SAS® Forecast Studio 2.1

User's Guide

The correct bibliographic citation for this manual is as follows: SAS Institute Inc. 2007. *SAS® Forecast Studio 2.1: User's Guide*. Cary, NC: SAS Institute Inc.

SAS® Forecast Studio 2.1: User's Guide

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

1st printing, October 2007

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

Part I

About SAS Forecast Studio

Chapter 1

What's New in SAS Forecast Studio 2.1

Contents

Overview

SAS Forecast Studio 2.1 has the following changes and enhancements:

- performance and scalability improvements
- project enhancements
- filtering enhancements
- new forecasting settings
- model enhancements
- general user-interface enhancements
- documentation enhancements

Performance and Scalability Improvements

The performance of project operations is improved in SAS Forecast Server 2.1. Project operations include the following:

- opening and closing large projects
- filtering and scrolling through large projects
- fitting models for large projects

Also, scalability for series operations, such as the following operations, has been improved.

- adding, editing, copying, and deleting models
- making a series active or inactive
- re-creating automatically-generated models, refitting all models, and selecting a forecast model for a series
- refreshing the current forecast model, updating the parameter values for a series, and refitting the current forecast

Project Enhancements

The following enhancements were made for projects:

- Several new SAS macro programs that can be run in batch mode enable your site administrator to perform the following tasks:
  - display all projects that are defined in the metadata
  - create a project
  - diagnose, select, fit, update, and forecast a project
  - copy a project
  - delete a project
  - rename a project
  - move a project to another location
  - export a project to an archive file
  - import a project from an archive file
  - register an existing project directory in the SAS metadata
  - remove a project from the SAS metadata
  - assign ownership of a project to a user
  - specify the default location for projects on a SAS workspace server
  - migrate a SAS Forecast Studio 1.4 project to SAS Forecast Studio 2.1
  - save a project to a new filename

For more information about these macros and how to use them, see the *SAS Forecast Server Administrator's Guide*.

- The following statistics of fit are now available from the output data sets that SAS Forecast Studio creates for a project:
  - mean absolute scaled error
  - minimum absolute error percent of standard deviation
  - maximum absolute error percent of standard deviation
  - mean absolute error percent of standard deviation
  - median absolute error percent of standard deviation
  - geometric mean absolute error percent of standard deviation
- Within SAS Forecast Studio, you can now do the following:
  - save a project. This replaces the functionality for archiving a project.
  - view the input data source when creating a project.

Filtering Enhancements

The following enhancements have been made to the filtering functionality:

- Filter results now appear in a filtered table in the navigation panel. Filters are no longer used to identify exceptions, so series that meet the filter criterion do not appear in red in the hierarchy.
- When creating a filter, you can select whether to apply the filter to some or all the levels of the hierarchy. You can also specify whether the series in the filter results must have overrides and notes.
- You can now create filters from model and series properties as well as statistics of fit. The new filter categories are the following:
  - reconciled statistics of fit
  - model statistics of fit
  - model properties
  - series properties
  - BY variables
- The following logical conditions are new:
  - greater than or equal to
  - less than or equal to
- The following statistics of fit are now available:

  - mean absolute scaled error
  - minimum absolute error percent of standard deviation
  - maximum absolute error percent of standard deviation
  - mean absolute error percent of standard deviation
  - median absolute error percent of standard deviation
  - geometric mean absolute error percent of standard deviation

New Forecasting Settings

The options in the Forecasting Settings dialog box have been reorganized to make it easier to find an option.

The following options are new or enhanced:

- In the Diagnostics panel, more outlier detection options are available for ARIMA models. These new options enable you to specify the significance level for an outlier and the maximum percentage of a series that can be outliers.
- In the Models panel, you can specify whether to fit only generated exponential smoothing models at the lowest levels of the hierarchy. Selecting this option improves the time for diagnosis and forecasting.
- In the Models panel, you no longer must select a generated ARIMA, exponential smoothing, or unobserved components model. Instead, you can use models from an external list.

Model Enhancements

The following enhancements were made to the modeling functionality:

- By default, no functional transformation is applied to the series. In previous releases, SAS Forecast Studio tried to automatically choose between no transformation and a logarithmic transformation.
- You can no longer add a model to all series in a project. However, you can add a model to an individual series by using the functionality in the Model View.
- In the Model View, you can open a report that describes the model specification. This read-only view enables you to see the options that were selected in the model specification.

General User-Interface Enhancements

The following enhancements have been made to the user interface:

- In the new navigation panel, you can switch between views of the hierarchy, filtered table, and system exceptions. The new system exceptions table identifies the series for which a forecast was not produced or SAS Forecast Studio could not fit a model from the model selection list.
- For reconciled projects, the filtered table will contain a column for the selection criterion for the reconciled forecasts. This new column is in addition to the column for the selection criterion for the statistical forecasts.
- At the top of the Forecasting View, the Model View, and the Series Analysis View, you can see the name of the series that you selected, the dependent variable for the project, and the statistics of fit for the series. If additional statistics of fit are available, you can view this information from the new Statistics of Fit dialog box.
- The Forecast Summary dialog box now contains analysis plots and tables for each level of the hierarchy. You can click these plots and tables for more detailed information.

Documentation Enhancements

The following changes were made to the *SAS Forecast Studio User's Guide*:

- The new "Working with SAS Forecast Studio Tasks in the SAS Add-In for Microsoft Office" chapter describes each of the available tasks and explains how you can use these tasks to work with SAS Forecast Studio projects.
- The new "Working in Batch Mode" chapter contains content that was previously in the "Working with Projects" chapter.

The following changes were made to the SAS Forecast Studio Help:

- To make it easier for you to find the content that you need while working in the product, the following chapters and appendixes were moved to the *SAS Forecast Studio User's Guide*:
  - "Preparing an Input Data Set for SAS Forecast Studio"
  - "Reserved Variable Names in SAS Forecast Studio"
  - "Output Data Sets in the Project Directory"
  - "Statistics of Fit"
  - "Default Model Selection Lists"
  - "Sample Reports in SAS Forecast Studio"

Chapter 2

Introduction to SAS Forecast Studio

Contents

Overview of SAS Forecast Studio

SAS Forecast Server has three components:

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

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

SAS Forecast Server is a large-scale automatic forecasting solution that enables organizations to produce huge quantities of high-quality forecasts quickly and automatically. To generate forecasts without programming, you can use SAS Forecast Studio, which is the client component of the SAS Forecast Server. 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*.

SAS Forecast Studio is a forecasting application that is 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.

Using this application, you can do the following tasks:

- Generate forecasts automatically by using models that ship with SAS Forecast Studio. You can generate forecasts by using a model selection list that you have 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 and specify how conflicts should be resolved.
- Export projects as SAS code for processing in a batch environment.
- Generate reports to share forecasting results with other people at your site.

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

Accessibility and Compatibility Features

SAS Forecast Studio 2.1 includes accessibility and compatibility features that improve the usability of the product for users with disabilities, with exceptions noted below. These features are related to accessibility standards for electronic information technology that were adopted by the U.S. Government under Section 508 of the U.S. Rehabilitation Act of 1973. If you have specific questions about the accessibility of SAS Forecast Studio, send them to accessibility@sas.com or call SAS Technical Support.

All known exceptions to accessibility standards are documented in the following table. SAS is committed to improving the accessibility and usability of our products. SAS currently plans to address these issues in a future release of the software.

Table 2.1 Accessibility Exceptions

Section 508 Accessibility Criteria	Support Status	Explanation
(c) A well-defined on-screen indication of the current focus shall be provided that moves along interactive interface elements as the input focus changes. The focus shall be programmatically exposed so that assistive technology can track the focus and any focus changes.	Supported with minor exceptions	In the New Project Wizard, the initial focus is not set correctly in the following steps: • In Step 3, the initial focus is set to the **Next** button. • In Step 6, the initial focus is set to the **Finish** button. Because the initial focus is set on a button, the previous text in the step is not read. SAS currently plans to address this and other accessibility issues in the New Project Wizard in a future release.
(d) Sufficient information about a user-interface element (including the identity, operation, and state of the element) shall be available to the assistive technology. When an image represents a program element, the information conveyed by the image must also be available in text.	Supported with exceptions	The following components could not be read by a JAWS screen reader: • Data in tables might not be read properly. This is a known issue in Java. • In the New Project Wizard and Advanced Reconciliation dialog box, the text for some of the options cannot be read. • In the Filter Properties dialog box and the SAS log, the screen reader cannot read the bitmap images that are used for some buttons. SAS currently plans to address these and additional capability issues with screen readers in a future release.

Table 2.1 *continued*

Section 508 Accessibility Criteria	Support Status	Explanation
(g) Applications shall not override user-selected contrast and color selections and other individual display attributes.	Supported with exceptions	When your operating system is set to high contrast, the following icons and controls might be difficult to discern from the background: • In the hierarchy, the expand (+) and collapse (-) icons for each branch of the hierarchy appear black on black in high contrast, so it is difficult to tell if a branch is fully expanded or collapsed. • In the New Project Wizard and the Variable Assignments dialog box, the arrow buttons that you use to assign variables to roles appear black on black in high contrast. • In the Model View and Series Analysis View, the minimize, maximize, and close buttons in the plot and table windows appear black on black in high contrast. SAS currently plans to address these and additional color and contrast issues in a future release. Until that time, low vision users who require high contrast might find the use of a screen magnifier with reverse video setting to be a sufficient accomodation.

Table 2.1 *continued*

Section 508 Accessibility Criteria	Support Status	Explanation
(l) When electronic forms are used, the form shall enable users of assistive technology to access the information, field elements, and functionality required for completion and submission of the form, including all directions and cues.	Supported with exceptions	To work with assistive technologies, software generally needs to support a visible focus indicator and work with a screen reader. SAS Forecast Studio has exceptions noted in both criteria (c) and (d) that limit its compliance with assistive technologies.

Using SAS Forecast Studio Help

About Help

To open Help, select **Help** → **SAS Forecast Studio Help** or click  in the toolbar. You can also view the help for a specific window by clicking the **Help** button. 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 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, select the topic in the table of contents, search, or index tab, and click  in the toolbar.

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

Chapter 3

Preparing an Input Data Set for SAS Forecast Studio

Contents

Understanding Time Series Data

Requirements for Time Series Data

To generate forecasts in SAS Forecast Studio, you need time series data. You might already have this time series data, or you might have transactional data. If you have transactional data, you can use the accumulation options in SAS Forecast Studio to convert the transactional data into a time series. For more information, see "Data Preparation Options" on page 87.

If you have time series data that you want to use in SAS Forecast Studio, the time series data must meet the following requirements:

- The data set contains one variable for each dependent variable.
- 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, or monthly.

Process for Creating Time Series Data

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" on page 16.

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

1. The data is accumulated to the appropriate time interval 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 "Working with Missing Values" on page 18.
3. The data is sorted 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.

Selecting the Input Data Set

Required Format

SAS Forecast Studio requires a date or datetime variable in the data set to generate forecasts. SAS Forecast Studio generates forecasts from timestamped 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:

- The time ID variable contains the date or datetime value of each observation.
- BY variables enable you to group observations into a hierarchy.
- Dependent variables are the variables to model and forecast.
- Independent variables are the explanatory, input, or indicator variables that are used to model and forecast the dependent variable.

- Reporting variables are not used for analysis but for reporting only.
- Indicator variables are used to signify any unusual events in the model, such as holidays and promotions. You can add an indicator variable to a SAS Forecast Studio project by assigning the variable to the Independent variables role or by creating an event.

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 Appendix A, "Reserved Variable Names in SAS Forecast Studio."

For more information about variable roles, see "Assign Data to Roles" on page 44.

Examples

Here are two examples of input data sets 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.

Working with Missing Values

If your data contains missing values in variables other than the time ID variable (such as 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 “Set Forecasting Options” on page 47.

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*.

Part II

Working in SAS Forecast Studio

Chapter 4

Using the Workspace

Contents

Overview of the Workspace

The SAS Forecast Studio workspace consists of the following components:

Toolbar
: displays some of the most commonly used SAS Forecast Studio options, so that you can quickly and easily manage your project.

Navigation Panel

enables you to choose the series that you want to appear in the view area. You can use the drop-down list at the top of the navigation panel to switch between the following panels:

- Hierarchy - displays the project data as a tree when you forecast the data hierarchically. You can expand and collapse the branches of the tree.
- Filtered table - displays the series that meet a selected filter criterion. You can select from predefined or user-defined filters.
- System exceptions - displays any system exceptions from SAS Forecast Studio when generating the forecasts.

View Area

displays the forecasting results for the selected series in the navigation panel. You can use the drop-down list at the top of the view area to switch between the following views:

- Forecasting View - displays the forecast plot and the data table for the selected series.
- Model View - enables you to view plots and tables 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 Analysis View - enables you to view various plots and tables for the selected series. The plots in the Series Analysis View use the historical data of the time series.

Series Properties pane

displays the statistical properties of the currently selected series.

Status bar

displays information about the current procedure, the user name, and the SAS Forecast Server Middle Tier 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 4.1 Buttons in the Toolbar

Button	Description
	enables you to create a new project
	enables you to manage existing projects
	reconciles the hierarchy

Table 4.1 (continued)

Button	Description
	enables you to reforecast the data for all series
	enables you to view a list of available reports
	opens SAS Forecast Studio Help

About the Hierarchy Panel

The Hierarchy panel displays a tree view of the hierarchy. This hierarchy is defined by the variables that you assign to the BY variables role when you create the project. You can expand or collapse the branches in the tree.

