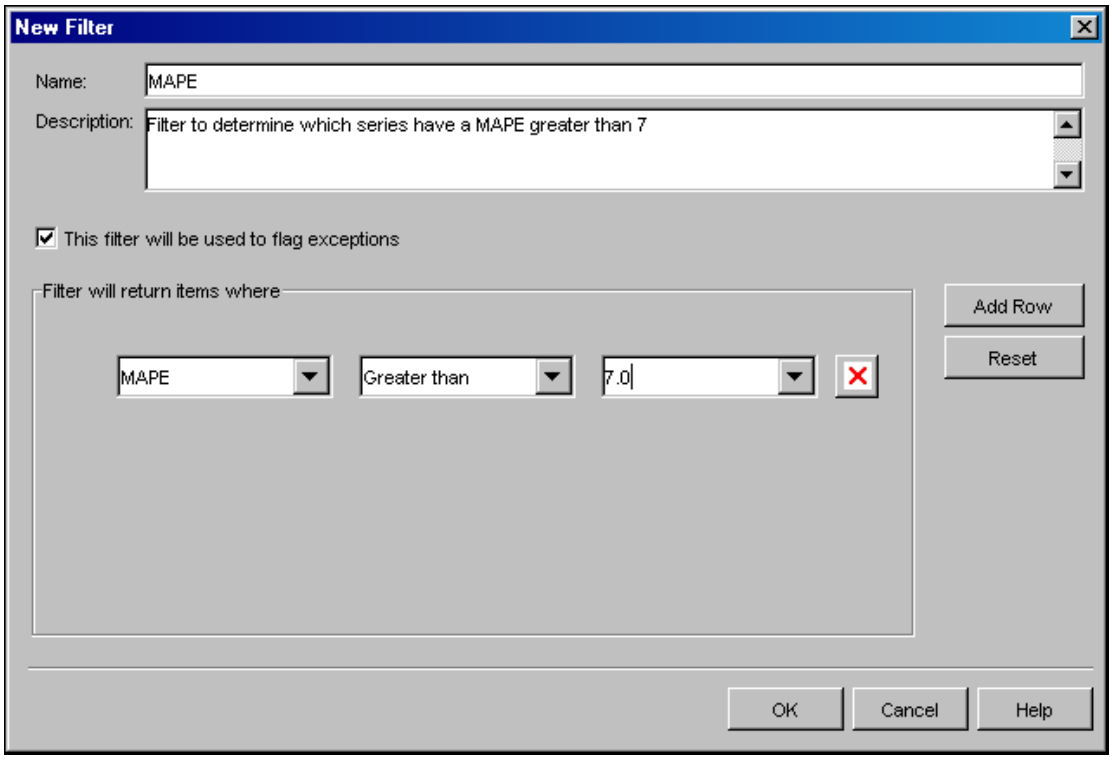

Figure 5.1. Completed New Filter Dialog Box for Filter Example

9. Click **OK** to add the filter to the list in the Filters dialog box.

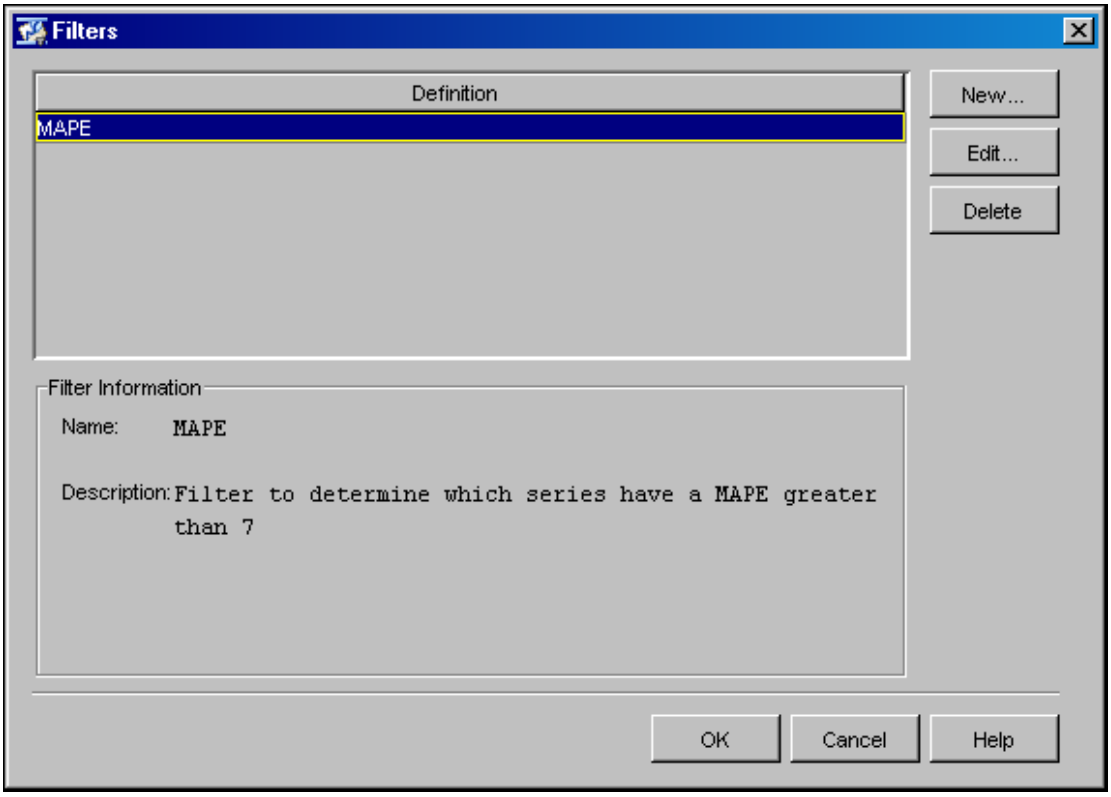

Figure 5.2. New Filter in Filters Dialog Box

10. Click **OK** in the Filters dialog box to apply the filter to the current project.

Since you selected the **This filter will be used to flag exceptions** check box, any series that has a MAPE greater than 7 appears in red in the **Hierarchy** tab and is identified with × in the **Table** tab.

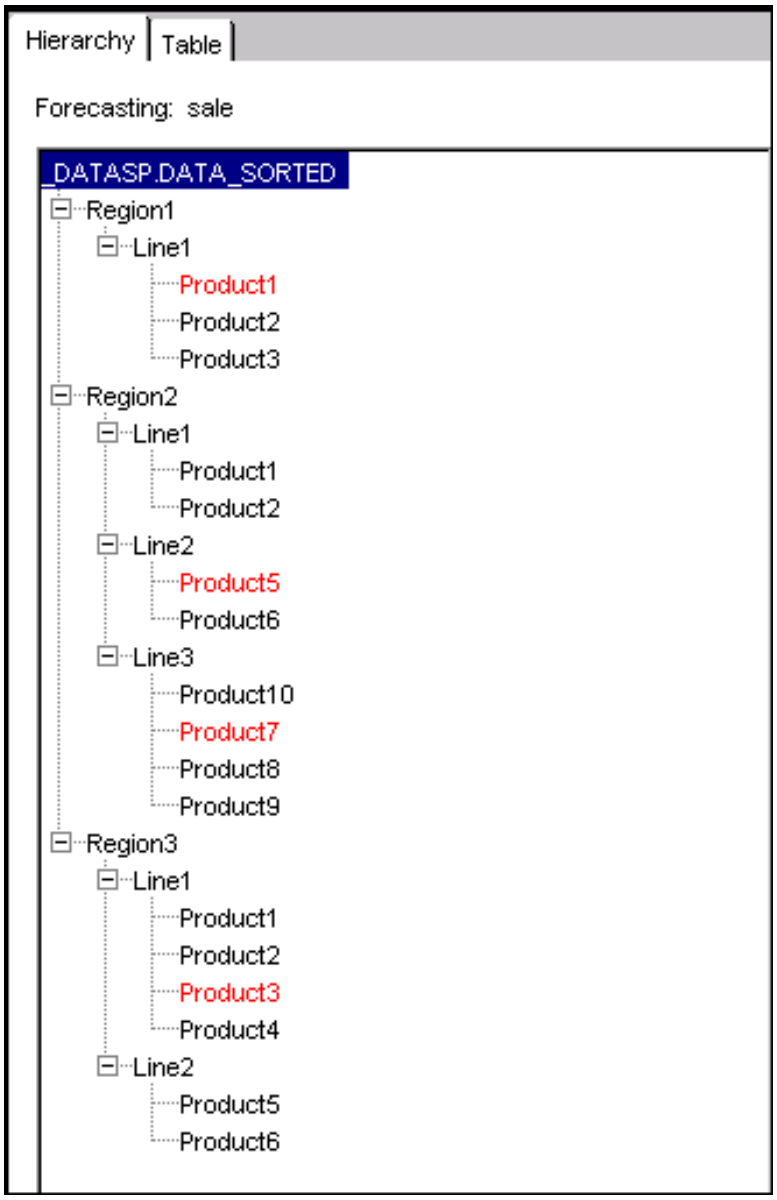

Figure 5.3. Part of Hierarchy Tab after Filter Is Applied

Chapter 6
Working with Events

Chapter Contents

Chapter 6
Working with Events

What Is an Event?

An event or calendar event is an incident that disrupts the normal flow of a process that generates the time series. Examples of events are holidays, retail promotions, and natural disasters. Defining an event enables you to model the effect that special events have on the dependent time series. When you define an event and apply it to a time series, it creates a dummy variable or an indicator variable based on the specified event type. The dummy variable indicates the occurrence of the event at any time period. This dummy variable is used as a regressor variable for time series modeling and forecasting.

For example, daily retail sales data follow a fairly steady pattern depending only on the day of the week except in the case of a special event such as holiday promotion. When you include events such as a one-time Christmas promotion, the forecasting model can predict the temporary increase in sales that occurs at those times and then the return to normal sales level afterward. Some yearly events that occur on a fixed date, such as Christmas, can be modeled as part of the regular seasonal model. Seasonal events that are most effectively modeled as regressors are those that occur on a different date each year such as Easter.

In SAS Forecast Studio, you define an event by specifying the event name, event type, a date or a time interval when the event occurs, and recurrence. The following table lists the event types that are available in SAS Forecast Studio.