NOTE: If the value for a BY variable is blank, then that value appears in the hierarchy as _ _. If you have several blank values in your data, then you could have multiple _ _ nodes in a hierarchy.

NOTE: You can modify the configuration of the hierarchy from the Hierarchy Settings dialog box. For more information, see "Modify the Hierarchy Settings" on page 53.

The following icons might appear in the hierarchy, depending on the changes that you have made to the project:

- If you add an override for a series, the  icon appears next to the series in the hierarchy. If you delete the override, then the override icon is removed from the hierarchy.
- If you add a note for a series, the  icon appears next to the series in the hierarchy.

Figure 4.1 Hierarchy panel

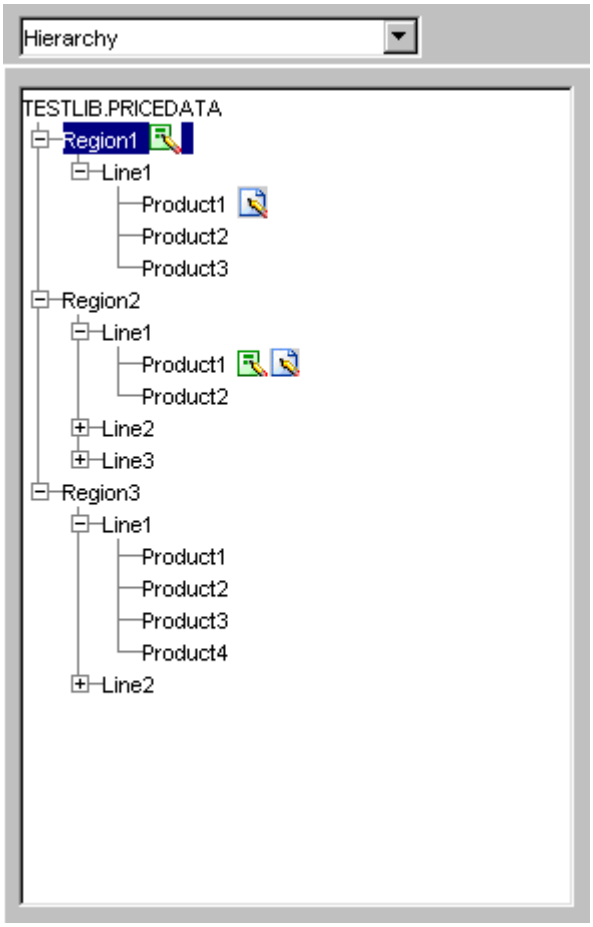

About the Filtered Table Panel

What Is the Filtered Table?

By default, the filtered table panel contains the following information:

- the series in the project that meet the criterion of the selected filter. You can filter the contents of the table by using the **Filter** drop-down list. SAS Forecast Studio automatically creates filters for all of the series and for each level of the hierarchy. However, you can create additional filters. These user-defined filters are added to the **Filter** drop-down list. After selecting a filter, you can refresh the contents in the table by clicking **Refresh**.

  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.

- the selection statistic of fit for the statistical forecast for each series

- the selection statistic of fit for the reconciled forecast for each series (if the project is reconciled)

For more information about how to customize the filtered table, see "Customizing the Filtered Table" on page 26.

By default, the series are sorted by the value of the statistic of fit for the statistical forecast. Click a column heading to sort the table in ascending order by that column. You can sort only by one column at a time.

Figure 4.2 Filtered Table panel

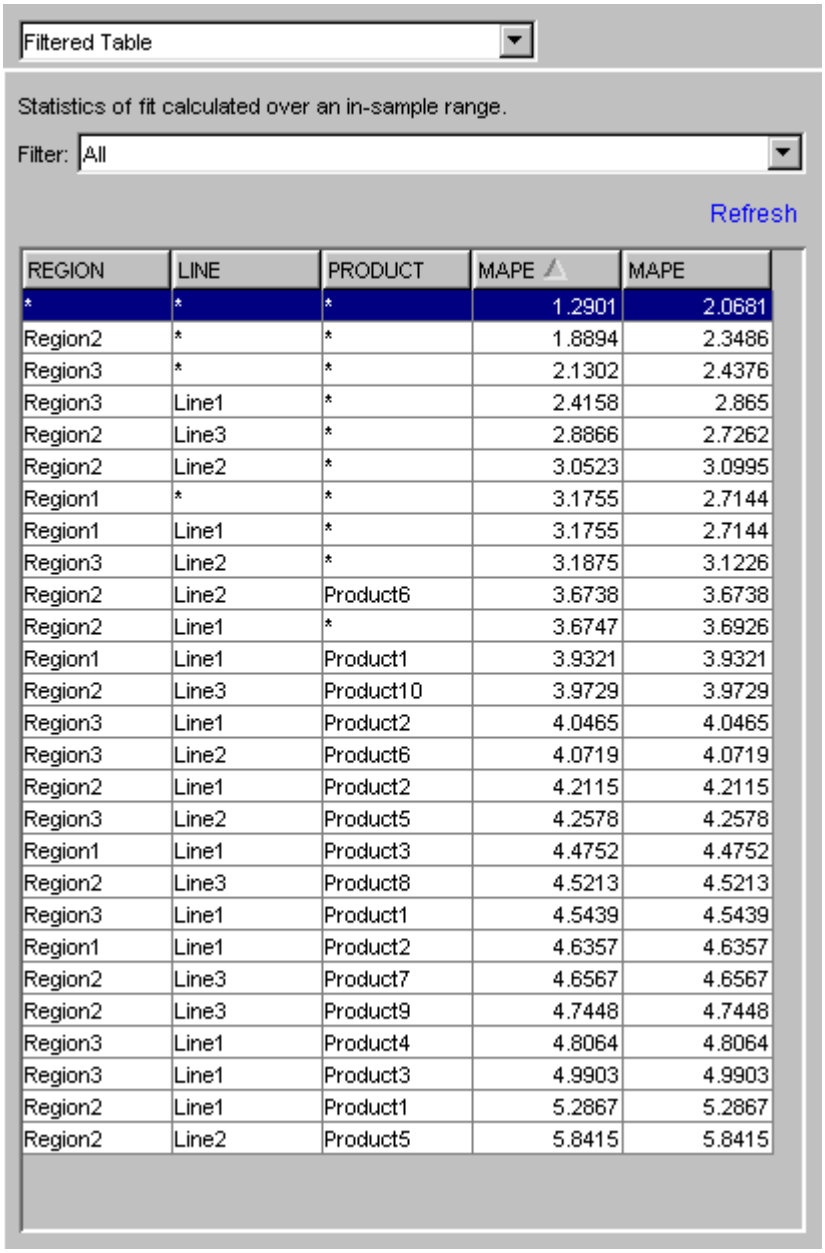

REGION	LINE	PRODUCT	MAPE	MAPE
*	*	*	1.2901	2.0681
Region2	*	*	1.8894	2.3486
Region3	*	*	2.1302	2.4376
Region3	Line1	*	2.4158	2.865
Region2	Line3	*	2.8866	2.7262
Region2	Line2	*	3.0523	3.0995
Region1	*	*	3.1755	2.7144
Region1	Line1	*	3.1755	2.7144
Region3	Line2	*	3.1875	3.1226
Region2	Line2	Product6	3.6738	3.6738
Region2	Line1	*	3.6747	3.6926
Region1	Line1	Product1	3.9321	3.9321
Region2	Line3	Product10	3.9729	3.9729
Region3	Line1	Product2	4.0465	4.0465
Region3	Line2	Product6	4.0719	4.0719
Region2	Line1	Product2	4.2115	4.2115
Region3	Line2	Product5	4.2578	4.2578
Region1	Line1	Product3	4.4752	4.4752
Region2	Line3	Product8	4.5213	4.5213
Region3	Line1	Product1	4.5439	4.5439
Region1	Line1	Product2	4.6357	4.6357
Region2	Line3	Product7	4.6567	4.6567
Region2	Line3	Product9	4.7448	4.7448
Region3	Line1	Product4	4.8064	4.8064
Region3	Line1	Product3	4.9903	4.9903
Region2	Line1	Product1	5.2867	5.2867
Region2	Line2	Product5	5.8415	5.8415

Customizing the Filtered Table

To specify what columns to display in the filtered table, complete the following steps:

1. Select **View** → **Select Filtered Table Columns**. The Filtered Table Columns dialog box opens.
2. From the drop-down list, select a category. You can choose from the following options:
   - **Reconciled (Rec.) statistic of fit**
   - **Model statistic of fit**
   - **Model property**
   - **Series property**

   The contents in the **Available** pane change depending on the category that you selected.
3. In the **Available** pane, select the statistics of fit or property to display in the table and click  . From the **Selected** pane, you can specify the order of the columns by using the up and down arrows. You can also delete columns.
4. Select whether to include a column for overrides or a column for notes in the table.

   These items are marked with special icons in the filtered table:
   - **Overrides** – They are identified in the series table by the  icon.
   - **Notes** – They are identified in the series table by the  icon.
5. Click **OK** to close the Filtered Table Columns dialog box and to apply your changes.

About the System Exceptions Panel

The System exceptions panel displays the series where one of the following problems occurred:

- A forecast was not produced. In the **Problems** column, these series are identified by "Failed to forecast".
- SAS Forecast Studio could not fit a model from the model selection list, and instead, the system used a model from the best model selection lists that ship with SAS Forecast Studio. In the **Problems** column, these series are identified by "Model list failed". (For more information about these model selection lists, see Appendix D, “Default Model Selection Lists.”

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.

About the Forecasting View

What Is the Forecasting View?

The Forecasting View is the default view when you open the project. If you open another view, you can return to the Forecasting View by selecting **View** → **Series Details** → **Forecasting View**.

In the Forecasting View, you can view the forecast for the selected series in the hierarchy or filtered table. At the top of the Forecasting View, you can see the name of the series that you selected, the dependent variable for the project, and the statistics of fit for the series. If additional information is available, click **More** to open the Statistics of Fit dialog box.

The Forecasting View is divided into three sections: a forecast plot, the data table, and a Notes section.

Figure 4.3 Forecasting View

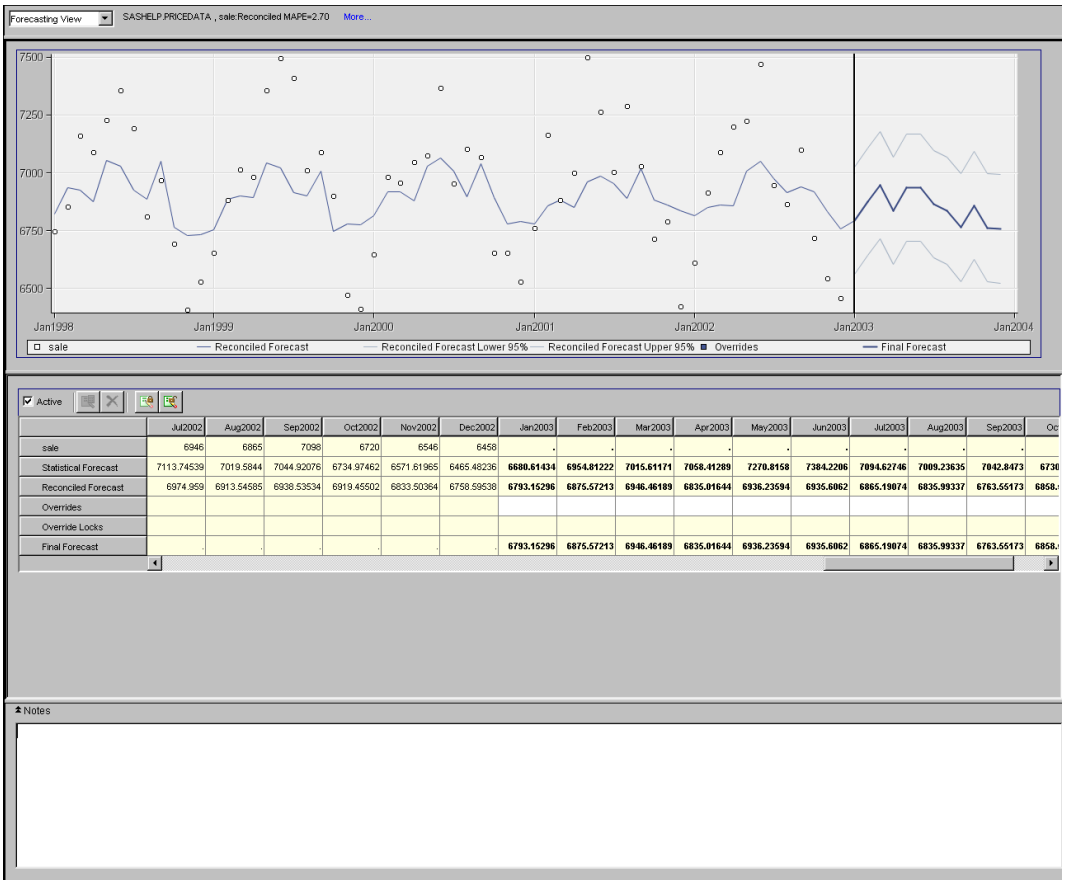

Components of the Forecasting View

About the Forecast Plot

The default content of the forecast plot depends on how you create your project.

If you forecast your data hierarchically and you select the **Reconcile hierarchy** check box when you configure the hierarchy, then the forecast plot shows the following items:

- historical data
- reconciled forecasts
- the confidence intervals for the reconciled forecasts
- any overrides that you have specified
- final forecasts
- a legend

If you do not forecast your data hierarchically or you forecast your data hierarchically but do not select the **Reconcile hierarchy** check box, then the forecast plot shows the following items:

- historical data
- the statistical forecast
- the confidence intervals for the statistical forecast
- any overrides that you have specified
- final forecasts
- a legend

You can select what to display in the forecast plot. For more information, see "Customizing the Forecasting View" on page 29.

About the Data Table

The data table appears immediately after the forecast plot in the Forecasting View.

When the **Active** check box is selected, a forecast is produced for the current series. If this check box is not selected, then no forecast is produced for the series, and the series is considered inactive. For more information, see "Specifying an Inactive Series" on page 94.

By default, the data table displays the following information:

- historical data. Historical values are displayed in regular font. Future values are displayed in bold. Missing values appear as . (a period).
- statistical forecast

- reconciled forecast
- overrides. You can specify overrides only for future time periods. The cells that you can edit have a white background.
- final forecast.

To customize the content of the data table, see "Customizing the Forecasting View" on page 29.

Figure 4.4 Data Table in the Forecasting View

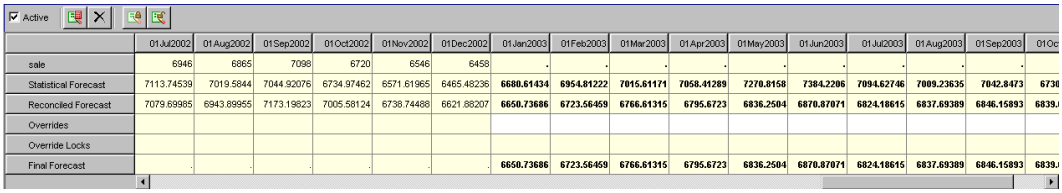

About the Notes Area

In the Notes box, you can add a note to the current series. The Notes area is hidden by default. To view the notes for the selected series, expand the **Notes** section.

NOTE: If the entire note is not visible, you can use the scroll bar to the right of the **Notes** box to view the rest of the text.

Working in the Forecasting View

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

- Manage overrides
- Add notes

Customizing the Forecasting View

Overview of the Forecasting Properties

You can choose to display the following items on the forecast plot or data table:

Option	Description
Historical data	the historical data that was in the input data set for the dependent variable. In the forecast plot, the historical data appears as white circles. In the data table, the historical data is listed in the row called *the name of the dependent variable*.
Selected model	the statistical forecast that SAS Forecast Studio generates for the dependent variable.
Selected model confidence intervals	the confidence intervals for the statistical forecast.
Aggregate event lifts	the aggregate effect of the events on the reconciled forecast for the selected series. In the forecast plot, the aggregate effect is represented by a green line. Each positive or negative change in the direction of the green line indicates an event lift.
Effect of outliers	the increase or decrease to the forecast when outliers are included.
Effect of reconciliation	the value that was added to or subtracted from the statistical forecast to determine the reconciled forecast. This value appears only in the data table.
Reconciled forecast	the total of the statistical forecast plus the effect of the reconciliation.
Reconciled forecast confidence intervals	the confidence intervals for the reconciled forecast.
Overrides	the override values that you specified for any time period. When you select this option, the overrides appear as dark blue circles in the forecast plot. In the data table, an Overrides row and an Override Locks row appear. The Overrides row displays the value of override. The Override Locks row contains a checkmark if you have specified to lock the override. For more information, see "Adding Overrides" on page 100 and "Lock and Unlock Overrides" on page 102.
Effect of overrides	the increase or decrease to the reconciled forecast to get the final forecast.

Row	Description
Final forecast	the combination of the reconciled forecast and the effect of the overrides, if you forecast your data hierarchically. If you did not forecast your data hierarchically or there is no reconciled forecast, then the value of the final forecast depends on the following: • If you did not specify an override, then the final forecast is equal to the statistical forecast. • If you specified an override, then the final forecast is equal to the override. **NOTE:** When you specify an override, then the value of the final forecast might not immediately reflect the override. • If you have forecast your data hierarchically, a warning message appears at the top of the Forecasting View to indicate whether the reconciliation is out of date and if you must reconcile the hierarchy. For more information, see "Understanding Hierarchy Reconciliation" on page 106. • If you did not forecast your data hierarchically or you did not select the **Reconcile hierarchy** check box when you configured the hierarchy, then the final forecast immediately reflects the value of the override.
Legend	a legend for the forecast plot

NOTE: The Reconciled Forecast, Override Locks, and the Effect of Overrides rows only appear in the data table if you have forecast your data hierarchically and if you selected the **Reconcile hierarchy** option when you created the project. For more information about how to configure the hierarchy when you create a project, see "Overview of the New Project Wizard" on page 42.

Customize the Contents of the Forecast Plot or Data Table

To specify what to display in the forecast plot or in the data table, complete the following steps:

1. Select **View** → **Edit Forecasting View**. The Customize Forecasting Properties dialog box opens.
2. Select what to display on the forecasting plot and what to display in the data table. For more information about the available options, see "Overview of the Forecasting Properties" on page 29.

   NOTE: Click **Reset** to return to the default values.
3. Click **OK** to close the Customize Forecasting Properties dialog box and apply your changes.

About the Model View

What Is the Model View?

In the Model View, you can view the model selection list for the selected series. You can also use plots and tables to compare how different models fit the data.

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

Figure 4.5 Model View

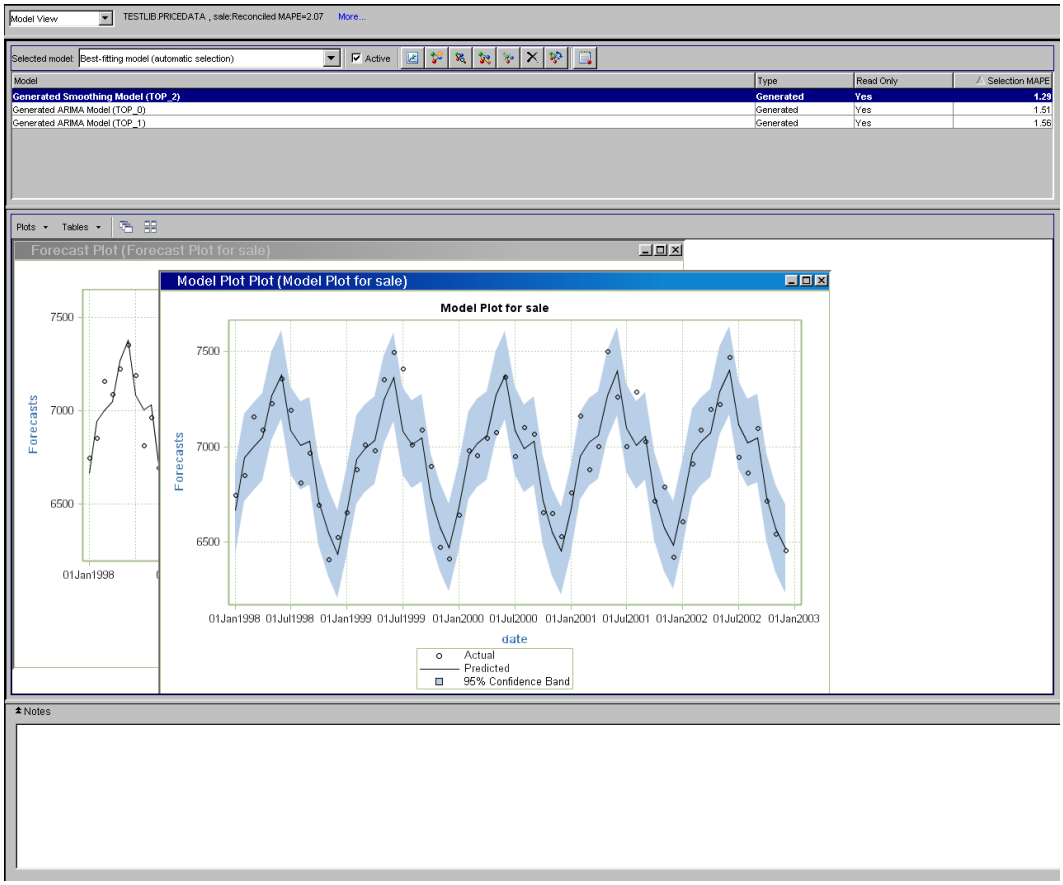

Components of the Model View

About the Model Selection List

The model selection list shows the models that have been fitted to this series. For each model, this list displays the model name, the model type, whether the model is read-only, and the fit criterion

for the model. The name of the forecast model is in bold. For more information, see "Using the Model Selection List" on page 191.

Figure 4.6 Model Selection List in Model View

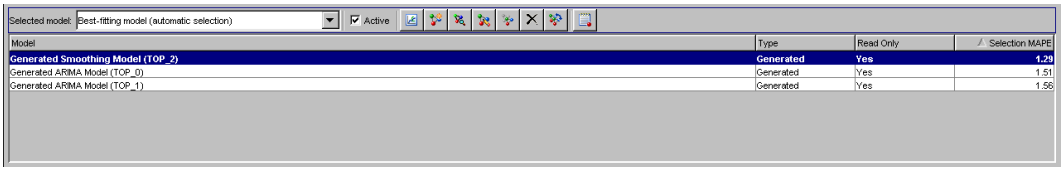

About the Model Plots and Tables

By default, the forecast plot for the dependent variable is displayed in the Model View. You can open additional model plots or additional model tables. You can use these plots and tables to compare how different models fit the data. For more information, see "View the Model Plots" on page 193 and "View the Model Tables" on page 194.

NOTE: For some series (such as a series with missing values), some of the plots or tables might not be available. If the current transformation cannot be displayed, then a message appears in the window for the plot or table that you selected.

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 View, you can choose from the following options:

-  to cascade the open plots and tables
-  to tile the open plots and tables

NOTE: SAS Forecast Studio saves your final configuration and uses that same configuration the next time that you open the Model View while working in the current project. When you close the current project, then these settings are lost.

Working in the Model View

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

- Select a model to forecast.
- Set a series as inactive.
- Reforecast a series.
- Edit the selected model.
- Create a new model.

- Copy an existing model.
- Compare models.
- Remove models from the model selection list.
- View model plots.
- View model tables.

About the Series Analysis View

What Is the Series Analysis View?

From the Series Analysis View, you can choose the time series plots and tables that you want to display for the series that is selected in the hierarchy or filtered table. The plots in the Series Analysis View use the historical data for the time series. You can also transform a time series from this window.

NOTE: If you have selected a series with all missing values, then you cannot create plots or tables in the Series Analysis View.

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

If you have several plots open simultaneously, you might want to cascade or tile the plots in the window. To organize the plots or tables in the Series Analysis View, you can choose from the following options:

-  to cascade the open plots and tables
-  to tile the open plots and tables

NOTE: For some series or time series transformations, some of the plots or tables might not be available. If the current transformation cannot be displayed, then a message appears in the window for the plot or table that you selected.

Figure 4.7 Series Analysis View with the Plots Tiled in the Window

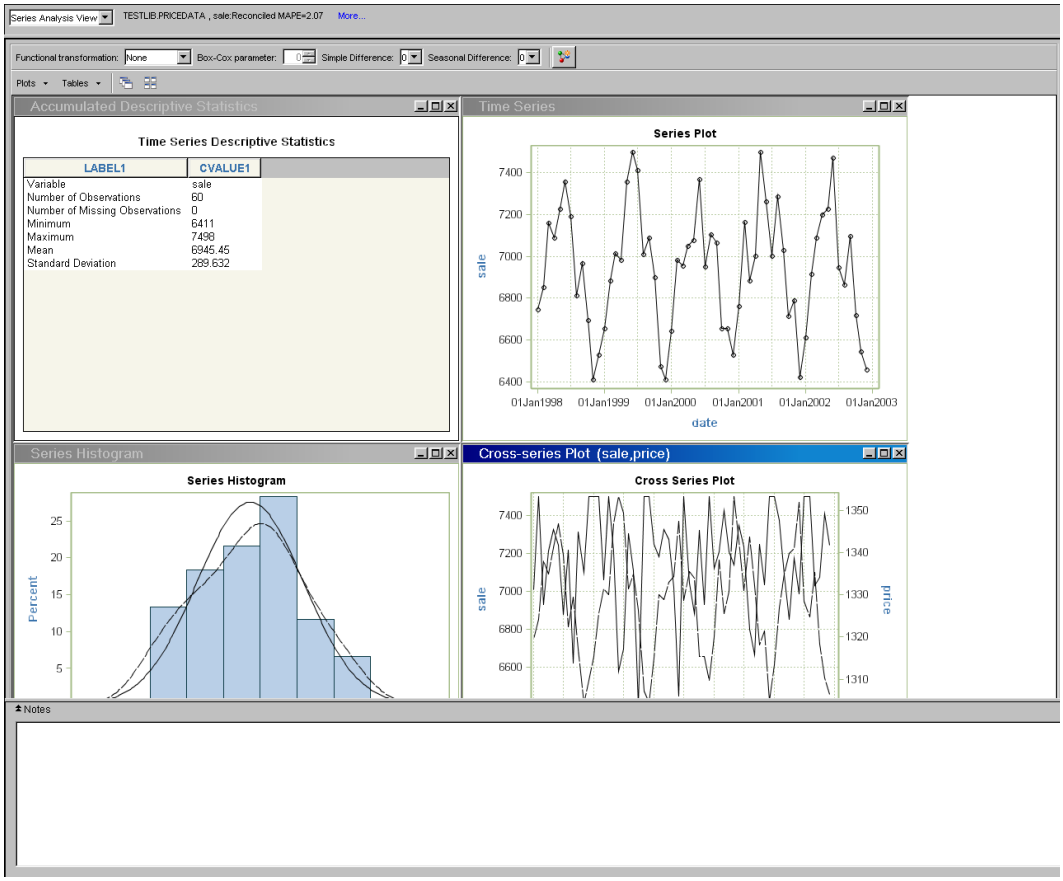

Working in the Series Analysis View

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.