Table 6.1. Description of Event Types

Event Type	Shape of Time Series	Description
Pulse		Temporary change in the magnitude of a time series process. The magnitude returns to the former level immediately after the change.
Level Shift		Persistent change in level of a time series process.
Ramp	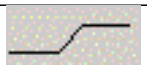	Persistent change in the trend or slope of a time series process.
Temporary Change	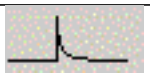	Temporary change in the magnitude. The magnitude decays to the former level after the change.
Combination	N/A	A combination of two or more of the simple event types.

Managing Events

The Events dialog box enables you manage your events.

You can open the Events dialog box in the following ways:

- When you are creating a project, you can create events when you specify additional steps in the New Project Wizard. For more information, see "Perform Additional Steps" on page 41.
- Once you have created a project, you can manage your events by selecting **Project** → **Events**.

The Event list displays a list of the events that you have created and their status. In the Event Information box, you can view the description for the selected event, the event type, and the current status of the event. The following options are available from the **Status** drop-down list: **Do not use**, **Try to use**, **Use if significant**, and **Force use**.

From the Events dialog box, you can do the following tasks:

- create a new event
- edit an existing event
- copy an event
- combine events
- delete an event

Create or Edit an Event

You can create a new event or edit an existing event, complete the following steps:

1. Select **Project** → **Events**. The Events dialog box opens.
2. If you are creating a new event, click **New**. If you are editing an existing event, select the event from the list and click **Edit**.

   The New Event or Edit *Event-name* dialog box opens.
3. Specify a name for the event. By default, the name for a new event is Event*n* where *n* is the lowest available integer value. The project name can be 32 characters long, and it must start with a letter (A-Z). Subsequent characters can be letters or numeric digits (0-9). Both upper- and lowercase letters are valid. For more information about SAS naming conventions, see *SAS Language Reference: Concepts*.
4. Optionally, specify a description for the event. Below the **Description** box is a time series plot that displays events and associated forecasts.
5. Click **Change** to specify an event type. For more information, see "Specify an Event Type" on page 61.

6. Specify the occurrence of the event. Click **Add** to add an occurrence. For more information, see "Add an Occurrence" on page 61. To remove the occurrence, select the occurrence name from the list and click **Remove**.
7. Click **Edit** to specify the recurrence. For more information, see "Specify Event Recurrence" on page 64.

   Note: You cannot specify the recurrence of an event if the occurrence dates are noncontiguous.
8. Optionally, click **View Code** to view the SAS code for the event.
9. Click **OK** to return to the Events dialog box. The new event appears in the list. For more information about this dialog box, see "Managing Events" on page 60.

Specify an Event Type

In SAS Forecast Studio, when you create a new event the default event type is Pulse. The Event Type dialog box enables you to change the event type to one of the following:

- **Pulse**
- **Level Shift**
- **Ramp**
- **Temporary Change**

After you have created two or more events, you can combine events. For more information, see "Combine Events" on page 65.

For more information about these event types, see "What Is an Event?" on page 59.

Add an Occurrence

You can specify the year, month, week, date, or hour that the event occurs. The **Occurrence** box displays a list of the occurrences for the event.

To add an occurrence, click **Add**. A dialog box opens to prompt you for more information about the occurrence. The dialog box that opens depends on the timestamp of your data.

- For daily data, select the day or date range that the event occurs in the Select Day dialog box.
- For hourly data, select the hour that the event occurs in the Select Hour dialog box.
- For weekly data, select the week that the event occurs in the Select Week dialog box.
- For monthly data, select the month that the event occurs in the Select Month dialog box.
- For all other time intervals, select the time period that the event occurs in the Select *Interval Types* dialog box.

Select the Day or Date Range

The Select Day dialog box enables you to specify the day(s) or date range for the event by completing the following steps:

1. Specify the year of the event.
2. Select a date, dates, or a date range in the calendar. Alternatively, you can select a predefined day from the list to the right of the calendar.
3. Click **Advanced** to specify the following advanced settings:
   - Indicate whether the event occurs on either the selected day or a number of periods before or after the selected day.
   - Specify the duration of the event in the **Duration** box.
4. Click **OK**.

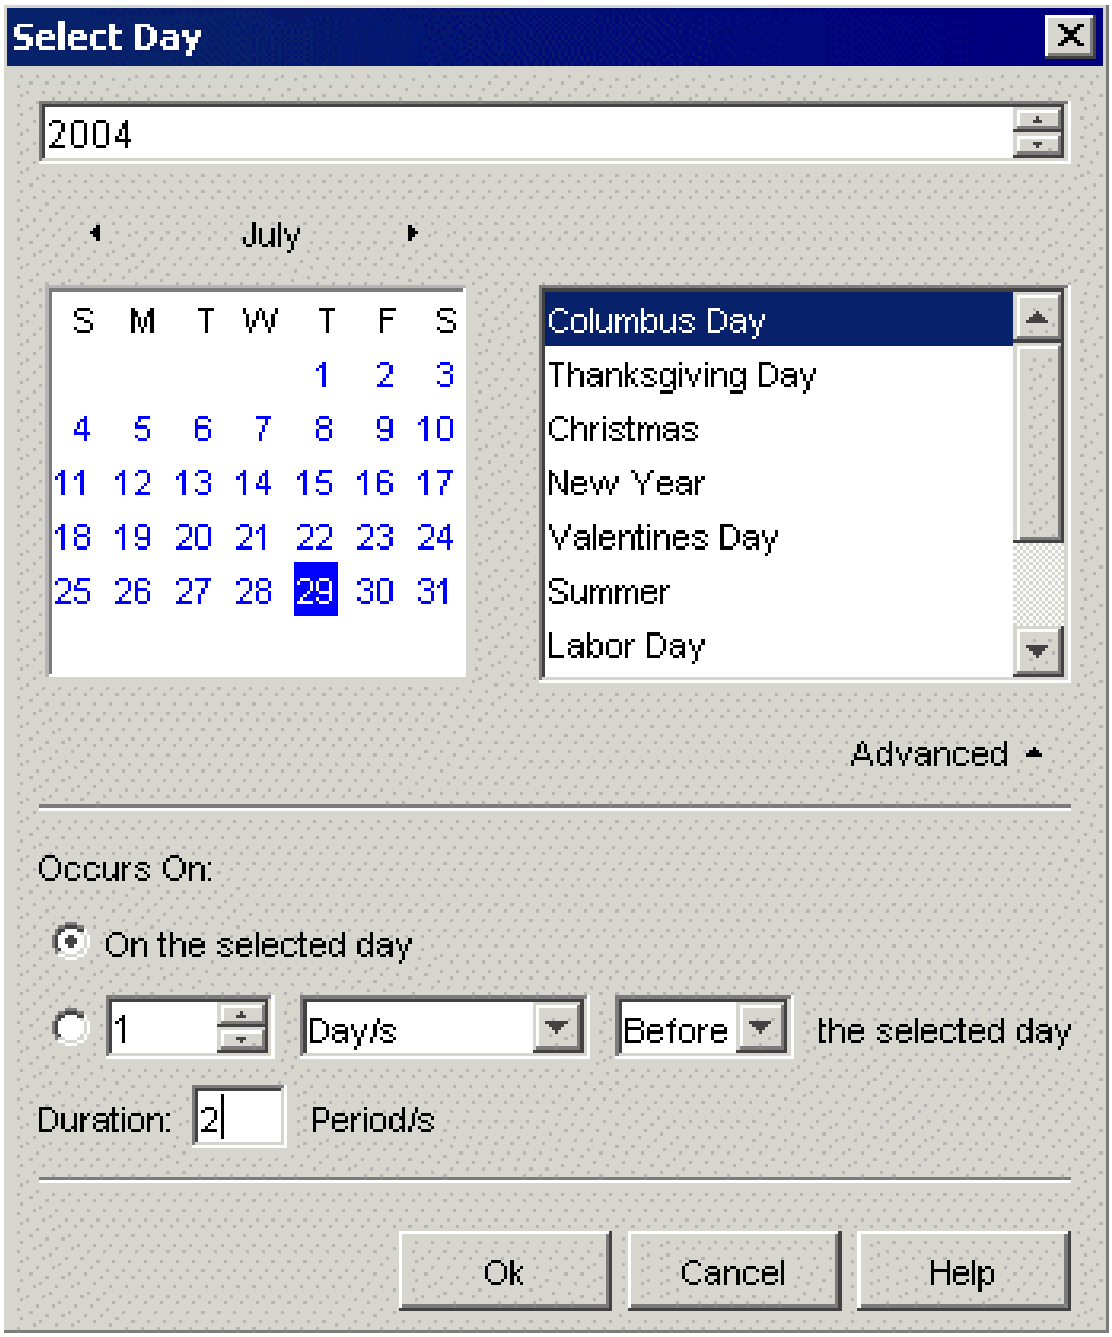

Figure 6.1. Select Day Dialog Box

Select the Hour

The Select Hour dialog box enables you to specify the hour of the event by completing the following steps:

1. Specify the year of the event. By default, the current year is used.

2. Select a date, dates, or a date range in the calendar. Alternatively, you can select a predefined day from the list to the right of the calendar.
3. Use the list to the right of the calendar to select the hour that the event occurs. If multiple days are selected, then the selected time applies to all the days.
4. Click **OK**.

Select the Week

The Select Week dialog box enables you to specify the week of the event by completing the following steps:

1. Specify the year of the event. By default, the current year is used.
2. Select the week or weeks from the list. Alternatively, you can select from the list of the event weeks that you have defined.
3. Click **OK**.

Select the Month

The Select Month dialog box enables you to specify the month of the event by completing the following steps:

1. Specify the year of the event. By default, the current year is used.
2. Select the month or months from the list.
3. Click **OK**.

Select the Time Period

The Select *Interval Type* dialog box enables you to create an event from all the time periods in the historical data, plus future time periods equal to three times the horizon. You can select more than one period. After you have selected the periods to include as occurrences, click **OK**.

Specify Event Recurrence

The Recurrence dialog box enables you to specify the event recurrence by completing the following steps:

1. Indicate when the event recurs by choosing one of the items from the **Recurrence** drop-down list. The options available from the drop-down list depend on the interval of your data. For a list of the available options, see the PROC HPFEVENTS documentation in the *SAS High-Performance Forecasting User's Guide*.
2. If available, specify the end period for the recurrence. You can end the recurrence after a specified number of occurrences or at a specified date and time.
3. Click **OK**.

Copy Events

To copy an event, complete the following steps:

1. Select **Project** → **Events**. The Events dialog box opens.
2. Select the event that you want to copy in the **Events** list.
3. Click **Copy**.

The copied event is added to the list of events. By default, this event is named *original-name*Copy*n*, where *original-name* is the name of the event that you copied and *n* is the lowest available integer value.

You can edit the event properties by using the Edit *Event-name* dialog box.

Combine Events

To combine events, complete the following steps:

1. Select **Project** → **Events**. The Events dialog box opens.
2. Select the events that you want to combine in the **Events** list.
3. Click **Combine**. The New Combination Event or Edit *Event Name* dialog box opens.
4. Specify a name for the new event.
5. Optionally, provide a brief description for the new event.
6. Select from the table the events that you are combining. This table displays the events that you are combining and their properties (such as name, event type, start date, and recurrence).

   To add an event, click **Add**. From the Select Event dialog box, choose the event that you want to add to the list and click **OK**.

   To delete an event, select the event from the list and click **Remove**.
7. Optionally, click **View Code** to view the SAS code for the combined event.
8. Click **OK**.

Select an Event

From the Select Event dialog box, select the event or events that you want to add to the list of events. Click **OK** to return to the Combine Events dialog box.

Delete an Event

To delete an event, complete the following steps:

1. Select **Project** → **Events**. The Events dialog box opens.
2. Select one or more events in the **Events** list.
3. Click **Delete**.

You are not prompted to confirm the deletion. To restore a deleted event, click **Cancel** in the Events dialog box.

Chapter 7
Generating Forecasts

Chapter Contents

Chapter 7
Generating Forecasts

Reforecasting the Data

After you initially create the project, you might change some parameters or create a new model. After these changes, you might want to reforecast your data. You can reforecast the current series or all of the series in the project.

To reforecast your data, complete the following steps:

1. To reforecast the current series, select the series in the **Hierarchy** tab or **Table** tab and select **Series** → **Reforecast** or click  in the toolbar. To reforecast all the series in the project, select **Project** → **Reforecast** or click  in the toolbar. The Reforecast dialog box opens.
2. Select how to reforecast the data. The following options are available:

   - **Automatic model generation and selection** - generates the new models, fits each model, and automatically chooses the best one. **Note:** This option is available only if you are reforecasting a series.
   - **Choose the best model from current model list(s)** - chooses the best model from all fitted models. This is the default.
   - **Re-estimate parameters for the current forecast model** - re-estimates the parameters for the current model and then reforecasts the data.
   - **Use existing parameters for the current forecast model** - reforecasts the data using the current parameters for the current model.

3. If you are reforecasting the entire project, you can select the **Use updated data if available** check box. SAS Forecast Studio checks to see if the data source has been updated, and if so, then uses the updated data to generate the forecast.
4. Click **OK** to reforecast your data.

   The Reforecast dialog box closes, and SAS Forecast Studio generates the forecasts. When finished, the Summary window opens. For more information, see "Viewing the Forecast Summary" on page 41.

Viewing the Forecast Options for the Project

All project options are initially set to a default value. To change an option for all series in the project, use the Project Options dialog box. To view the forecast options for a project, select **Project** → **Forecasting Options**. The Project Options dialog box opens.

The project options are divided into the following tabs:

- **Automatic Model Selection** includes the options for generating or selecting the candidate models for your project.
- **Other** includes options for data preparation.

An option is used in the project if the check box in the column to the left of the Value column is selected. The value for that option is displayed in the Value column. You can change these option values by clicking in the Value column for the option that you want to change.

To save your changes and exit the Project Options window, click **OK**. After you have made your changes, SAS Forecast Studio reforecasts your data.

Using the Options on the Automatic Model Selection Tab

The following options are available from the **Automatic Model Selection** tab in the Project Options window:

- **Create ARIMA model** includes the ARIMA models that SAS Forecast Studio generates. This option is selected by default.
- **Create ESM** includes the exponential smoothing models (ESM) that SAS Forecast Studio generates. This option is selected by default.
- **Create UCM** includes the unobserved component models (UCM) that SAS Forecast Studio generates.
- **Fit existing model selection list** enables you to select an existing model from the model selection list.
- **Detect Outliers** specifies whether SAS Forecast Studio should detect outliers in the data.

Using the Options on the Other Tab

The following options are available from the **Other** tab in the Project Options window:

Missing Value Interpretation
replaces the missing values in the data with a value that you specify. You can choose from the following options:

- **Missing** - the missing values are set to missing.
- **Average** - the missing values are set to the accumulated average value. This is the default value.
- **Minimum** - the missing values are set to the accumulated minimum value.

- **Median** - missing values are set to the accumulated median value.
- **Maximum** - missing values are set to the accumulated maximum value.
- **First** - missing values are set to the accumulated first nonmissing value.
- **Last** - missing values are set to the accumulated last nonmissing value.
- **Previous** - missing values are set to the previous accumulated nonmissing value. Missing values at the beginning of the accumulated series remain missing.
- **Next** - missing values are set to the next accumulated nonmissing values. Missing values at the end of the accumulated series remain missing.

Leading/Trailing Zeroes

specifies how beginning and/or ending zero values are interpreted in the accumulated time series. You can choose from the following options:

- **None** - the beginning and/or ending zeros are unchanged. This is the default value.
- **Left** - beginning zeros are set to missing.
- **Right** - ending zeros are set to missing.
- **Both** - both beginning and ending zeros are set to missing.

Leading/Trailing Missing Values

specifies how missing values are removed from the accumulated time series. You can choose from the following options:

- **None** - the missing values are kept.
- **Left** - the beginning missing values are removed.
- **Right** - the ending missing values are removed.
- **Both** - both beginning and ending missing values are removed. This is the default.

End periods to omit

specifies the number of observations before the end of the data that the multistep forecasts are to begin. You cannot specify a negative value. The default value is 0.

Specify start date

specifies the datetime value at the beginning of the data. If the first value of the time ID variable is greater than the start date, the series is prepended with missing values. If the first value of the time ID variable is less than the end date, the series is truncated.

Specify end date

specifies the datetime value at the end of the data. If the last value of the time ID variable is less than the end date, the series is extended with missing values. If the last value of the time ID variable is greater than the end date, the series is truncated.

Minimum observations for a trend model

specifies minimum observations for a trend model.

Minimum observations for a non-Mean model

specifies observations for a non-mean model.

Horizon

specifies the number of periods into the future for which multi-step forecasts are made. The larger the horizon value, the larger the prediction error variance at the end of the horizon. If SAS Forecast Studio detects seasonality in the data, then the default value is determined by the seasonality. If no seasonality is detected, then the default value is 12.

Do not allow negative forecasts

specifies whether negative forecasts are allowed. If this option is selected, then the series is log transformed. This option is selected by default.

Intermittency Test

specifies a number greater than 1 that is used to determine whether or not a time series is intermittent. This option is selected by default, and the default value is 1.25.

Seasonality Test

specifies the options for the seasonality test. You can specify no seasonality test or specify a significance probability value for the test. Series with strong seasonality have small test probabilities. A significance probability value of 0 always implies seasonality. A significance probability value of 1 always implies no seasonality. The default is a significance probability value of .01.

Holdout sample (%)

specifies the holdout sample as a percent of the series. This value cannot be negative.

Maximum holdout sample (periods)

specifies the holdout sample as a number of periods. This value cannot be negative.

Selection Criterion

specifies the selection criterion (also called the statistic of fit) for the model. You can select only one value from the drop-down list. The following values are available:

- **AADJRSQ** - Amemiya's Adjusted R-Square
- **ADJRSQ** - Adjusted R-Square
- **AIC** - Akaike Information Criterion

- **APC** - Amemiya's Prediction Criterion
- **DFE** - Degrees of Freedom for Error
- **GMAPE** - Geometric Mean Percent Error
- **GMAPPE** - Geometric Mean Predictive Percent Error
- **GMASPE** - Geometric Mean Symmetric Percent Error
- **GMRAE** - Geometric Mean Relative Absolute Error
- **MAE** - Mean Absolute Error
- **MAPE** - Mean Absolute Percent Error
- **MAPPE** - Symmetric Mean Absolute Predictive Percent Error
- **MAXERR** - Maximum Error
- **MAXPE** - Maximum Percent Error
- **MAXPPE** - Maximum Predictive Percent Error
- **MAXRE** - Maximum Relative Error
- **MAXSPE** - Maximum Symmetric Percent Error
- **MDAPE** - Median Percent Error
- **MDAPPE** - Median Predictive Percent Error
- **MDASPE** - Median Symmetric Percent Error
- **MDRAE** - Median Relative Absolute Error
- **ME** - Mean Error
- **MINERR** - Minimum Error
- **MINPE** - Minimum Percent Error
- **MINPPE** - Minimum Predictive Percent Error
- **MINRE** - Minimum Relative Error
- **MINSPE** - Minimum Symmetric Percent Error
- **MPE** - Mean Percent Error
- **MPPE** - Mean Predictive Percent Error
- **MRAE** - Mean Relative Absolute Error
- **MRE** - Mean Relative Error
- **MSE** - Mean Square Error
- **MSPE** - Mean Symmetric Percent Error
- **N** - sample size
- **NMISSA** - Number of Missing Actual Values
- **NMISSP** - Number of Missing Predicted Values
- **NOBS** - Number of Observations
- **NPARMS** - Number of Parameters
- **RMSE** - Root Mean Square Error
- **RSQUARE** - R-Squared
- **RWRSQ** - Random Walk R-Square
- **SBC** - Schwarz Bayesian Information Criterion
- **SMAPE** - Symmetric Mean Absolute Percent Error
- **SSE** - Sum of Square Error

- **SST** - Total Sum of Squares
- **TSS** - Total Corrected Sum of Squares for the dependent variable
- **UMSE** - Unbiased Mean Square Error
- **URMSE** - Unbiased Root Mean Square Error

Confidence Limits

specifies the confidence level for the series. By default, this confidence level is 95%.