NOTE: Not all plots might be available for the current series. If a plot contains content that does not apply to the current series, then SAS Forecast Studio does not open this plot when you select the plot option from the drop-down list. For example, if you do not have any independent or reporting variables in your project, then you cannot create a cross-series plot. If you remove an independent or reporting variable from your project, then any cross-series plots that depend on that variable will close.

To open a time series plot, complete the following steps:

1. Click **Plots**.
2. From the drop-down list, select from the following options:

- **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-correlation Function** - displays correlations between the dependent variable and a second variable. From the Cross 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 Cross Series window, you choose the second variable from a list of independent and reporting variables.
- **Cycle Component** - displays a plot of the cycle component that is estimated by decomposing or filtering 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 that is 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 the seasonally adjusted time series that are expressed as percent change.
- **Seasonal-Irregular Component** - displays a plot of the seasonal-irregular components that are estimated by decomposing the time series.
- **Seasonal Component** - displays a plot of the seasonal components that are displayed season after season.
- **Seasonal Cycles** - displays a plot of the seasonal cycles that are 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 that are estimated by decomposing the time series.
- **Trend-Cycle Component** - displays a plot of the trend and cycle components that are estimated by decomposing the time series.
- **Trend Component** - displays a plot of the trend component that is estimated by decomposing or filtering 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.

NOTE: In SAS Forecast Studio, you can copy a graph and save it to an external file. To export a graph, right-click on the graph and select **Copy** from the pop-up menu. Then paste the graph in the image editor of your choice.

View the Time Series Tables

From the Series Analysis View, you can select the tables to display.

NOTE: Not all tables might be available for the current series. If a table contains content that does not apply to the current series, then SAS Forecast Studio does not open this table when you select the table option from the drop-down list.

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

1. Click **Tables**.
2. From the drop-down list, 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 the number of observations, minimum, maximum, sum, and standard deviation for 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 when 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 Series Analysis View, you can transform a time series and the results.

To transform a time series, you can specify one or more of the following options:

- **Functional transformation** - specifies the transformation for the time series. You can choose from the following options:

- **Box-Cox parameter** - specifies the Box-Cox transformation where the parameter number is between -5 and 5. This option is available only if you select **Box-Cox** as the functional transformation.
- **Simple Difference** - specifies the simple difference.
- **Seasonal Difference** - specifies the seasonal difference.

SAS Forecast Studio immediately updates all of the open plots to show the specified transformation.

About the Status Bar

The status bar is located at the bottom of the main workspace. The status bar contains the following information:

Left Side	Right Side
The status of the current action in SAS Forecast Studio appears. For example, if you are generating forecasts, then the status bar displays the name of the procedure that is currently executing.	• A warning icon appears if there are errors in the log from the last code submission, and you have not yet opened the SAS log to view these errors. Click  to open the SAS log. • Your user ID. • The name of the SAS Forecast Server Middle Tier and the port that you are connected to.

Detaching Windows

You might want to have multiple Forecasting Views, Model Views, or Series Analysis Views open during a single session. To open the current view in a separate window, select **Window** → **Detach**. To close a window, click  . When SAS Forecast Studio reforecasts the project, it automatically closes any detached windows.

NOTE: Detached windows are read-only. Some functionality (such as adding overrides or adding notes in the Forecasting View) is not available from detached windows.

Chapter 5

Working with Projects

Contents

What Is a Project?

SAS Forecast Studio organizes the files that are created when you generate a forecast into a project. Each project is saved on the SAS Workspace Server. (This is the same server where the SAS processing occurs.)

By default, the project files are saved in the following locations:

Windows operating environments
!ROOT\SAS\ForecastStudio\Projects*project-name*

UNIX operating environments
!ROOT/SAS/ForecastStudio/Projects/*project-name*

The name of the project directory is the name that you specify when you create a project.

For a complete list of the directories that are part of an individual project, see the *SAS Forecast Server Administrator's Guide*.

Managing Your Projects

You can manage your projects using the Projects dialog box. For each project, you can view the project name and the date that the project was last saved. To open the Project dialog box, select **File** → **Projects** or click  in the toolbar.

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

- Create a new project. Click **New**. The New Project Wizard opens. For more information, see "Overview of the New Project Wizard" on page 42.
- Open an existing project. Select the project and click **Open**. You can have only one project open at a time.

  NOTE: Dimmed projects can be opened only by the person who created the project. To share a project, select the **Allow other users to view and edit this project** check box in the Properties dialog box.
- View the properties of a project. Select the project from the list and click **Properties**.
- Delete a project. Select from the list the project that you want to delete and click **Delete**. You are prompted to confirm your selection. If you are certain that you want to delete the project, click **OK**.

  NOTE: You should not delete projects by deleting the project directory on the SAS Workspace Server. When you delete a project directory, not all of the metadata that is associated with the project is removed. Consequently, the project still appears in the Projects dialog box, but you cannot select it. Therefore, it is recommended that you delete projects by using SAS

Forecast Studio. For more information about the project directory, see "What Is a Project?" on page 40.

Opening a SAS Forecast Studio 1.4 Project

Changes to Projects in SAS Forecast Studio 2.1

You can open a SAS Forecast Studio 1.4 project in SAS Forecast Studio 2.1. When you open a 1.4 project in SAS Forecast Studio 2.1, you will see the following changes:

- For time periods where the actual value of the independent variable is missing, you might see a difference in the forecasts for the dependent variable. When the actual value of the independent variable is missing, SAS Forecast Studio forecasts the value of the independent variable. In 2.1, improvements were made to the modeling of stochastic input variables. These improvements could result in different forecasts for the independent variables, which could result in different forecasts of the dependent variables.
- For filters, you will see the following changes:
  - In 1.4, the results from a filter were prepared on the client-side. In 2.1, these results are prepared on the server. Changing where the results are prepared was done to help improve scalability. However, you might see slower results when you a refresh a filter, reapply a filter, or sort the filter results. The results might be especially slow if the filter results contain a large number of series; the filter is complex and contains conditions from many different categories; or the results that are displayed in the filtered table of the navigation panel contain columns from many different categories.
  - Filter definitions are different in 2.1. For example, in 2.1, you can filter on model properties that you could not filter on in 1.4. Any filters that you defined in 1.4 are automatically converted to the 2.1 format.
- By default, no functional transformation is applied to the series. In 1.4, SAS Forecast Studio tried to automatically choose between no transformation and a logarithmic transformation.

Updating the Project Version

When you open a SAS Forecast Studio 1.4 project in SAS Forecast Studio 2.1, the Update Project Version dialog box opens.

When updating a project, follow these steps:

1. (Optional) Specify a name and location for the project. By default, SAS Forecast Studio uses the current name and location. To prevent the older project from being overwritten, you might want to rename or specify a different location for the updated project.
2. Specify how to update the statistical models for each series. You can choose from the following options:

- **Diagnose: create new generated models, refit all models, and select a model** - specifies that SAS Forecast Studio should create the forecasts by performing a non-destructive diagnose.
- **Select: refit all models, and select a model** - specifies that SAS Forecast Studio performs the model selection, estimates the parameters of the selected model, and generates the forecasts. This is the default.

3. (Optional) Select the **Reset all series to "best fitting model (automatic selection)"** check box.

   When the project is created, SAS Forecast Studio automatically selects the best-fitting model as the forecast model for each series. However, using the **Selected model** drop-down list in the Model View, you can change the forecast model for a series.

   This option is available only when you have selected a model other than the best-fitting model from the **Selected model** drop-down list. If you select this option, then SAS Forecast Studio selects the best-fitting model and ignores any models that you might have selected for a series. For more information about how you select a model, see "Specifying the Forecast Model" on page 195.

4. (Optional) Specify whether the project should use any updated data, if it is available. If you select the **Use updated data if available** check box, then SAS Forecast Studio determines whether the data source has been updated, and if so, updates the project data and reforecasts the existing series based on the option that you selected.

5. Click **OK**. SAS Forecast Studio automatically converts any 1.4 projects to be compatible with 2.1. If you try to open this project again in SAS Forecast Studio 1.4, the project might not work as expected.

Creating a New Project

Overview of the New Project Wizard

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 **New** in the Projects 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. Name your project.
2. Select the data set.
3. Assign data to roles.
4. Configure the hierarchy (optional).
5. Set forecasting options.
6. Finalize the project.

Name Your Project

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

To specify the project name and description, 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. The project name can also contain an underscore. For more information about SAS naming conventions, see *SAS Language Reference: Concepts*.
2. (Optional) Provide a brief description of the project. A description helps other users better understand the function of your project.
3. (Optional) 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 creator of the project can open it.
4. Click **Next** to move to the next step in the wizard.

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 when SAS Forecast Studio creates the project. If you click **No**, then you must rename the new project before you can continue.

Select the 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. Double-click the name of the library to view the data sets within that library. Select the data set that you want to use. To view the contents of the selected data set, click **View**.

   To return to a previous level, click  .

   NOTE: You should not create a project using a data set in the Work library. The Work library is a temporary directory that exists only for the current session. Any data that is within this library is deleted when you close SAS Forecast Studio. If you create a project using a data set in the Work library, then SAS Forecast Studio will be unable to find this data the next time that you open the project. Also no files in the Work library should be modified while SAS Forecast Studio is running.

2. Click **Next** to move to the next step in the wizard.

NOTE: Before you can use a data set in SAS Forecast Studio, it must be in the required format. For more information, see Chapter 3, "Preparing an Input Data Set for SAS Forecast Studio."

Assign Data to Roles

In the New Project Wizard, you can assign the variables in the data to roles in the project. After the project is created, you cannot change the time ID variable, BY variables, and dependent variables. However, you can modify or add independent and reporting variables to the project. For more information, see "Modify Variable Assignments" on page 49.

To assign the data to roles, 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.

     SAS Forecast Studio analyzes the variable assigned to the time ID role to detect the time interval of the data. If the time interval cannot be detected from the variable that you assign, then you are prompted to select another variable as the time ID or to manually specify the time interval and other time ID properties. For more information, see "Troubleshooting the Time ID Variable" on page 49.

   - **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 detected, verify 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 any explanatory, input, predictor, or causal factor variables. You can assign only numeric variables to this role.
- **Reporting Variables** - specifies the variables that you want to include in the reports. These variables are not used to generate the forecasts.

For more information, see "How to Assign a Variable to a Role" on page 48.

2. Click **Add Events** to add an event to your project. For more information, see "Managing Events" on page 72.

3. Click **Adjustments** to add an adjustment variable. For more information, see "Managing Adjustment Variables" on page 51.

4. Click **Next** to move to the next step.

   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 49.

   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 Appendix A, "Reserved Variable Names in SAS Forecast Studio."

Configure the Hierarchy

For every project, you must assign at least one dependent variable. The options that are available from this step differ depending on the remaining variables that you assign for the project.

- If you assign at least one BY variable, then you can select whether to forecast the data hierarchically.
- If you assign more than one dependent variable or you did not assign any BY variables, then the following options are not available. If you do not want to forecast your data hierarchically, then you can move to the next step in the wizard. If you do want to forecast your data hierarchically, then you need to return to the Assign Data to Roles step to change your variable assignments.

To configure a hierarchy, complete the following steps:

1. Specify whether to forecast hierarchically. You can choose from the following options:

- **Do not forecast a hierarchy** - enables you to forecast the data without creating a hierarchy. Selecting this option enables you to assign multiple variables to the dependent variable role.
- **Forecast 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, SAS Forecast Studio sorts the dependent variable into individual series in a hierarchy using the BY variables. If you choose to forecast hierarchically, then only one dependent variable can be forecast.

2. 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.

3. In the table, specify the statistic for aggregating data across the levels of the hierarchy. By default, all the dependent variables, independent variables, adjustment variables, and reporting variables are listed in the table. (The reporting variables are listed as auxiliary in the Type column.) In the Method column, select the aggregation statistic to use for each variable. The options that are available depend on the variable. For example, if you are aggregating data for the dependent variable, then only the **Sum of values** and **Average of values** options are available.

   For more information about the aggregation methods, see "Understanding Aggregation and Accumulation" on page 82.

4. If you want to reconcile the hierarchy, select the **Reconcile hierarchy** check box and then specify the following options:

   - From the **Reconciliation method** drop-down list, select the reconciliation method that you want to use. You can choose from the following options:
     - **Top Down** - aggregates the data from the lowest levels in the forecast and then uses these values to generate the forecasts at the highest level. SAS Forecast Studio then uses this forecast and the disaggregation method that you specified to reconcile 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.
     - **Bottom Up** - uses the data at the lowest level of the hierarchy to generate the forecasts. These forecasts are then used to reconcile the 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** - aggregates the data from the lower levels and then uses these values to generate the forecasts for the middle level. SAS Forecast Studio uses the forecasts at the middle level to reconcile the forecasts for both the higher and lower levels.