Create index files

indexes the series in the project so that it is faster to navigate between series in the application. If you only want to run the project code in batch code, do not select this option.

Working with Overrides

Add Overrides

You can add overrides by using the Override Calculator.

The final forecast changes based on the override value. Series that have an associated override are idenitified in the **Hierarchy** tab with the Override icon.

You might want to add a note to explain each override. For more information, see "Adding Notes" on page 76.

Working with the Override Calculator

You can open the Override Calculator in the following ways:

- From the Forecasting View, click **Override Calculator** under the Data Table.
- Select a series in the **Hierarchy** or **Table** tab, and select **Series** → **Override Calculator**.

Note: You can only override future values; for this reason, the Override Calculator is not available when the value for the **Horizon** is less than the number of **End periods to omit**. For more information about these options, see the **Other** tab in the Project Options window. For more information, see "Viewing the Forecast Options for the Project" on page 69.

When the Override Calculator opens, the table contains the reconciled forecast (or statistical forecast if the data was not forecast hierarchically), manual override, and final forecast for all periods in the horizon.

To create an override, complete the following steps:

1. Select the time periods for the override. You can select multiple cells for any one row.

2. Specify the changes to make to the selected values. You can choose from the following options.

   - Select the first radio button to add the overrides on a per period basis. Using the drop-down lists and text box, specify the number of units or percentage to increase or decrease the current value. In the first drop-down list, select **Increase** or **Decrease**. In the text box, specify the value and in the second drop-down list, specify whether this value is in units or percent.

     The value of the override is calculated based on the row you select.

     – If you select a cell in the Reconciled Forecast (or Statistical Forecast) row, then the override values are based on the specified increase or decrease over the reconciled forecast.
     – If you select a cell in the Manual Override row, then the new override value is calculated based on the specified increase or decrease over the existing manual override values.

     For example, you want to create an override that increases the current value by 10%. Select a cell in the Manual Override row. In the first drop-down list, select **Increase**. In the text box, type `10`, and in the second drop-down list, select **%**. The existing manual override values will increase by 10%.
   - Select **Set to** if you want to apply a single override value to all the selected periods. Specify this override value in the text box. How the override value is applied depends on the option that you choose. You can choose from the following options.

     – **Set each selected cell to this value**
     – **Spread this value equally across selected cells**
     – **Set proportionally across selected periods**

3. Click **Apply** to apply the override.

Remove Overrides

You can remove the manual overrides for the project or the series.

- To remove the overrides for the project, select **Project** → **Remove Overrides**.
- To remove the overrides for the selected series, select **Series** → **Remove Overrides**.

In the warning that appears, click **OK** to remove the overrides. Click **Cancel** to keep the current overrides.

Adding Notes

From the Forecasting View, you can add notes for the statistical forecast or manual overrides.

Note: Because detached windows are read-only, you cannot add a note if the Forecasting View is detached.

To enter a note, complete the following steps:

1. Expand the Notes section of the window.
2. In the text box, type your note. This note cannot exceed 2000 characters.

Series that have an associated note are idenitified in the **Hierarchy** tab and Data Table with the Note  icon. The notes persist when you close the project.

Chapter 8
Working with Models

Chapter Contents

Chapter 8
Working with Models

Create a New Model

To create a new model, complete the following steps:

1. Select **Project** → **Add Model to All Series**. The New Model dialog box opens.
2. When adding a model you can select from the following options:

   - **Select from model repository** - enables you to select a new model from the model repository.
   - **ARIMA** - analyzes and forecasts equally spaced univariate time series data, transfer function data, and intervention data using the AutoRegressive Integrated Moving-Average (ARIMA) model.
   - **Subset (factored) ARIMA** - is similar to the ARIMA model, except that it uses a general specification of the options. In this model, you can specify which lags have parameters.
   - **Exponential smoothing** - uses exponential smoothing models. You can choose from many types of exponential smoothing models.
   - **Unobserved components** - analyzes and forecasts equally spaced univariate time series data using unobserved component models or state space models. These models can include events and independent variables.
   - **Intermittent demand** - is used for time series that have a large number of values that are zero or other constant values. Intermittent time series occur when the demand for an item is intermittent. Because many time series models are based on weighted-summations of past values, they bias the forecast toward zero; so these models will not work for intermittent time series data.
   - **Multiple regression** - creates a multiple linear regression model with autocorrelated errors. You can specify which independent variables and events to include in the model.
   - **Moving average** - creates a specific type of simple moving average model. In the moving average model, you can specify the number of periods for the moving average.
   - **Curve fitting** - creates a curve fitting model that enables you to identify trends and relationships in your time series data.
   - **Random walk** - creates a specific type of ARIMA (0,1,0) model. In the random walk model, you can specify whether to include the drift and seasonal terms in the model.

3. Click **OK** to create the model that you selected. The model specification dialog box for the selected model opens. After you create the model, it is fit to every series in the project.

Select a New Model from Repository

To select a new model for the statistical forecast from the model repository, select the **Select model from repository** option in the New Model dialog box. The Select Model from Repository dialog box opens.

Select the model that you want to use from the list and click **OK**.

When you select a model, the model name, description, and system name appear in the Model Information box. You can sort the models alphabetically by clicking the Model column heading.

Create an ARIMA Model

To open the New ARIMA Model window, select the **ARIMA** option in the New Model dialog box.

From the New ARIMA Model window, you can specify the autoregressive and moving average polynomials of an ARIMA model. You can also optionally specify one or more input variables. The input variables can be any of the following:

- variables from the input data set
- predefined variables, such as linear or quadratic trends
- events that you have created

First, specify a name and optionally a description for the new model. The Details box lists the system name for the model. You cannot edit this text.

Use the selection pane on the left to navigate among the following options:

- specification
- inputs
- estimation

To view the SAS code for the model, click **View Code**.

For more information about the ARIMA model, see the HPFARIMASPEC procedure in the *SAS High-Performance Forecasting User's Guide*.

Specification Options for ARIMA Models

In the New ARIMA model window, select **Specification** in the selection pane to specify a basic nonseasonal (ARIMA(p, d, q)) or seasonal (ARIMA(p, d, q)(P, D, Q)) model.

Using these specification options, you can specify the following:

- the transformation to apply to the dependent series from the **Dependent Transformation** drop-down list. You can select from the following transformations:
  - no transformation
  - logarithmic transformation
  - square-root transformation
  - logistic transformation
  - Box-Cox transformation where the parameter number is between -5 and 5. Specify this parameter value in the Box-Cox Transformation window. For more information, see "About the Box-Cox Transformation Window" on page 81.

  Note: All transformation values (except no transformation) require that the dependent series is positive.
- the autogressive (p and P), differencing (d and D), and moving average (q and Q) orders for your model.

  Specify the nonseasonal and seasonal autogressive orders by specifying the values of p and P. Specify the nonseasonal and seasonal differencing orders by specifying the values of d and D. Specify the nonseasonal and seasonal moving average orders by specifying the values of q and Q.

  The ARIMA model has nonseasonal and seasonal components (p d q)(P D Q). If you set p=4, then you are implying autoregressive orders (1 2 3 4) in the nonseasonal component of the model. For a simple nonseasonal ARIMA model, you do not have to specify the orders for the seasonal component (P D Q). The season length (such as monthly, daily, etc.) is implied by the Time ID variable in the project.

  Note: The seasonal options are only available if the data is seasonal.
- whether an intercept term is included in the model. To include an intercept term, select the **Intercept** check box.

  Note: If you specify a differencing order, then a dialog box appears asking if you want to suppress the intercept. It is recommended that you suppress the intercept when a differencing order is specified.

About the Box-Cox Transformation Window

The Box-Cox Transformation window appears when you select Box-Cox as the transformation to apply to the dependent series for the model. In the **Set parameter** box, specify the Box-Cox transformation where the parameter number is between -5 and 5. The default value is 0. Click **OK** to return to the model specification dialog box.

Inputs for ARIMA Models

You can specify independent variables, predefined variables, events, and outliers as inputs to the model.

In the New ARIMA model window, select the input type in the selection pane. Each variable, event, or outlier is listed in the table. You can specify whether to use the variable, event, or outlier or ignore it.

To specify the input series options for the independent variable or event, select the variable or event from the table and click **Options**. The Input Series Options dialog box opens. For more information, see "Specifying the Input Series Options" on page 93.

To create a new event, click **New Event** on the Events page. After the event is created, it is added to the list of events and used for the current model. It also appears in the list of events for the project, but it is not used at the project level.

Estimation for ARIMA Models

In the New ARIMA model window, select **Estimations** in the selection pane. You can specify the following estimation options:

- **Method** - specifies the estimation method to use. The following methods are available: maximum likelihood (ML), unconditional least-squares (ULS), and conditional least-squares (CLS).
- **Convergence criterion** - specifies the convergence criterion. Convergence is assumed when the largest change in the estimate for any parameter is less than the specified value.
- **Number of iterations** - specifies the maximum number of iterations allowed.
- **Delta** - specifies the perturbation value for computing numerical derivatives.
- **Singularity criterion** - specifies the criterion for checking singularity. If a pivot of a sweep operation is less than the specified value, the matrix is deemed singular.

You can also select the **Restrict parameters to stable values** option. This option specifies that the autoregressive and moving-average parameter estimates for the noise part or disturbance polynomial of the model are restricted. These estimates are restricted to the stationary and invertible regions, respectively.

Create a Subset (Factored) ARIMA Model

To open the New Subset ARIMA Model window, select the **Subset (factored) ARIMA** option in the New Model dialog box.

From the New Subset ARIMA Model window, you can specify a general ARIMA model. You can specify the autoregressive and moving average polynomials of arbitrary complexity. You can also specify a general differencing order. The Subset (Factored) ARIMA model specification window is similar to the ARIMA

Specification window, except that it uses a more general specification of the ARIMA options.

Specify a name and optionally a description for the new model. The Details box lists the system name for the model. You cannot edit this text.

Use the selection pane on the left to navigate among the following options:

- specification
- inputs
- estimation

To view the SAS code for the model, click **View Code**.

For more information about the Subset ARIMA model, see the HPFARIMASPEC procedure in the *SAS High-Performance Forecasting User's Guide*.