NOTE: Because some hierarchies have more than one middle level, you need to specify the level that you want to use. If you have more than one middle level, then the option name is **Middle Out - *name-of-level***.

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

  For an example of how forecasts are generated based on the reconciliation method and disaggregation method, see "Example: How Forecasts Are Generated for a Hierarchy" on page 54.
- Click **Advanced** to specify the advanced reconciliation settings.

NOTE: If you clear the **Reconcile hierarchy** check box, then the **Reconciliation method**, **Disaggregation method**, and **Advanced** options are not available.

5. Click **Next** to move to the next step.

Set Forecasting Options

To set the forecasting options, complete the following steps:

1. Specify the forecasting options to use for the project. These options are identical to the options in the Forecasting Settings dialog box. For more information about these options, see "Working with the Forecasting Settings for the Project" on page 84.
2. Click **Next** to move to the next step.

Finalize the Project

To finish the project, complete the following steps:

1. 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.
2. Click **Finish**. You can monitor the progress of the forecasts from the status bar and review a summary of the run from the Forecast Summary dialog box after SAS Forecast Studio is finished. For more information, see "Reviewing a Summary of the Forecast Run" on page 93.

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 Projects dialog box.
- After SAS Forecast Studio is running, select **File** → **Projects** to open the Projects dialog box. For more information, see "Managing Your Projects" on page 40.

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

Working with Project Variables

How to Assign a Variable to a Role

In the New Project Wizard or the Variables dialog box, you assign the variables in the data set to a role. By default, SAS Forecast Studio lists the name of each variable. To view the label for a variable, place your cursor over the variable name. The variable label appears in the tooltip.

To assign a variable to a role in the New Project Wizard or the Variables dialog box, select the variable in the **Data set variables to assign** box and click  for that role.

To remove a variable from a role, click  .

NOTE: After the project is created, you cannot change the time ID variable, BY variables, and dependent variables. However, you can modify or add independent and reporting variables to the project.

Icons beside the variable names indicate the type of variable:

Icon	Description
	numeric variable
	character variable
	date/time variable

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

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

Modify Variable Assignments

When you create a project, you assign variables to roles in the New Project Wizard. After the project is created, you cannot change variable assignments that you made to the time ID, BY variable, and dependent variable roles. However, you can modify or add independent and reporting variables to the project by using the Variables dialog box. For more information about each of the available roles for a project, see "Assign Data to Roles" on page 44.

To modify the variable assignments, complete the following steps:

1. Select **Project** → **Variables** to open the Variables dialog box.
2. Assign each variable in the **Data set variables to assign** box to the appropriate role. For more information, see "How to Assign a Variable to a Role" on page 48.
3. Click **OK**.

NOTE: When you add an independent variable, SAS Forecast Studio performs a nondestructive diagnosis. SAS Forecast Studio preserves any user-created models. Any system-generated models are re-created. When you remove an independent variable, SAS Forecast Studio performs a destructive diagnosis. SAS Forecast Studio does not preserve any user-created models. Both the user-created and system-generated models are re-created.

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 Chapter 3, "Preparing an Input Data Set for SAS Forecast Studio".

Common errors include the following:

- SAS Forecast Studio is unable to determine the time interval. When SAS Forecast Studio is unable to determine the time interval for the data, then the Specify Time ID Interval dialog box opens. You can use this dialog box to select another variable as the time ID variable or to manually set the time interval.

  Here are the reasons why SAS Forecast Studio cannot determine the time interval:

  - The values of the time ID variable are not evenly-spaced values at a specific time interval. For example, transactional data is recorded at no particular time interval.

    In the Specify Time ID Interval dialog box, select the **Manually set the time interval and other time ID properties** option and specify the time interval to use. Then SAS Forecast Studio can aggregate the data using that time interval.

    NOTE: If the specified time interval is too small, then a large number of observations could have missing values. For example, you select WEEK as the interval, but there is no data for many of the weeks. A larger interval, such as MONTH, should be selected.
  - The values of the time ID variable are evenly-spaced values at a specific time interval, but SAS Forecast Studio cannot detect the time interval. For example, SAS Forecast Studio cannot detect the time interval if the data represents a multi-weekday interval, such as the WEEKDAY3. interval.

    In the Specify Time ID Interval dialog box, select **Manually set the time interval and other time ID properties** option and specify the time interval to use.
  - The wrong variable has been selected as the time ID variable.

    In the Specify Time ID Interval dialog box, select the **Select another variable as a time ID** option. SAS Forecast Studio returns you to the Assign Variables to Roles step in the New Project Wizard, so you can select a different variable to use as the time ID.

- The values of the time ID variable are in an unsupported format, such as data that contains SAS time values. SAS Forecast Studio supports only SAS date and datetime values.

  If your data is hourly or uses a smaller time interval, then the values in the time ID variable must use the datetime format (which is the number of seconds since January 1, 1960). If the data uses a SAS time format, then SAS Forecast Studio does not allow you to select that variable as the time ID. If the time ID variable has time values (the number of seconds since midnight) but no format, then SAS Forecast Studio assumes that the time ID values are dates (the number of days since January 1, 1960).

- 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 not aggregated over that variable.

Working with Adjustment Variables

What Are Adjustments?

Systematic variations and deterministic components are included in time series data. You can specify adjustment variables to identify the data that should be excluded before a statistical analysis.

Using SAS Forecast Studio, you can specify adjustments in the following ways:

- before generating the forecasts. After the time-stamped data has been accumulated and interpreted, you might need to adjust the time series before generating forecasts. By adjusting the

time series for any known systematic variations or deterministic components, the underlying time series process can be more easily identified and modeled.

Examples of systematic adjustments are exchange rates, currency conversions, and trading days. Examples of deterministic adjustments are advanced bookings and reservations and contractual agreements.

- after the forecasts have been generated. You might need to adjust the statistical forecast to return the forecasts to the metric used in the original data.

Generally, the adjustments before and after generating the forecasts are operations that are inverses of each other.

Here are two examples of how adjustments are used.

- To adjust a time series for an exchange rate, you divide the time series by the exchange rate. You then analyze and forecast the adjusted time series without regard to the exchange rate. After this analysis, you adjust the forecasts by multiplying the values by the exchange rate.
- To adjust a time series for advanced bookings, you subtract the advanced bookings from the time series. You then analyze and forecast the adjusted time series without regard to the advanced bookings. After this analysis, you adjust the forecasts by adding the advanced bookings.

Managing Adjustment Variables

When you create a project, you can add and remove adjustment variables from the project by using the Adjustments dialog box. You can also edit and specify the order of existing adjustment variables.

In the Adjustments dialog box, you can view the list of adjustment variables, the dependent variables being adjusted, and a description.

NOTE: You cannot add, edit, change the order, or remove adjustment variables after you have created the project, but you can view the adjustment variables and the properties of these variables by selecting **Project→ Variables**. The Variables dialog box opens. Click **Adjustments** to open the Adjustments dialog box.

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

- add an adjustment variable
- edit an adjustment variable
- change the order of the adjustment variables
- delete an adjustment variable

Add an Adjustment Variable

You can add adjustment variables only when you create a project. After you have created the project, you can view your adjustment variables from the Adjustments dialog box. For more information, see "Managing Adjustment Variables" on page 51.

To add an adjustment variable when you create a project, complete the following steps:

1. In the Assign Data to Roles step of the New Project Wizard, click **Adjustments**. The Adjustments dialog box opens.
2. In the Adjustments dialog box, click **New**. The Adjustment Properties dialog box opens.
3. Select the adjustment variable. The variable assigned as the adjustment variable cannot have any other role in the project.
4. (Optional) Specify a description of the adjustment variable.
5. Specify the operations to perform before and after forecasting. You can choose from minimum, maximum, add, subtract, multiply, and divide. By default, no operation is selected.
6. Select the check box of the dependent variable to adjust. You must select at least one variable to adjust.
7. Click **OK** to return to the Adjustments dialog box.

Edit an Adjustment Variable

You can edit adjustment variables only when you create a project. After you have created the project, you can view your adjustment variables from the Adjustments dialog box. For more information, see "Managing Adjustment Variables" on page 51.

To edit an adjustment variable when you create a project, complete the following steps:

1. In the Assign Data to Roles step of the New Project Wizard, click **Adjustments**. The Adjustments dialog box opens.
2. In the Adjustments dialog box, select the adjustment variable that you want to edit and click **Edit**. The Edit *name-of-adjustment-variable* dialog box opens.
3. Modify the options that you want to change and click **OK**. The options in the Edit *name-of-adjustment-variable* dialog box are identical to when you created the adjustment variable. For more information, see "Add an Adjustment Variable" on page 52.

Order the Adjustment Variables

You can specify the order of the adjustment variables only when you create a project. After you have created the project, you can view your adjustment variables from the Adjustments dialog box. For more information, see "Managing Adjustment Variables" on page 51.

To order the adjustment variables when you create a project, complete the following steps:

1. In the Assign Data to Roles step of the New Project Wizard, click **Adjustments**. The Adjustments dialog box opens.
2. In the Adjustments dialog box, select the adjustment variable that you want to reorder and click the up or down arrow. The variable that you selected moves up or down one position in the list.

Delete Adjustment Variables

You can delete adjustment variables only when you create a project. After you have created the project, you can view your adjustment variables from the Adjustments dialog box. For more information, see "Managing Adjustment Variables" on page 51.

To order the adjustment variables when you create a project, complete the following steps:

1. In the Assign Data to Roles step of the New Project Wizard, click **Adjustments**. The Adjustments dialog box opens.
2. In the Adjustments dialog box, select the adjustment variable that you want to delete and click **Delete**. You are prompted to confirm this deletion. Click **OK** to delete the variable.

Working with the Hierarchy

Modify the Hierarchy Settings

When you create a project, you can configure the hierarchy in the New Project Wizard. After the project is created, you can modify the hierarchy by using the Hierarchy Settings dialog box.

To modify the hierarchy, complete the following steps:

1. Select **Project** → **Hierarchy Settings** to open the Hierarchy Settings dialog box.
2. Specify the options for the hierarchy. The options that are available from the Hierarchy Settings dialog box are the same options that appear in the New Project Wizard. For more information about the hierarchy options, see "Configure the Hierarchy" on page 45.

   NOTE: If you change the aggregation method for a variable, then SAS Forecast Studio performs a nondestructive diagnosis. SAS Forecast Studio preserves any user-created models. Any system-generated models are re-created.
3. Click **OK**.

Example: How Forecasts Are Generated for a Hierarchy

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 5.1 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	Aggregation Forecast	Disaggregation Forecast
Product Line	Aggregation Forecast	Disaggregation Forecast	Forecast	Disaggregation Forecast
Product	Forecast	Disaggregation Forecast	Disaggregation Forecast	Disaggregation Forecast

Save a Project

To save a project, complete the following steps:

1. Select **File** → **Save Project As**. The Save Project dialog box opens.
2. Specify a name for the project. By default, the name is *project-name*_Copy_*n*, where *n* is the lowest available integer.
3. Click **OK**.

   NOTE: If a project with that name already exists, you are prompted to overwrite the existing project.

The new project is saved to the SAS Forecast Studio directory.

Update the Project Data

Before reforecasting the data, you can check to see whether 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 Update Project Data dialog box opens.
2. Select the action that you want to perform. You can choose from the following options:
   - **Diagnose: create new generated models, refit all models, and select a model** - specifies that SAS Forecast Studio should create the forecasts by performing a non-destructive diagnose.
   - **Select: refit all models, and select a model** - specifies that SAS Forecast Studio performs the model selection, estimates the parameters of the selected model, and generates the forecasts. This is the default.
   - **Fit: refresh the current forecast model, updating the parameter values** - specifies that SAS Forecast Studio estimates the parameters using the current model. No model selection is performed.
   - **Forecast: refresh the current forecast model, using the same parameter values** - specifies that SAS Forecast Studio reforecasts using the current models and parameter estimates. No parameter estimation occurs.
   - **Do not update the project data**.
3. (Optional) For the diagnose and select options, select the **Reset all series to "best fitting model (automatic selection)"** check box.

   SAS Forecast Studio then selects the best-fitting model and ignores any models that you might have selected for a series. For more information about how you select a model, see "Specifying the Forecast Model" on page 195.
4. Click **OK**.

Based on the option that you select, SAS Forecast Studio updates the project data and reforecasts any existing series.

NOTE: If you choose any option other than **Diagnose: create new generated models, refit all models, and select a forecast model**, then any new series are forecast using the model selection list for the Best Smoothing Model, which ships with SAS Forecast Studio. For more information, see Appendix D, "Default Model Selection Lists".

Viewing Properties

View Project Properties

To view the project properties, complete the following steps:

1. Select **File** → **Project Properties**. The Project Properties dialog box opens.
2. To add a project description, type the description in the **Description** box.
3. If you want to share this project with other users at your site, select the **Allow other users to view and edit this project** check box. This option is not selected by default. When this check box is not selected, only the person who created the project can open it. Only the project creator or a user with administrative privileges can change this option.
4. Click **OK** to save your changes and to close the Project Properties dialog box.

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 filtered table, and select **View** → **Show 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 navigation panel.

You can see the following information for the selected series:

- the name of the series
- start date of the first observation in the series
- end date of the last observation in the series
- number of values
- number of missing values
- mean
- standard deviation
- minimum value
- maximum value

Figure 5.1 Hierarchy Panel with Series Properties Pane

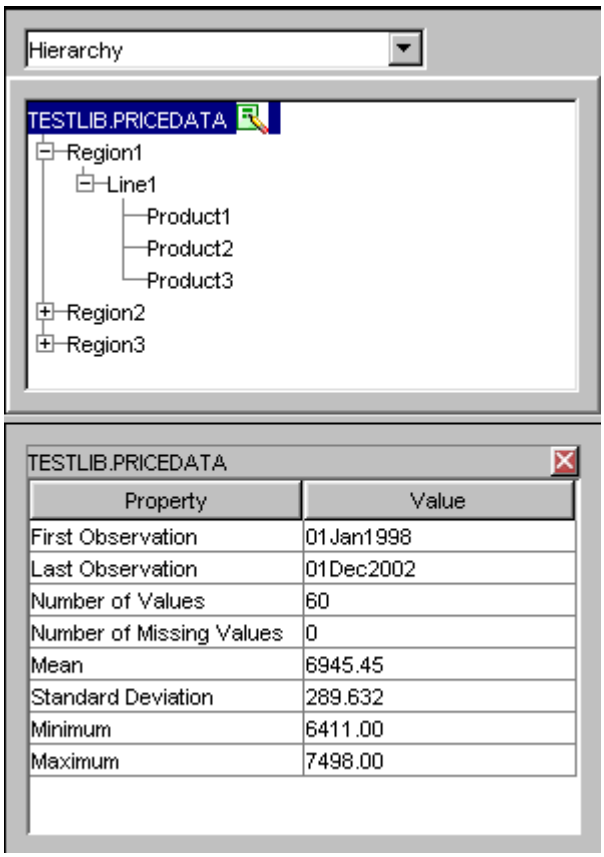

Working with the SAS Log

About the SAS Log

The SAS log is a file that contains a record of the SAS statements that you enter as well as messages about the execution of your program. The log lists any errors that occurred while the data was being processed. The log also contains warnings and messages that help you troubleshoot these errors.

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

View the SAS Log

To view the SAS log that SAS Forecast Studio creates when generating the forecast, select **Tools** → **Log**.

You might want to save this information to a file so that you can refer to it later. For more information, see "Save the SAS Log" on page 58.

Save the SAS Log

To save the SAS log, complete the following steps:

1. Select **Tools** → **Log**. The SAS Log dialog box opens.
2. Click  . The Save dialog box opens.
3. Select the location where you want to save this file and click **Save**.

Search the SAS Log

To search the SAS log, complete the following steps:

1. Select **Tools** → **Log**. The SAS Log dialog box opens.
2. Click  . The Find/Replace dialog box opens.
3. Specify the text that you want to search for. You can also specify the following:
   - whether to match the whole word and the case
   - the direction to search
   - whether to search the code, the comments, or both

   Click **Find Next** to view the first item in the log that matches the search results. To move to the next item, click **Find Next**.

Working with the SAS Code

About the SAS Code

In SAS Forecast Studio, you can view the SAS code for tasks, such as creating a model or event, or for forecasting the entire project. For more information about how to view the code for a project, see "View the Project Code" on page 59.

From the SAS Code dialog box, you view a read-only version of the SAS code that runs for that object or project. You can expand and collapse the code to view all or only portions of the code. You can also save the code to an external file. For more information, see "Save the Code to a File" on page 59.

View the Project Code

To view the project code that SAS Forecast Studio has generated, complete the following steps:

1. Select **Tools** → **Code**. The SAS Code dialog box opens.
2. Select the version of the code that you want to view from the **Task** drop-down list. The following options are available:
   - **Re-create Project** - displays the code that SAS Forecast Studio uses to re-create the project. This option includes the code for all the data preparation and forecasting steps.
   - **Forecast** - displays the code that SAS Forecast Studio uses to reforecast the model and parameters. No parameter estimation occurs. This option does not include the code for the data preparation, diagnose, and fit model steps.
   - **Fit Models** - displays the code that SAS Forecast Studio uses to estimate the parameters using the model that you specified and then to create a forecast. No model selection is performed. This option does not include the code for the data preparation and diagnose steps.
   - **Diagnose Project** - displays the code that SAS Forecast Studio uses to rediagnose the project while preserving all customizations that are created by user-defined models. This option does not include the code for the data preparation steps.
   - **Select Models** - displays the code that SAS Forecast Studio uses to perform model selection, to estimate the parameters of the selected model, and to produce forecasts. This code does not include the data preparation, diagnose, fit, and forecasting steps. This option is the default.
   - **Reset Project** - displays the code that SAS Forecast Studio uses to rediagnose all series in the project and re-create the model repositories. This option does not include the code for the data preparation steps.
3. Select the **Include code to import new data** check box to display the code for importing new data. If you select this option, then any new data that is available is used in the calculations. If you do not select this option, then only the data that was used when the project was created is used in the calculations. This option is not available if you select **Re-create Project** from the **Task** drop-down list.

Save the Code to a File

You can save the code that SAS Forecast Studio generates for a project in order to run the project in batch mode.

To save the SAS Forecast Studio code to a file, complete the following steps:

1. Select **Tools** → **Code**. The SAS Code dialog box opens.
2. Select the version of the project code that you want to view. For more information, see "About the SAS Code" on page 58.
3. In the SAS Code dialog box, click **Save As**. The Save As dialog box opens.
4. Select the location where you want to save this file and click **Save**.

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 Project**. After you close the current project, you can create a new project or open an existing project.

NOTE: If you try to create a new project or open an existing project while you have a project open, then SAS Forecast Studio automatically closes the current project.

To close SAS Forecast Studio, select **File** → **Exit**.

Chapter 6

Working with Filters

Contents

What Is a Filter?

Filters enable you to specify criteria to subset your data. In SAS Forecast Studio, you create filters to help you identify forecasts that do or do not meet a specific criterion. To view the results from a filter, use the filtered table view in the navigation panel.

For an example of how to create a filter, see "Examples of Filters" on page 64.

Managing Filters

SAS Forecast Studio can generate large numbers of forecasts. Generally, filters help you identify those forecasts that are less desirable, but you can also use filters to see which series have desirable forecasts.

After you create a project, you can manage filters by using the Filters dialog box. To open this dialog box, select **Tools** → **Filters**.

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

- create a filter

- edit a filter
- delete a filter

Create a Filter

To create a filter, complete the following steps:

1. Select **Tools** → **Filters**. The Filters dialog box opens.
2. Click **New**. The Filter Properties dialog box opens.
3. Specify a valid SAS name for the filter. This name can be 32 characters long, and it must start with a letter (A-Z). Subsequent characters can be letters, numeric digits (0-9), or an underscore (_). Both upper- and lowercase letters are valid. For more information about SAS naming conventions, see *SAS Language Reference: Concepts*.
4. (Optional) Specify a description.
5. Specify the hierarchy level for the filter. By default, the filter is applied to all the levels in the hierarchy. If you select a specific level in the hierarchy, then the results from the filter contain only series in that level of the hierarchy.

   NOTE: The **Hierarchy level** drop-down list is available only if you forecast your data hierarchically.
6. Specify the filter by using the columns in the table. The following table lists the options that are available. The values of the **Property**, **Condition**, and **Value** columns depend on the value of the **Category** column.

Category	Property	Condition	Value
Reconciled (Rec.) statistic of fit	All of the statistics of fit	=, <>, >, >=, <, <=	A text box where you can specify a value for the statistic of fit.
Model statistic of fit	All of the statistics of fit	=, <>, >, >=, <, <=	A text box where you can specify a value for the statistic of fit.

Category	Property	Condition	Value
Model property	Dependent transformation Events Independent variables Model family Model name Outliers Seasonal model Trend model	=	The available value depends on the model property that you selected: • If you selected **Dependent transformation** as the property, then you can select the type of dependent transformation to use. • If you selected **Model family** as the property, then you can select the model family to use or whether to filter for inactive models. • If you selected **Model name** as the property, then you can specify the name of the model. • If you choose any other model property, then the drop-down list contains **Yes** and **No**. The default is **No**.
Series property	All of the series properties	=, <>, >, >=, <, <=	A text box where you can specify a value for the statistic of fit.
BY variable	All of the BY variables, if you forecast your data hierarchically	=	A drop-down list that contains the values for the selected BY variable.

7. Click  if you need to add additional filter expressions.

   To remove a row, click  .

8. Specify whether the series must have overrides, notes, or both.
9. Click **OK** to save the new filter. This filter now appears in the Filters dialog box.

For an example of how to create a filter, see "Examples of Filters" on page 64.

Edit a Filter

To edit an existing filter, complete the following steps:

1. Select **Tools** → **Filters**. The Filters dialog box opens.
2. Select the filter that you want to edit and click **Edit**. The Filter Properties dialog box opens.
3. Edit the values of the filter and click **OK** to save the changes to your filter.

Delete a Filter

To delete a filter, complete the following steps:

1. Select **Tools** → **Filters**. The Filters dialog box opens.
2. Select the filter that you want to delete and click **Delete**.

View the Results from a Filter

To view the series that meet the filter criterion, complete the following steps:

1. In the navigation panel, select **Filtered Table** from the drop-down list.
2. From the **Filter** drop-down list, select the filter that you want to apply. SAS Forecast Studio automatically updates the contents of the table. The resulting rows in the table show the series that meet the filter criterion.

Examples of Filters

Example 1: Creating a Simple Filter

In this example, you want to know which series has a Mean Absolute Percent 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 Filter Properties dialog box opens.
3. In the **Name** box, type `MAPE`.

4. In the **Description** box, type `Filter to identify which series have a MAPE greater than 7`.
5. From the **Hierarchy level** drop-down list, select **All**.
6. In the table, specify the following values for each column:
   - In the **Category** column, select **Reconciled (Rec.) statistic of fit**.
   - In the **Property** column, select **Mean absolute percent error**.
   - In the **Condition** column, select **>**.
   - In the **Value** column, type `7` and press ENTER.
     The final window should look like the following:

Figure 6.1 Completed Filter Properties Dialog Box for Filter Example

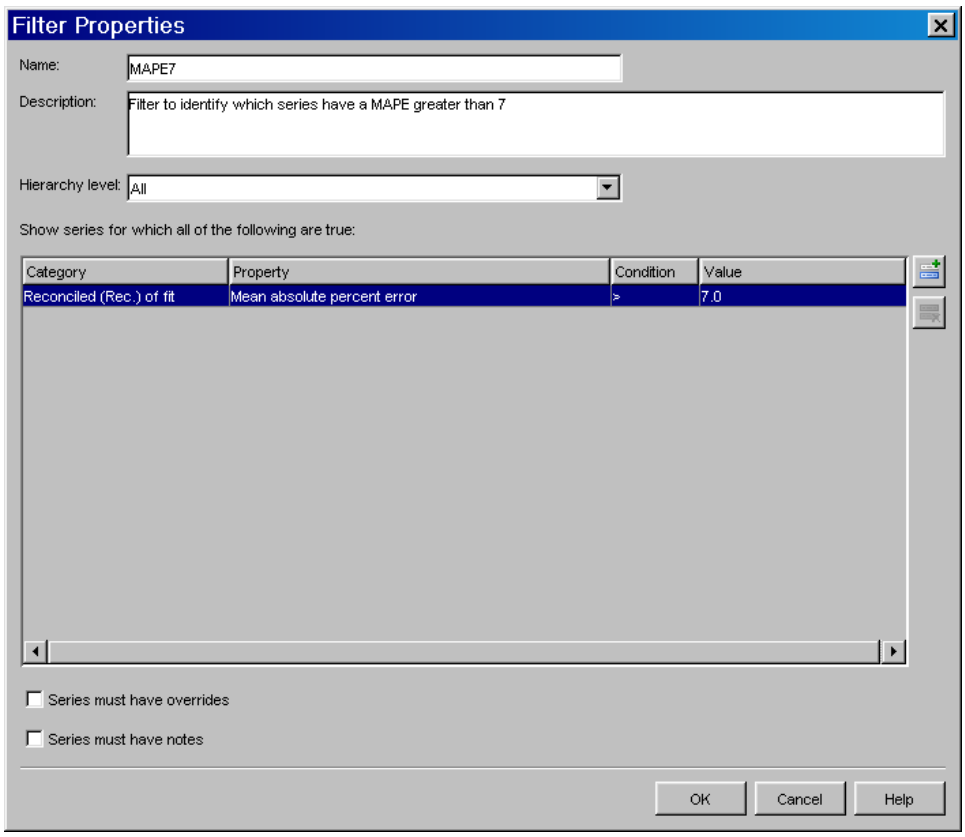

7. Click **OK** to add the filter to the list in the Filters dialog box.

Figure 6.2 New Filter in Filters Dialog Box

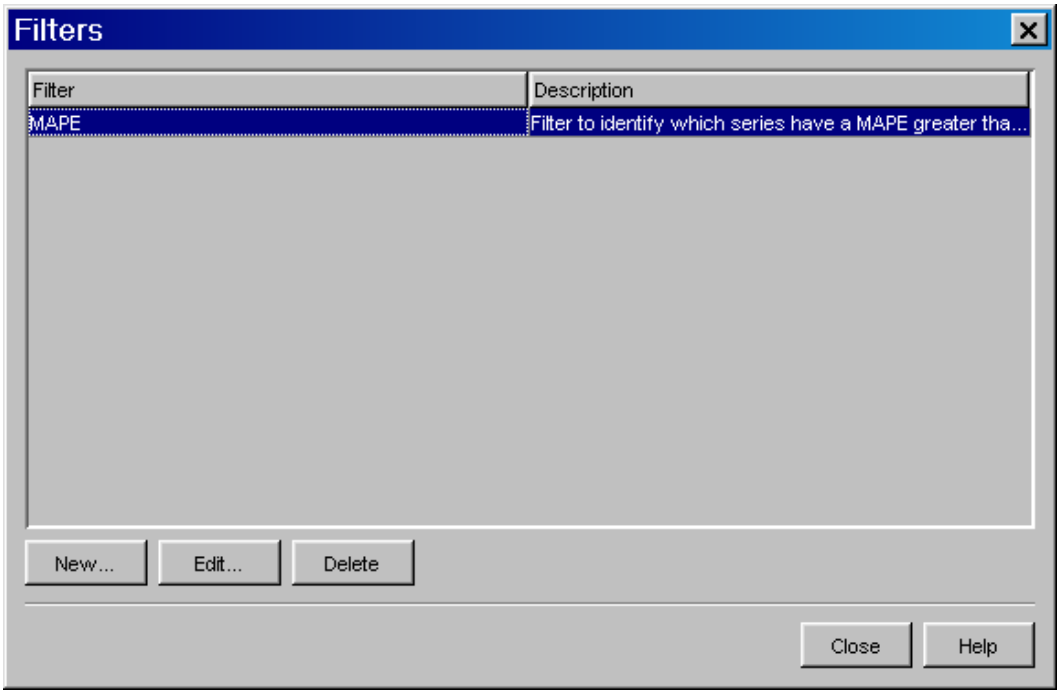

8. Click **Close** in the Filters dialog box to return to the project.

To view the series that meet this filter criterion, complete the following steps:

1. In the navigation panel, select **Filtered Table** from the drop-down list.
2. From the **Filter** drop-down list, select **MAPE**.