Specification Options for Subset ARIMA Models

In the New Subset ARIMA model window, select **Specification** in the selection pane to access the following options:

- the transformation to apply to the dependent series from the **Dependent transformation** drop-down list. You can select from the following transformations:
  - no transformation
  - logarithmic transformation
  - square-root transformation
  - logistic transformation
  - Box-Cox transformation where the parameter number is between -5 and 5. Specify this parameter value in the Box-Cox Transformation window. For more information, see "About the Box-Cox Transformation Window" on page 81.
- whether an intercept is included in the model. If you specify a differencing term, then a dialog box appears asking if you want to suppress the intercept.
- the autogressive (p), differencing (d), and moving average (q) orders for your model.

  Use the following syntax for each of these options:

d=*order*
d= *(order1, order2, ...)*
specifies the differencing orders for the dependent series. For example, d= (1 12) specifies that the series be differenced using the operator $(1 - B)(1 - B^{12})$. The differencing orders can be positive integers.

p= *order*
p= (*lag, ..., lag*) **...** (*lag, ..., lag*)
p= (*lag, ..., lag*)<s_1> **...** (*lag, ..., lag*)<s_k>
specifies the autoregressive part of the model. By default, no autoregressive parameters are fit.

p=(l_1, l_2, ..., l_k) defines a model with autoregressive parameters at the specified lags. p= *order* is equivalent to P=(1, 2, ..., *order*).

A concatenation of parenthesized lists specifies a factored model. For example, p=(1,2,5)(6,12) specifies the autoregressive model

$$(1 - \phi_{1,1}B - \phi_{1,2}B^2 - \phi_{1,3}B^5)(1 - \phi_{2,1}B^6 - \phi_{2,2}B^{12})$$

Optionally, you can specify *multipliers* after the parenthesized lists. For example, p=(1)(1)12 is equivalent to P=(1)(12), and P=(1,2)4(1)12(1,2)24 is equivalent to p=(4,8)(12)(24,48). These multipliers can either be positive integers or they can be "s," indicating a placeholder that will be substituted later.

q= *order*
q= (*lag, ..., lag*) **...** (*lag, ..., lag*)
q= (*lag, ..., lag*)<s_1> **...** (*lag, ..., lag*)<s_k>
specifies the moving-average part of the model. By default, no moving-average parameters are fit.

The manner of specification of the moving-average part is identical to the specification of the autoregressive part described in the p= option.

Inputs for Subset ARIMA Models

You can specify independent variables, predefined variables, and events as inputs to the model.

In the New Subset ARIMA Model window, select the input type in the selection pane. Each variable or event is listed in the table. You can specify whether to use the variable or event or ignore it.

To specify the input series options for the independent variable or event, select the variable or event from the table and click **Options**. The Input Series Options dialog box opens. For more information, see "Specifying the Input Series Options" on page 93.

To create a new event, click **New Event** on the Events page. After the event is created, it is added to the list of events and used for the current model. It also appears in the list of events for the project, but it is not used at the project level.

Estimation for Subset ARIMA Models

In the New Subset ARIMA model window, select **Estimations** in the selection pane. You can specify the following options:

- **Method** - specifies the estimation method to use. The following methods are available: maximum likelihood, unconditional least-squares, and conditional least-squares.
- **Convergence criterion** - specifies the convergence criterion. Convergence is assumed when the largest change in the estimate for any parameter is less than the specified value.
- **Number of iterations** - specifies the maximum number of iterations allowed.
- **Delta** - specifies the perturbation value for computing numerical derivatives.
- **Singularity criterion** - specifies the criterion for checking singularity. If a pivot of a sweep operation is less than the specified value, the matrix is deemed singular.

You can also select the **Restrict parameters to stable values** option. This option specifies that the autoregressive and moving-average parameter estimates for the noise part of the model are restricted. These estimates are restricted to the stationary and invertible regions, respectively.

Create an Exponential Smoothing Model

To open the New Smoothing Model window, select the **Exponential smoothing** option in the New Model dialog box.

In the New Smoothing Model window, you can create an exponential smoothing model. For more information about the Exponential Smoothing model, see the HPFESMSPEC procedure in the *SAS High-Performance Forecasting User's Guide*.

To create an exponential smoothing model, complete the following steps:

1. Specify a name and optionally a description for the new model. The Details box lists the system name for the model. You cannot edit this text.
2. Select from the **Smoothing method** drop-down list the forecasting method to be used to forecast the time series. The following forecasting models are available:
   - **Simple** - Simple (Single) Exponential Smoothing. This is the default. The model equation for simple exponential smoothing is as follows:

     $$Y_t = \mu_t + \epsilon_t$$

- **Double** - Double (Brown) Exponential Smoothing
  The model equation for double exponential smoothing is as follows:

  $$Y_t = \mu_t + \beta_t t + \epsilon_t$$

- **Linear** - Linear (Holt) Exponential Smoothing
  The model equation for linear exponential smoothing is as follows:

  $$Y_t = \mu_t + \beta_t t + \epsilon_t$$

- **DampTrend** - Damped Trend Exponential Smoothing
  The model equation for damped-trend linear exponential smoothing is as follows:

  $$Y_t = \mu_t + \beta_t t + \epsilon_t$$

- **Best** - Best Candidate Smoothing Model
- **BestN** - Best Candidate Nonseasonal Smoothing Model
- **Bests** - Best Candidate Seasonal Smoothing Model
  The model equation for seasonal exponential smoothing is as follows:

  $$Y_t = \mu_t + s_p(t) + \epsilon_t$$

  For more information about these models, see the "Forecast Process" chapter in the *SAS High-Performance Forecasting User's Guide*.

3. Select from the **Dependent Transformation** drop-down list the dependent series transformation to be applied to the time series. You can select from the following transformations:

   - no transformation. This is the default.
   - logarithmic transformation
   - square-root transformation
   - logistic transformation
   - Box-Cox transformation where the parameter number is between -5 and 5. Specify this parameter value in the Box-Cox Transformation window. For more information, see "About the Box-Cox Transformation Window" on page 81.
   - an automatically chosen transformation based on the model selection criteria.

4. Select the selection criterion to use to select from several candidate models. This option is available only when the method is **Best**, **BestN**, or **Bests**. The default selection criterion is the project default.
5. Specify how to calculate the forecasts. Forecasts can be based on the mean or median. By default the mean value is used.
6. Specify whether to restrict all weights to the 0 to 1 range. You can also edit the smoothing weights.

To view the SAS code for the model, click **View Code**.

Edit the Smoothing Weights

You can restrict the values and set the initial value for the following smoothing weights:

- combined level/trend weight parameter
- level weight parameter
- trend weight parameter
- damping weight parameter
- season weight parameter

Note: Only the parameters that are relevant for a model are available. The parameters that are not relevant are not available.

If you select the **Do not perform estimation** check box, then the model parameters are fixed values. To use this option, you must specify the initial values for all the model parameters that are available. By default, the model parameters are optimized.

To restrict the weights to stable values, select the **Restrict weights to stable values** check box.

To return to the default values, click **Reset**.

Create an Unobserved Components Model

To open the New Unobserved Component Model window, select the **Unobserved Components** option in the New Model dialog box.

From the New Unobserved Component Model window, you can create an unobserved component model. Specify a name and optionally a description for the new model. The Details box lists the system name for the model. You cannot edit this text.

Use the selection pane on the left to navigate among the following options:

- specification
- inputs

To view the SAS code for the model, click **View Code**.

For more information about the Unobserved Components model, see the HPFUCMSPEC procedure in the *SAS High-Performance Forecasting User's Guide*.

Specification Options for Unobserved Component Models

In the New Unobserved Component Model window, the following options are available under the Specification heading:

- transformation
- irregular component
- trend component
- seasonal component
- block seasonal components
- cycle components
- autoregressive component
- dependent lag

Transformations for the Unobserved Component Model

In the New Unobserved Component Model window, select **Transformation** under the Specification heading. From the **Dependent transformation** drop-down list, select the transformation for the model.

You can select from the following transformations:

- no transformation. This is the default.
- logarithmic transformation
- square-root transformation
- logistic transformation
- Box-Cox transformation where the parameter number is between -5 and 5. Specify this parameter value in the Box-Cox Transformation window. For more information, see "About the Box-Cox Transformation Window" on page 81.

Irregular Component for the Unobserved Component Model

In the New Unobserved Component Model window, select **Irregular Component** under the Specification heading. The irregular component corresponds to the overall random error in the model. To create an irregular component, complete the following steps:

1. Select the **Irregular component** check box.
2. Optionally, select the **Define initial variance** check box and specify a value for the initial variance. This value is used as the initial value during the parameter estimation process.
3. If you want the variance to be set at the initial value, select the **Do not perform estimation** check box. This option is not available unless you specify a value for the initial variance.

Trend Component for the Unobserved Component Model

In the New Unobserved Component Model window, select **Trend Component** under the Specification heading. The level component and the slope component combine to define the trend component for the model. If you specify both a level and slope component, then a locally linear trend is obtained. If you omit the slope component, then a local level is used.

To specify a component, complete the following steps:

1. Select the check box for the **Level** or **Slope** component.

   Optionally, select the **Define initial variance** check box and specify the initial value for the variance.
2. If you want the variance to be set at the initial value, select the **Do not perform estimation** check box. This option is not available unless you specify a value for the initial variance.

Seasonal Components for the Unobserved Component Model

In the New Unobserved Component Model window, select **Seasonal Components** under the Specification heading. The table lists each component and its specified values. For a model, you can specify a maximum of three seasonal components.

From this window, you can complete the following tasks.

- To add a new seasonal component, click **New**. The New or Edit Seasonal Component dialog box opens.
- To edit an existing seasonal component, select the component from the list and click **Edit**. The New or Edit Seasonal Component dialog box opens.
- To delete a component, select the component from the list and click **Delete**.

Create or Edit a Seasonal Component

To create or edit a seasonal component, complete the following steps:

1. Specify the type of seasonal component. A seasonal component can be one of two types: Dummy or Trigonometric (which is the default).
2. Specify the season length. The season length can be any integer larger than or equal to two. An example of a season length is 4, which corresponds to quarterly seasonality. The default value is determined by the variable that you assigned to the Time ID role.
3. Specify the initial value for the disturbance variance.
4. If you want the disturbance variance to be set at the initial value, select the **Do not perform estimation** check box. This option is not available unless you specify a value for the initial variance.