Here are the results:

Figure 6.3 Results in Filtered Table after Applying the MAPE Filter

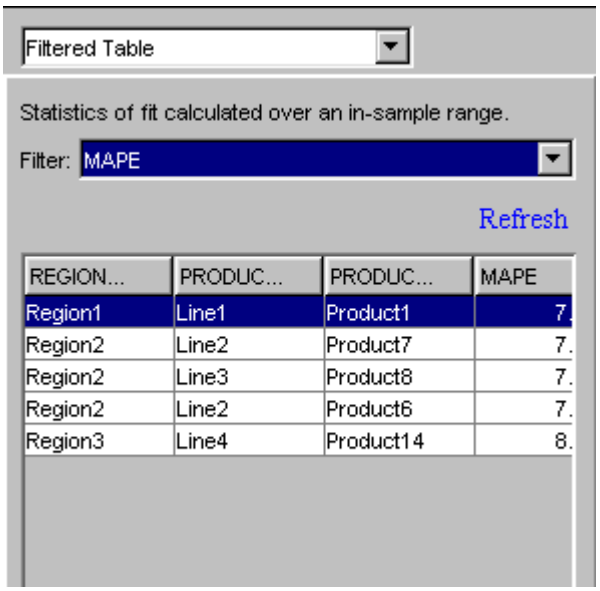

Example 2: Creating a Filter with Multiple Expressions

In this example, you want to know which series has a Mean Absolute Percent Error (MAPE) greater than 7 and a Minimum value (Min) less than 300.

To create this filter, complete the following steps:

1. Select **Tools** → **Filters**. The Filters dialog box opens.
2. Click **New**. The Filter Properties dialog box opens.
3. In the **Name** box, type `MAPE7MIN300`.
4. In the **Description** box, type `Filter to identify which series have a MAPE greater than 7 and a Minimum value less than 300.`
5. From the **Hierarchy level** drop-down list, select **All**.
6. In the table, specify the following values for each column:
   - In the **Category** column, select **Reconciled (Rec.) statistic of fit**.
   - In the **Property** column, select **Mean absolute percent error**.
   - In the **Condition** column, select **>**.
   - In the **Value** column, type `7` and press ENTER.
7. Click  . A new row appears in the Filter Properties dialog box. SAS Forecast Studio automatically joins the filter conditions by AND.
8. In the new row, specify the following values for each column:

- In the **Category** column, select **Series property**.
- In the **Property** column, select **MIN**.
- In the **Condition** column, select **<**.
- In the **Value** column, type `300` and press ENTER.

The final window should look like the following:

Figure 6.4 Completed Filter Properties Dialog Box for Filter Example

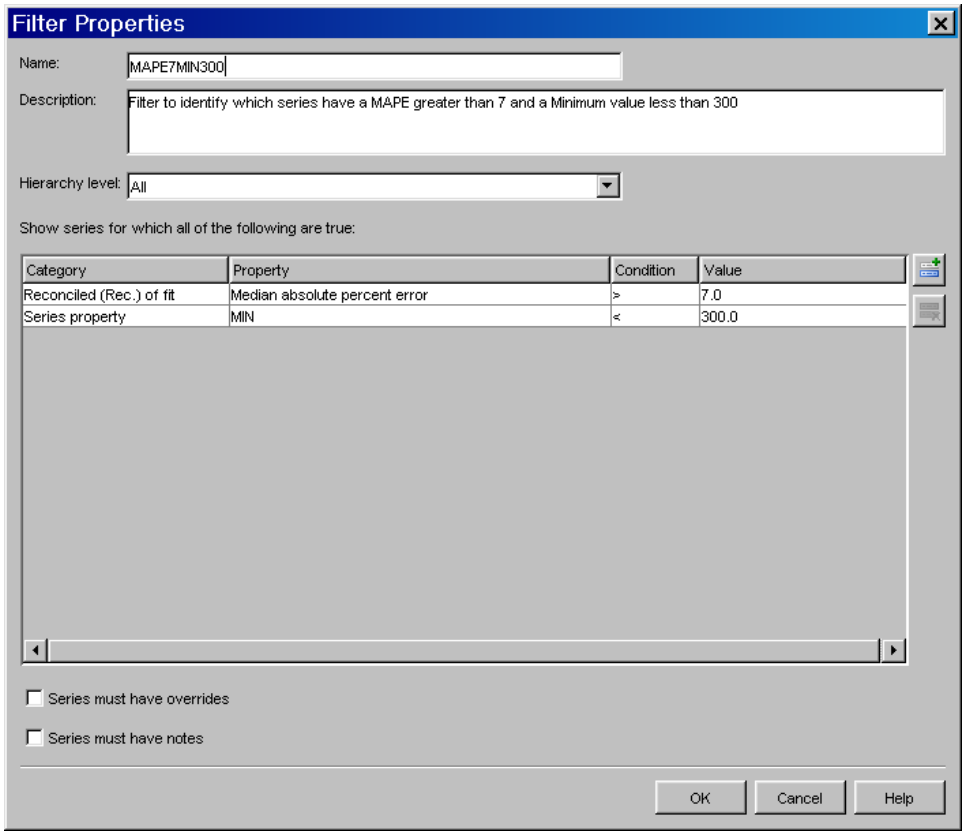

9. Click **OK** to add the filter to the list in the Filters dialog box.
10. Click **Close** in the Filters dialog box to return to the project.

The filter that you created is now available from the filtered table in the navigation panel. From the **Filter** drop-down list, select **MAPE7MIN300** to view the series that meet these filter criteria.

Figure 6.5 Results in the Filtered Table after Applying the MAPE7MIN300 Filter

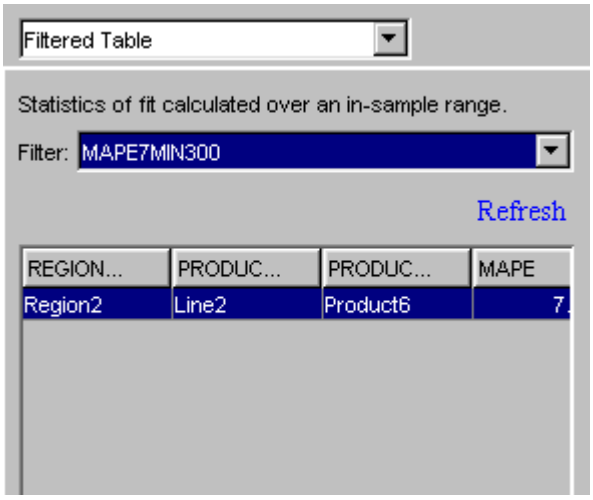

Chapter 7

Working with Events

Contents

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 create and apply an event to a time series, SAS Forecast Studio creates an indicator variable (also called a dummy variable) based on the specified event type. The indicator variable indicates the occurrence of the event at any time period. This indicator variable is used as a regressor variable for time series modeling and forecasting.

For example, daily retail sales data follows a fairly steady pattern depending only on the day of the week except in the case of a special event such as a holiday promotion. When you include events such as a one-time New Year's Day 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 New Year's Day, 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 Thanksgiving.

NOTE: Before the forecasting model can predict the effects of a future event, SAS Forecast Studio must have data on the effects of past occurrences of the event. For example, if you create an event for next Christmas, SAS Forecast Studio must have data that explains how retail sales were impacted

by previous Christmas events. Without this data, SAS Forecast Studio cannot predict the effect of the next Christmas event on retail sales.

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:

Event Type	Shape of the Event	Description
pulse	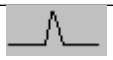	Temporary change in the magnitude of a time series process. The magnitude returns to the former level immediately after the change.
level shift	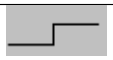	Persistent change in the level of a time series process.
ramp		A sudden growth or decay in the trend or slope of a time series process. This change in growth could persist or end when the time series reaches a new level.
temporary change	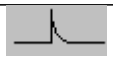	Temporary change in the magnitude. The magnitude decays to the former level after the change.

Managing Events

The Events dialog box enables you to 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 assign variables to roles in the New Project Wizard. For more information, see “Assign Data to Roles” on page 44.
- After you have created a project, you can manage your events by selecting **Project** → **Events**.

Figure 7.1 Events Dialog Box

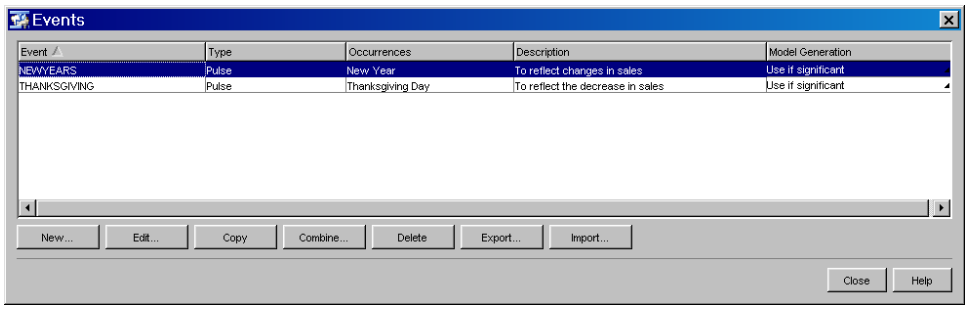

The Events dialog box lists the events that you have created. For each event, you can view the event type, occurrences, and description. You can also see whether the event is included in the generation of the models.

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
- export all the events
- import one or more events
- specify how the event is used when SAS Forecast Studio generates the models

Create an Event

To create a new event, complete the following steps:

1. Select **Project** → **Events**. The Events dialog box opens.
2. Click **New**. The Event Properties dialog box opens.

Figure 7.2 Event Properties Dialog Box

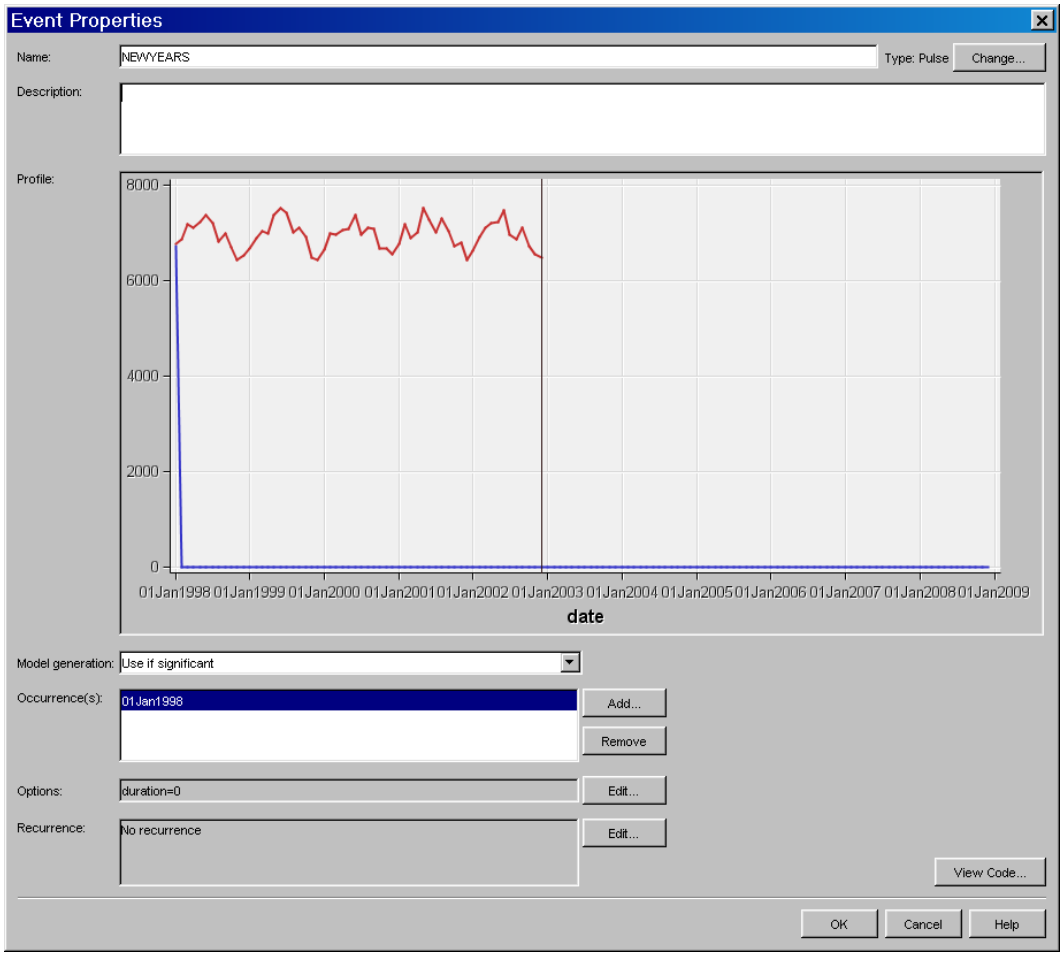

3. Specify a valid SAS name for the event. By default, the name for a new event is EVENT*n* where *n* is the lowest available integer value. The event name can be 32 characters long, and it must start with a letter (A-Z). Subsequent characters can be letters or numeric digits (0-9). Both upper- and lowercase letters are valid. For more information about SAS naming conventions, see *SAS Language Reference: Concepts*.