Block Seasonal Component for the Unobserved Component Model

In the New Unobserved Component Model window, select **Block Seasonal Components** under the Specification heading. The table lists each component and its specified values. For a model, you can specify a maximum of three block seasonal components.

From this window, you can complete the following tasks.

- To add a new block seasonal component, click **New**. The New or Edit Block Seasonal Component dialog box opens.
- To edit an existing block seasonal component, select the component from the list and click **Edit**. The New or Edit Block Seasonal Component dialog box opens.
- To delete a component, select the component from the list and click **Delete**.

Create or Edit a Block Seasonal Component

To create or edit a block seasonal component, complete the following steps:

1. Specify the type of block seasonal component. A seasonal component can be one of two types: Dummy (which is the default) or Trigonometric .
2. Specify the block size. The block size can be any integer larger than or equal to 2. An example of a block size is 24 for the number of hours in a day when a day is being used as the block.

   Note: The block seasonal component requires the values for the block size and the number of blocks.
3. Specify the number of blocks. The number of blocks can be any integer larger than or equal to 2. An example of the number of blocks is 7 for the number of days in a week.

   If you specify 24 as the block size and 7 as the number of blocks, then the block seasonal specification will have 7 blocks each of size 24.
4. If the first measurement is not at the start of a block, then select the **Define offset** check box. You can specify the first measurement within the block. This value must be between one and the block size.
5. To supply an initial value for the variance, select the **Define initial variance** check box. Specify the initial value for the disturbance variance.
6. If you want the disturbance variance to be set at the initial value, select the **Do not perform estimation** check box. This option is not available unless you specify a value for the initial variance.

Cycle Components for the Unobserved Component Model

In the New Unobserved Component Model window, select **Cycle Components** under the Specification heading. The table lists each component and its specified values.

You can specify up to 50 cycles in a model. By default, the cycle components are estimated from the data. You can optionally create additional cycle components.

From this window, you can complete the following tasks.

- To add a new cycle component, click **New**. The New or Edit Cycle Component dialog box opens.
- To edit an existing cycle component, select the component from the list and click **Edit**. The New or Edit Cycle Component dialog box opens.
- To delete a cycle component, select the component from the list and click **Delete**.

Create or Edit a Cycle Component

A default name is assigned to each cycle component. You cannot edit this name.

To specify a cycle component, complete the following steps:

1. To specify an initial cycle period to use during the parameter estimation process, select the **Define initial period** check box and then specify the initial value in the box. This value must be larger than 2.
2. To specify an initial damping factor to use during the parameter estimation process, select the **Define initial damping** check box and then specify the initial value in the box. You can specify any value between 0 and 1 (excluding 0 but including 1).
3. To specify an initial value for the disturbance variance parameter to use during the parameter estimation process, select the **Define initial variance** check box and then specify the initial value in the box. This value must be greater than or equal to 0.

For the period, damping, and variance values, you can select the **Do not perform estimation** check box. This option fixes the values of the component parameters to those specified. This option is not available unless you specify a value for the initial variance.

Autoregressive Component for the Unobserved Component Model

In the New Unobserved Component Model window, select **Autoregressive Component** under the Specification heading.

To specify an autoregressive component, complete the following steps:

1. Select the **Autoregressive component** check box.
2. Optionally, to specify an initial damping value to use during the parameter estimation process, select the **Define initial damping** check box, and then specify a value in the box. This value must be between -1 and 1, including -1 but excluding 1.

3. Optionally, to specify an initial variance to use during the parameter estimation process, select the **Define initial variance** check box, and then specify a value in the box. This value must be greater than or equal to 0.

For the damping and variance values, you can select the **Do not perform estimation** check box. This option fixes the values of the component parameters to those specified. This option is not available unless you specify a value for the initial damping or variance.

Dependent Lag for the Unobserved Component Model

To specify the forecast variable lags to be included as predictors in the model, select **Dependent Lag** under the Specification heading in the New Unobserved Component Model window.

To create a lag-dependent variable, complete the following steps:

1. Select the **Lag dependent variable** check box.
2. Specify the lag values. By default, the number of lags is 0.

   The syntax for specifying a lag value is the following:

   LAGS= *order*
   LAGS= (*lag, ..., lag***) ... (***lag, ..., lag***)**
   LAGS= (*lag, ..., lag***)<**s_1**> ... (***lag, ..., lag***)<**s_k**>**

   This is a required option in this statement. LAGS=(l_1, l_2, ..., l_k) defines a model with specified lags of the forecast variable included as predictors. LAGS= *order* is equivalent to LAGS=(1, 2, ..., *order*).

   A concatenation of parenthesized lists specifies a factored model. For example, LAGS=(1)(12) specifies that the lag values, 1, 12, and 13, corresponding to the following polynomial in the backward shift operator, be included in the model

   $$(1 - \phi_{1,1}B)(1 - \phi_{2,1}B^{12})$$

   Note that, in this case, the coefficient of the thirteenth lag is constrained to be the product of the coefficients of the first and twelfth lags.

   You can also specify a multiplier after a parenthesized list. For example, LAGS=(1)(1)12 is equivalent to LAGS=(1)(12), and LAGS=(1,2)4(1)12(1,2)24 is equivalent to LAGS=(4,8)(12)(24,48).

3. To specify the starting values for the coefficients of the lagged forecast variable, select the **Define initial phi values** check box and then specify the values in the text box.
4. To set the parameter values to the initial value, select the **Do not perform estimation** check box.

Inputs for the Unobserved Component Model

You can specify independent variables, predefined variables, and events as inputs to the model.

In the New Unobserved Component Model window, select the input type in the selection pane. Each variable or event is listed in the table. You can specify whether to use the variable or event or ignore it.

To specify the input series options for the independent variable or event, select the variable or event from the table and click **Options**. The Input Series Options dialog box opens. For more information, see "Specifying the Input Series Options" on page 93.

To create a new event, click **New Event** on the Events page. After the event is created, it is added to the list of events and used for the current model. It also appears in the list of events for the project, but it is not used at the project level.

Specifying the Input Series Options

In the ARIMA, Subset (factored) ARIMA, and Unobserved Components models, you can specify whether an input (independent variables or events) enters the model as a simple regressor or if it enters the model after some modifications, such as lagging or differencing. You can specify these modifications by using the Input Series Options dialog box.

From this dialog box, the following modification options are available:

- **Transformation** - specifies the transformation to be applied to the time series. You can select from the following transformations:
  - no transformation. This is the default.
  - logarithmic transformation
  - square-root transformation
  - logistic transformation
  - Box-Cox transformation where the parameter number is between -5 and 5. Specify this parameter value in the Box-Cox Transformation window. For more information, see "About the Box-Cox Transformation Window" on page 81.

  Note: The transformation is applicable only if the input variable is positive.
- **Specify lagging order** - specifies the delay, or lag, order for the input series. The syntax for this option is the following:
- **Specify differencing orders** - specifies the differencing orders for the input series. Use the following syntax when specifying a differencing order:

  DIF= *order*
  DIF= *(order1, order2, ...)*

- Specify whether you want an ordinary or dynamic regression.

  If you specify a dynamic regression for an ARIMA or subset ARIMA model, then you can specify the numerator (NUM) and denominator (DEN) polynomial of the transfer function. You cannot specify these options for an Unobserved Components model.

  The syntax to use for these polynomials is the following:

  NUM= *order*
  NUM= (*lag, ..., lag***) ... (***lag, ..., lag***)**
  NUM= (*lag, ..., lag***)<**s_1**> ... (***lag, ..., lag***)<**s_k**>**

  DEN= *order*
  DEN= (*lag, ..., lag***) ... (***lag, ..., lag***)**
  DEN= (*lag, ..., lag***)<**s_1**> ... (***lag, ..., lag***)<**s_k**>**

  For information about how to specify the polynomial order, see the autogressive (p) option in "Specification Options for Subset ARIMA Models" on page 83.

Once you have finished specifying these options, click **OK** to return the model dialog box.

Create an Intermittent Demand Model

To open the New Intermittent Demand Model window, select the **Intermittent Demand** option in the New Model dialog box.

In the New Intermittent Demand Model window, you can create the intermittent demand model that you choose. Specify a name and optionally a description for the new model. The Details box lists the system name for the model. You cannot edit this text.

Use the selection pane on the left to navigate among the following options:

- setup
- specification

To view the SAS code for the model, click **View Code**.

Setup for the Intermittent Demand Model

To specify which intermittent demand method to use, select **Setup** in the selection pane in the New Intermittent Demand Model dialog box. You can specify the following options:

- the model type. You can choose among the following types:
  - **Automatically select model type** - enables SAS Forecast Studio to determine the appropriate model to use based on the selection criterion that you select from the drop-down list. The model selection is determined by how well the method fits (in sample) or predicts (holdout sample) the demand series component by treating the demand index as a time index.
  - **Croston's method** - models and forecasts each component independently and then combines the two forecasts.
  - **Average Demand model** - forecasts the intermittent time series by treating the average demand series as a time series based on the demand index. Optimal smoothing parameters can be estimated and predictions for average demand can be computed using nonseasonal exponential smoothing methods as well as their transformed versions.
- the base demand. The demand for the intermittent time series typically depends on a base value. You can have SAS Forecast Studio automatically determine the base demand based on the characteristics of the time series, or set the base demand yourself. The default base value is 0, but it can be any constant value.

Specification Options for an Intermittent Demand Model

In the New Intermittent Demand Model window, you can create an average demand model, a demand intervals model, and a demand sizes model. The demand intervals and demand sizes models use Croston's method.

For the model that you want to create, in the selection pane select **Specification** under the model heading to access these options.

The following specification options for the intermittent demand models are available:

- the forecasting method to be used to forecast the time series. The following forecasting models are available:
  - **Simple** - Simple (Single) Exponential Smoothing. This is the default.
    The model equation for simple exponential smoothing is as follows:

    $$Y_t = \mu_t + \epsilon_t$$

  - **Double** - Double (Brown) Exponential Smoothing
    The model equation for double exponential smoothing is as follows:

    $$Y_t = \mu_t + \beta_t t + \epsilon_t$$

  - **Linear** - Linear (Holt) Exponential Smoothing
    The model equation for linear exponential smoothing is as follows:

    $$Y_t = \mu_t + \beta_t t + \epsilon_t$$

  - **DampTrend** - Damped Trend Exponential Smoothing
    The model equation for damped-trend linear exponential smoothing is as follows:

$$Y_t = \mu_t + \beta_t t + \epsilon_t$$

  - **BestN** - Best Candidate Nonseasonal Smoothing Model

  For more information about these models, see the "Forecast Process" chapter in the *SAS High-Performance Forecasting User's Guide*.

- the dependent transformation to be applied to the time series. The following transformations options are available:
  - no transformation. This is the default.
  - logarithmic transformation
  - square-root transformation
  - logistic transformation
  - Box-Cox transformation where the parameter number is between -5 and 5. Specify this parameter value in the Box-Cox Transformation window. For more information, see "About the Box-Cox Transformation Window" on page 81.
  - an automatically chosen transformation based on the model selection criteria.
- the selection criterion to use to select from several candidate models. This option is available only when the method is **BestN** (Best Candidate Nonseasonal Smoothing Model). The default selection criterion is the project default.
- how to calculate the forecasts. Forecasts can be based on the mean or median. By default the mean value is used.
- whether to restrict all weights to the 0 to 1 range, do not restrict the weights, or specify the weights. You can also edit the smoothing weights. For more information, see "Edit the Smoothing Weights" on page 87.

Create a Curve Fitting Model

To open the Curve Fitting Model window, select the **Curve Fitting** option in the New Model dialog box.

In the Curve Fitting Model window, you can create a curve fitting model with a linear, quadratic, or exponential trend.

To create a curve fitting model, complete the following steps:

1. Specify a name and optionally a description for the model. The Details box lists the system name for the model. You cannot edit this text.
2. Select the **Log transform dependent variable** check box if you want to log transform the dependent variable.
3. Select the curve fitting method. You can select from the following options:

- **Linear trend** - creates a model with a linear time trend, with $X_t = t - c$.
- **Quadratic trend** - creates a model with a quadratic time trend, with $X_t = (t - c)^2$.

4. Select the **Log transform curve component** check box if you want to log transform the curve component.
5. Click **OK** to save the model.

To view the SAS code for the model, click **View Code**.

Create a Moving Average Model

To open the New Moving Average Model window, select the **Moving Average** option in the New Model dialog box.

In the New Moving Average Model window, you can create a moving average model that you choose. The formula for the Moving Average model with width *k* is the following:

$$Y_t = [y_{(t-1)} + ... + y_{(t-k)}]/k + error$$

In ARIMA notation, this model is ARIMA(k, 0, 0) with no intercept and with the autogressive parameters (AR) fixed: $AR = 1/k, 1/k, ..., 1/k$.

To create a moving average model, complete the following steps:

1. Specify a name and optionally a description for the new model. The Details box lists the system name for the model. You cannot edit this text.
2. Select the **Log transform dependent variable** check box to log transform the dependent variable.
3. In the **Window (periods)** box, specify the number of periods for the moving average. The default value is 3.

To view the SAS code for the model, click **View Code**.

Create a Multiple Regression Model

To open the New Multiple Regression Model window, select the **Multiple Regression** option in the New Model dialog box.

You use the New Multiple Regression Model window to create a multiple regression model. Specify a name and optionally a description for the new model. The Details box lists the system name for the model. You cannot edit this text.

Use the selection pane on the left to navigate among the following options:

- independent variables
- events

To view the SAS code for the model, click **View Code**.

Independent Variables for the Multiple Regression Model

In the New Multiple Regression Model window, select **Independent Variables**.

To specify the independent variable options, complete the following steps:

1. Select the **Log transform dependent variable** check box to log transform the dependent variable.
2. Select the **Log transform independent variables** check box to log transform the independent variables.
3. Use the drop-down list in the Status column to specify which independent variables should be included in the model. By default, all of the independent variables are ignored.

Events for the Multiple Regression Model

In the New Multiple Regression Model window, select **Events** to specify the events to use as inputs to the model.

Use the drop-down list in the Status column to specify which events should be included in the model. By default, all of the events are ignored.

To create a new event, click **New Event** on the Events page. After the event is created, it is added to the list of events and used for the current model. It also appears in the list of events for the project, but it is not used at the project level.

Create a Random Walk Model

To open the Random Walk Model window, select the **Random Walk** option in the New Model dialog box.

In the Random Walk Model window, you can create a random walk model. If you use the default settings, then you can create an ARIMA(0, 1, 0) model with no intercept. The formula for this model is $y_t = y_{(t-1)} + error$.

You can also create the following Random Walk models:

Random Walk with Drift

$y_t = const + y_{(t-1)} + error$,
or in ARIMA notation ARIMA(0, 1, 0)

Seasonal Random Walk without Drift

ARIMA(0, 1, 0)(0, 1, 0)s with no intercept

Seasonal Random Walk with Drift

ARIMA(0, 1, 0)(0, 1, 0)s

To create a random walk model, complete the following steps:

1. Specify a name and optionally a description for the new model. The Details box lists the system name for the model. You cannot edit this text.
2. Select the **Log transform** check box to log transform the dependent variable. When you select this option, the model uses the data for the log transformed series.
3. Select the **Trend** check box to include a trend term in the model.
4. Select the **Drift** check box to include a drift term in the model.
5. Select the **Seasonal** check box to include a seasonal term in the model.

To view the SAS code for the model, click **View Code**.

View a Model

You can view the forecast plots and tables for a model by using the Model Analysis View. For more information, see "About the Model Analysis View" on page 25.

Edit a Model

You can edit a model from the model selection list, the model repository, Data Table in the Forecasting View, or the Model Analysis View.

Note: When you edit a model that is used by multiple series, then you need to refit the model to all of the series that use that model.

The following is an example in which such a refit is necessary: You have an ARIMA model named MYMODEL for one of the series in the project. While editing an exponential smoothing model for another series, you rename that model MYMODEL. You are asked if you want to overwrite the model with the same name, and you click **OK**.

Now the MYMODEL model in the first series is an exponential smoothing model instead of an ARIMA model, but the MAPE value is still for the ARIMA model. To see the updated MAPE value, you have to refit the model. Because other series could be using the same model, you should refit the model to all of the series in the project. To do so select **Project > Reforecast**. The Reforecast Project dialog box opens. Select **Re-estimate parameters for the current forecast model** and click **OK**. All of the models in the project are refitted.

To edit a model from the model selection list, complete the following steps:

1. Select the series in the **Hierarchy** tab or **Table** tab and select **Series → Model List**. The Model List dialog box opens.
2. Select the model that you want to edit from the list and click **Edit**. The model specification dialog box for that model opens.

To edit a model from the model repository, complete the following steps:

1. Select the series in the **Hierarchy** tab or **Table** tab and select **Project** → **Model Repository**. The Model Repository dialog box opens.
2. Select the model that you want to edit from the list and click **Edit**. The model specification dialog box for that model opens.

To edit a model from the Data table in the Forecasting View, click **Statistical Forecast** in the table. The model specification dialog box for that model opens.

To edit a model from the Model Analysis View, complete the following steps:

1. Select **View** → **Model Analysis View**. The Model Analysis View opens.
2. Open the model that you want to edit and click **Edit**. The model specification dialog box for that model opens.

For more information about each of the model specification dialog boxes, see the following topics:

- "Create an ARIMA Model" on page 80
- "Create a Curve Fitting Model" on page 96
- "Create an Exponential Smoothing Model" on page 85
- "Create an Intermittent Demand Model" on page 94
- "Create a Moving Average Model" on page 97
- "Create a Multiple Regression Model" on page 97
- "Create a Random Walk Model" on page 98
- "Create a Subset (Factored) ARIMA Model" on page 82
- "Create an Unobserved Components Model" on page 87

Compare Models

If you have created several models for the current series, you might want to compare the statistics of fit for these models. Through this comparison, you can determine which model that you want to use.

To compare the statistics for the models, select **Series** → **Compare Models** or click **Compare** in the Model List dialog box. The Compare Models dialog box opens.

By default, each model is a row in the table and the statistics of fit are in columns. For each statistic of fit, the cell for the best model is shaded in green.

From this dialog box, you can view the following statistics for each of the models:

- **AADJRSQ** - Amemiya's Adjusted R-Square
- **ADJRSQ** - Adjusted R-Square
- **AIC** - Akaike Information Criterion
- **APC** - Amemiya's Prediction Criterion
- **DFE** - Degrees of Freedom for Error
- **GMAPE** - Geometric Mean Percent Error
- **GMAPPE** - Geometric Mean Predictive Percent Error
- **GMASPE** - Geometric Mean Symmetric Percent Error
- **GMRAE** - Geometric Mean Relative Absolute Error
- **MAE** - Mean Absolute Error
- **MAPE** - Mean Absolute Percent Error
- **MAPPE** - Symmetric Mean Absolute Predictive Percent Error
- **MAXERR** - Maximum Error
- **MAXPE** - Maximum Percent Error
- **MAXPPE** - Maximum Predictive Percent Error
- **MAXRE** - Maximum Relative Error
- **MAXSPE** - Maximum Symmetric Percent Error
- **MDAPE** - Median Percent Error
- **MDAPPE** - Median Predictive Percent Error
- **MDASPE** - Median Symmetric Percent Error
- **MDRAE** - Median Relative Absolute Error
- **ME** - Mean Error
- **MINERR** - Minimum Error
- **MINPE** - Minimum Percent Error
- **MINPPE** - Minimum Predictive Percent Error
- **MINRE** - Minimum Relative Error
- **MINSPE** - Minimum Symmetric Percent Error
- **MPE** - Mean Percent Error
- **MPPE** - Mean Predictive Percent Error
- **MRAE** - Mean Relative Absolute Error
- **MRE** - Mean Relative Error
- **MSE** - Mean Square Error
- **MSPE** - Mean Symmetric Percent Error
- **N** - sample size

- **NMISSA** - Number of Missing Actual Values
- **NMISSP** - Number of Missing Predicted Values
- **NOBS** - Number of Observations
- **NPARMS** - Number of Parameters
- **RMSE** - Root Mean Square Error
- **RSQUARE** - R-Squared **Note**: You can have negative R-square values with time series models. A negative R-square value means that your current model is worse than a model of fitting the mean.
- **RWRSQ** - Random Walk R-Square
- **SBC** - Schwarz Bayesian Information Criterion
- **SMAPE** - Symmetric Mean Absolute Percent Error
- **SSE** - Sum of Square Error
- **SST** - Total Sum of Squares
- **TSS** - Total Corrected Sum of Squares for the dependent variable
- **UMSE** - Unbiased Mean Square Error
- **URMSE** - Unbiased Root Mean Square Error

Select the Model to Forecast

You can select the forecast model from the model selection list or from the Model Analysis View.

To select the forecast model from the model selection list, complete the following steps:

1. Select the series in the Hierarchy or Table tab and select **Series** → **Model List**. The Model List dialog box opens.
2. In the Forecast column, select the forecast model for the series.
3. Click **OK** to reconcile the project.

To select the forecast model from the Model Analysis View, complete the following steps:

1. Select **View** → **Model Analysis View**. The Model Analysis View opens.
2. Open the model that you want to use for the forecast and click **Select as Forecast**.
3. Click **Close** to return to the Forecasting View.

Using the Model Selection List

What Is the Model Selection List?

The model selection list shows all of the models that have been defined for the current project. This is different from the model repository that shows all of the user-defined and system-defined models.

To view the model selection list, select the series in the Hierarchy or Table tab and select **Series** → **Model List**. The Model List dialog box opens.

Note: You can also open this dialog box by clicking **Model List** in the Forecasting View or by clicking  in the toolbar.

Working with the Model Selection List

In the Model List dialog box, you manage the model selection list. If a model in the list failed to fit, then the value in the SOF (Statistic of Fit) column is set to missing. In the Forecast column, select the model that will be the forecast model for the series. If you select a new model, the project is reconciled when you click **OK**.

To view more information about a model, select the model in the list. In the Model Information box, you can see the time plot for the model, the model name, the description, and the details (or system name) for the model.

In this window, you can complete the following tasks:

- To create a new model, click **New**. The New Model dialog box opens. For more information, see "Create a New Model" on page 79.
- To view an existing model, select the model from the list and click **View**. The Model Analysis View opens. For more information, see "View a Model" on page 99.
- To edit an existing model, select the model from the list and click **Edit**. The model specification dialog box for that model opens. For more information, see "Edit a Model" on page 99.
- To copy a model, select the model name from the table and click **Copy**. The new model is added to the list with the name *model-name*Copy*n* where *model-name* is the original model name and *n* is the lowest integer value available.
- To compare the models, click **Compare**. The Compare Models dialog box opens. For more information, see "Compare Models" on page 100.
- To delete model from the list, select the model name and click **Remove**.
- To view the code for a model, select the model from the list and click **View Code**. The SAS Code window opens. For more information, see "View the SAS Code" on page 47.

Using the Model Repository

What Is the Model Repository?

The model repository shows all of the user-defined and system-defined models. This is different from the model selection list that shows only the models for the current series.

To view all of the models in the model repository, select **Project** → **Model Repository**. The Model Repository dialog box opens.

Working with the Model Repository

In the Model Repository dialog box, you manage the models in the model repository.

To view more information about a model, select the model in the list. In the Model Information box, you can see the model name, the description, and the details (or system label) for the model.

From this dialog box, you can complete the following tasks:

- To create a new model, click **New**. The New Model dialog box opens. For more information, see "Create a New Model" on page 79.
- To edit a model, select the model from the table and click **Edit**. The model specification dialog box for that model opens. For more information, see "Edit a Model" on page 99.
- To copy a model, select the model from the table and click **Copy**. The copied model is named *model-name*Copy*n* where *model-name* is the name of the original model and *n* is the lowest available integer value.
- To delete a model, select the model from the table and click **Delete**.
- To import a model, click **Import**.

Chapter 9
Generating Reports

Chapter Contents

Chapter 9
Generating Reports

Overview of Reports

You can use reports to share forecasting results with other people at your site. The reports in SAS Forecast Studio are created from stored processes. A stored process is a SAS program that is designed to be stored on a central server and accessed by numerous remote clients. Stored processes consist of two distinct parts: the SAS code and the stored process definition that resides on a metadata server.

Including parameters in your code increases your ability to reuse code and also enables the code to be customized at run time. Stored process parameters typically correspond to macro variables that are included in the SAS code. Macro variables are usually preceded by an ampersand (&).

SAS Forecast Studio ships with some default stored processes. These stored processes are saved in the Stored Process Manager in SAS Management Console. You can run a stored process to view a report or create your own report by writing a stored process.

View a Report

Reports are created by running a SAS stored process. By default, SAS Forecast Studio ships with some sample reports. You might want to create your own reports. For more information, see "Create a Report" on page 108.

To view a report for the current forecast, complete the following steps:

1. Select **Tools** → **Reports and Stored Processes**. The Reports and Stored Processes window opens.
2. From the tree, select the stored process for the report that you want. When you select a stored process from the list, the Information box displays the name of the stored process, a brief description (if one was provided), and the date that the stored process was created.
3. Click **Run**. If your stored process code includes parameters that need to be customized at run time, the Enter Parameters dialog box opens. If no parameters are included, then this dialog box will not appear. The report opens in a new Web browser window.

Create a Report

The reports for SAS Forecast Studio are created using stored processes. These stored processes are written in SAS code and must be registered in SAS Management Console.

For more information, see the "Additional Administration Tasks" in the *SAS Forecast Studio Administrator's Guide*.

Glossary

aggregation statistic
: the mathematical operation used to combine forecasts across levels. The levels where the aggregation statistic is used depend on the reconciliation method. See also reconciliation method.

autocorrelation
: the internal correlation between observations in a time series. Autocorrelation coefficient values range from -1 to +1. When autocorrelation coefficient values are positive, deviations from the mean of one sign tend to be followed by deviations of the same sign. When they are negative, deviations of one sign tend to be followed by deviations of the opposite sign.

autocorrelation function (ACF) plot
: a plot of the autocorrelation coefficients across different values of time lags. This plot enables you to determine whether seasonality exists in the time series.

bottom-up method
: a type of reconciliation method that uses the data at the lowest level of the hierarchy to generate forecasts. These forecasts are then used to generate the forecasts for the higher levels in the hierarchy. See also middle-out method, reconciliation method, top-down method.

confidence limits
: the upper and lower values of a confidence interval. There is a percentage of confidence (typically 95%) that the true value of the parameter being estimated lies within the interval.

controllable independent variable
: a causal factor where the future values are under the control of the organization producing the forecasts.

deterministic independent variable
: a causal factor where the future values are known with certainty.

dummy variable
: a numeric variable whose value is either 1 or 0. In SAS Forecast Studio, dummy variables are used to indicate whether or not unusual events occur. The variable takes on the value of 1 during the event and 0 otherwise.

event
: an incident that disrupts the normal flow of a process that generates the time series. Examples of events are holidays, retail promotions, and natural disasters.

event repository
: a storage location that contains information about calendar events and includes a brief description of each event.

fit criterion
: a statistical value that is used to evaluate how well a forecasting model performs by comparing the actual data to the predictions.

forecast
: a numerical prediction of a future value for a time series.

holdout sample
: the number of periods of the most recent data that should be excluded from the parameter estimation.

horizon
: the number of periods into the future for which predictions are made.

inverse autocorrelation
: the autocorrelation of an autoregressive model remodeled as a moving average model.

inverse autocorrelation function (IACF) plot
: a plot of the inverse autocorrelation coefficients across different values of time lags. This plot is useful for detecting over-differencing in the model.

level shift
: a persistent change in the magnitude of a time series curve.

middle-out method
: a type of reconciliation method that adds the data from the lower levels and then uses these values to generate the forecasts for the middle level. Because some hierarchies have more than one middle level, you need to specify the level that you want to use. The forecasts at the middle level are used to generate forecasts for both the higher and lower levels. See also bottom-up method, reconciliation method, top-down method.

model selection criterion
: the statistical value that is used for forecast model selection.

model selection list
: a list of candidate model specifications. You can choose which model specification is best suited to forecast a particular time series.

partial autocorrelation
: the internal correlation between observations in a time series that has the effect of all intervening lags removed.

partial autocorrelation function (PACF) plot
: a plot of the partial autocorrelation coefficients across different values of time lags. This plot is useful for identifying the order of an autoregressive model.

project hierarchy
: the order of the variables that you have assigned to the BY variables role. An example of a hierarchy is Region > Product Line > Product Name.

pulse
: a temporary change in the magnitude of a time series curve. The magnitude returns to the former level immediately after the change.

ramp
: a persistent change in the trend or slope of a time series curve.

reconciliation method
: the level in the hierarchy that contains the data that you want to forecast. The

following reconciliation methods are available: bottom-up method, middle-out method, and top-down method.

residual
: the difference between an observed data value and its predicted value.

seasonal adjustment
: the process of seasonality from time series data.

seasonality
: a regular change in time series data values that occurs at the same point in each time cycle.

statistic of fit
: See fit criterion.

stochastic independent variable
: a causal factor where the future values are not known with certainty.

time series data
: time-stamped data collected over time at a particular frequency. Some examples of time series data are Web visits per hour, sales per month, and calls per day.

top-down method
: a type of reconciliation method that adds up the data from the lowest levels in the forecast and then uses this value to generate the forecasts. The forecast at the highest level is used to generate the forecasts for lower levels in the hierarchy. See also bottom-up method, middle-out method, reconciliation method.

transactional data
: time-stamped data collected over time at no particular frequency. Some examples of transactional data are point-of-sale data, inventory data, call center data, and trading data.

white noise
: a series of random fluctuations in the values of a data series. A white noise model has only a mean or constant parameter. A series is described as white noise if its spectral density function is constant.

Index

R

S

T

U

V

W

Your Turn

If you have comments or suggestions about *SAS® Forecast Studio 1.2: User's Guide,* please send them to us on a photocopy of this page or send us electronic mail.

For comments about this book, please return the photocopy to

SAS Publishing
SAS Campus Drive
Cary, NC 27513
E-mail: **yourturn@sas.com**

For suggestions about the software, please return the photocopy to

SAS Institute Inc.
Technical Support Division
SAS Campus Drive
Cary, NC 27513
E-mail: **suggest@sas.com**