4. (Optional) Specify a description for the event. The description is limited to 256 characters.

5. Click **Change** to specify an event type. From the Event Type dialog box, you can choose from the following event types:

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

   For more information about these event types, see "What Is an Event?" on page 71.

6. Specify whether an event is used for all of the models that SAS Forecast Studio generates. From the **Model Generation** drop-down list, you can select from the following options:

- **Do not use** - does not include the event in the model.
- **Try to use** - specifies to include the event in the model if feasible and if the resulting model is stable and an improvement over the simple model.
- **Use if significant** - specifies to include the event in the model if feasible, if the resulting model is stable, an improvement over the simple model, and if the event is significant.
- **Force use** - specifies to include the event in the model if feasible.

7. Click **Add** to add an occurrence. From the Select Occurrences dialog box, select the time period for the occurrence and click **OK**. The **Occurrence** box in the Event Properties dialog box now displays a list of the occurrences for the event.

   NOTE: The options that are available depend on the timestamp of your data.

   To remove the occurrence, select the occurrence name from the list and click **Remove**.

8. To specify the options for the occurrence, click **Edit** next to the **Options** box. From the Event Options dialog box, you can specify the duration of the occurrences and the time of the occurrences.

9. To specify the recurrence, click **Edit** next to the **Recurrence** box. From the Event Recurrence dialog box, you can specify how frequently the event will recur and how long the recurrence will last.

   NOTE: Before you can specify a recurrence, you must specify an occurrence. However, you cannot specify a recurrence if you have selected two or more occurrences or if you have selected a holiday for the event occurrence.

10. (Optional) Click **View Code** to view the SAS code for the event.

11. Click **OK** to save the event. The new event appears in the Events dialog box.

Edit an Event

To edit an event, complete the following steps:

1. Select **Project** → **Events**. The Events dialog box opens.

2. Select the event from the list and click **Edit**. The Event Properties dialog box opens. For more information about the available options, see "Create an Event" on page 73.

   NOTE: After you have created an event, you cannot change the name of the event from the Event Properties dialog box. If you want to rename an event, you must create a copy of the event and specify a new name. After you have created this copy, you might want to delete the original event. For more information, see "Copy Events" on page 76.

3. Click **OK** to save your changes.

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 from the 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.

After you have copied an event, the Event Properties dialog box automatically opens. You can edit the name of a copied event only at the time that you copy an event. For more information, see "Edit an Event" on page 75.

Import Events

You might want to use events that are saved in a repository that is external to SAS Forecast Studio. You can do this by importing the event.

To import an event, complete the following steps:

1. Select **Project** → **Events**. The Events dialog box opens.
2. Click **Import**. The Import Events dialog box opens.
3. Select the data set that contains the events that you want to import into your project.

   If an event with the same name already exists in your project, you are prompted whether to override the event with the same name.
4. Click **OK**.

The imported event appears in the list of events for the project.

Export All the Events

By exporting events, you can share events with other SAS Forecast Studio users at your site. Exporting events enables you to create a SAS data set that contains event data and is saved to an events

repository. Other SAS Forecast Studio users can then import these events into projects that are running on the same SAS Workspace Server as the events repository.

To export an event, complete the following steps:

1. Select **Project** → **Events**. The Events dialog box opens.
2. Click **Export**. The Save As dialog box opens.
3. To view the libraries on a server, double-click the name of the server.
4. Select the library where you want to save the SAS data set that contains the event data. You must have write permission to save the events data set to the library that you select. You cannot save to the Sashelp or Sasuser libraries.

   NOTE: To return to a previous level, click  .
5. Specify a name for the data set in the **Data set name** text box.
6. Click **Save**.

NOTE: When you export an event, two SAS data sets are created. The first data set (called *data-set-name*) contains the variables for your event. The second data set (called *data-set-name*2) contains the information about whether the event should be included in the model generation. Only the *data-set-name* data set is visible from SAS Forecast Studio when you import an event. However, both data sets are automatically imported when you select the *data-set-name* data set. For more information, see "Import Events" on page 76.

Combine Events

To combine events, complete the following steps:

1. Select **Project** → **Events**. The Events dialog box opens.
2. Select the event(s) that you want to include in a combined event, and click **Combine**. The Combination Event Properties dialog box opens.

Figure 7.3 New Combination Event Dialog Box

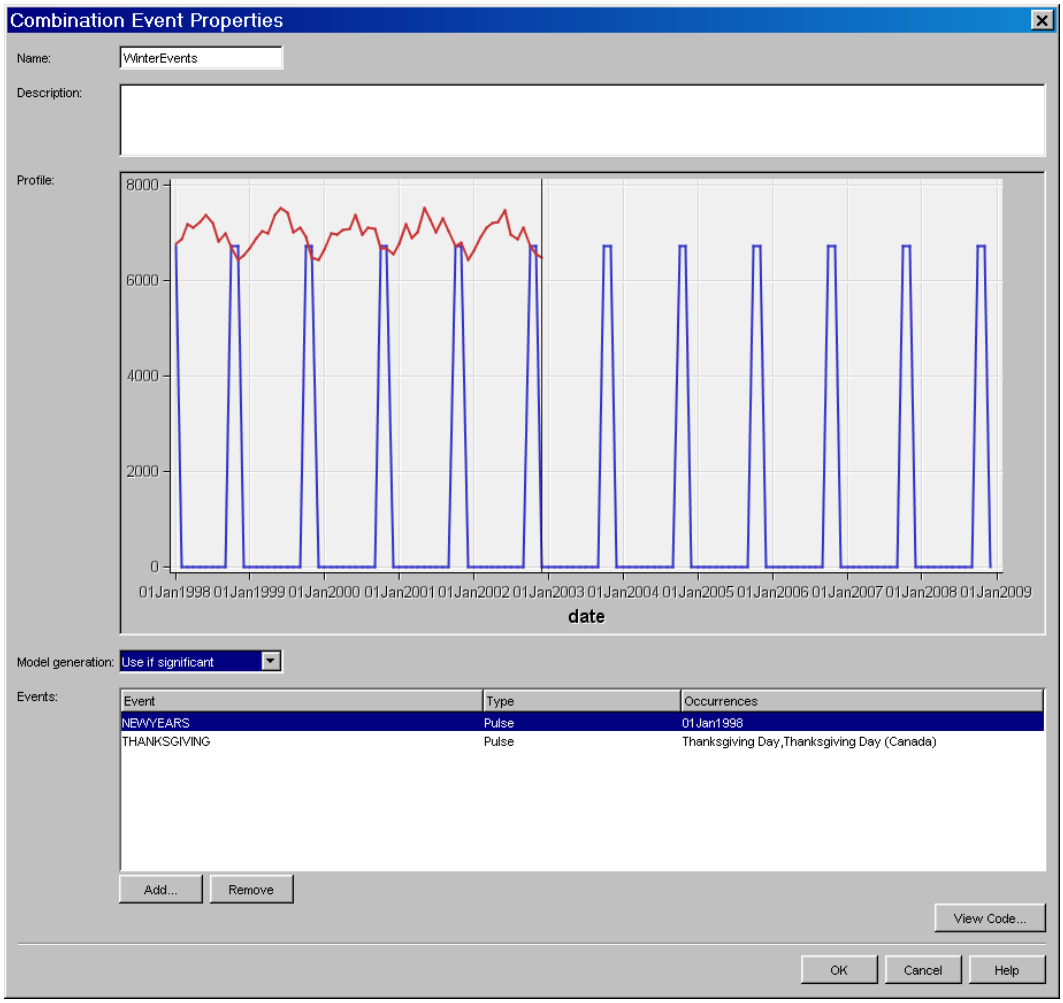

3. Specify a name for the new combined event.
4. (Optional) Provide a brief description for the new event.
5. Specify the events to include in the combination event. By default, the **Events** box lists the events that you selected in the Events dialog box. These events are used to create the combination event.
   - To add an event, click **Add**. From the Select Event dialog box, you can select the event(s) that you want to add to this list.
   - To delete an event, select the event from the list and click **Remove**.
6. (Optional) Click **View Code** to view the SAS code for the combined event.
7. Click **OK**. The new combined event appears in the Events dialog box.

To edit a combined event, select the combined event from the **Events** list and click **Edit**.

Specify How the Event Is Used in Model Generation

When you create an event, you specify whether to include the event when SAS Forecast Studio generates the models. You can change the status of the event or multiple events from the **Model Generation** column in the Events dialog box.

The status of the event that appears in the Events dialog box is used for all the models that SAS Forecast Studio generates. To include an event in the model specification for an individual model, see "Include Events in User-Defined Models" on page 183.

To specify the status of an event in model generation, complete the following steps:

1. Select **Project** → **Events**. The Events dialog box opens.
2. Select the event(s) whose status you want to change.
3. From the drop-down list in the **Model Generation** column, select one of the following options:
   - **Do not use** - does not include the event in the model.
   - **Try to use** - specifies to include the event in the model if feasible and if the resulting model is stable and an improvement over the simple model.
   - **Use if significant** - specifies to include the event in the model if feasible, if the resulting model is stable, an improvement over the simple model, and if the event is significant.
   - **Force use** - specifies to include the event in the model if feasible.

   NOTE: If you selected multiple events, then specifying the value of the **Model Generation** column for one event changes the value for all the selected events.

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 from the list.
3. Click **Delete**. You are prompted to confirm the deletion. Click **Yes** if you want to delete this event. Click **No** to keep the event.

Chapter 8

Generating Forecasts

Contents

Understanding Forecasting Concepts

Understanding SAS Time Intervals

You specify the time ID variable when you create the project using the New Project Wizard. After the project has been created, you cannot change the time ID variable. For more information, see "Assign Data to Roles" on page 44.

SAS Forecast Studio analyzes the variable assigned to the time ID role to detect the time interval of the data. SAS assumes that all of the values in the time ID variable are either date or date-time values and distinguishes between the values by their magnitude. This assumption fails if you have dates extending beyond July 21, 2196, or date-times before January 1, 1960.

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 (timestamped 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 length if there are numerous gaps (missing values) in the data. In this case, SAS Forecast Studio supplies the missing values. A validation routine checks the values of the time ID to determine whether they are spaced according to the interval that you specified.

Time intervals are specified in SAS by using character strings. Each of these strings is formed according to a set of rules that enables you to create an almost infinite set of attributes. For each time interval, you can specify the type (such as monthly or weekly), a multiplier, and a shift. You can specify a greater time interval than that found in the input data. A smaller interval should not be used, because a small interval will generate a large number of observations.

Understanding Aggregation and Accumulation

Difference between Aggregation and Accumulation

Aggregation is the process of combining more than one time series to form a single series. Aggregation combines data within the same time interval. For example, you can aggregate data into a total or average. In the New Project Wizard, you specify the aggregation options if you selected to forecast hierarchically.

The following examples explain when you might want to use an aggregation method.

- Your data set contains the sales for a group of products. If you want to know the total sales for a category, you would choose **Sum of values** as the aggregation method.
- Your data contains the price of each product. If you want to know the average price for a product line, you would choose **Average of values** as the aggregation method.

Accumulation can be either of the following:

- the process of converting a time series that has no fixed interval into a time series that has a fixed interval (such as hourly or monthly)
- the process of converting a time series that has a fixed interval into a time series with a lower frequency time interval (such as hourly into daily)

Accumulation combines data within the same time interval into a summary value for that time period.

Because this process is not dependent on whether you have a hierarchy, you might need to accumulate data regardless of whether you forecast your data hierarchically.

The Accumulation and Aggregation Methods

NOTE: The difference between accumulation and aggregation is the dimension along which each method is applied. The following equations focus on accumulation, but the same equations also apply to the aggregation methods that you specify when you configure the hierarchy.

Let $\mathbf{R} = \{\mathbf{r}_q\}_{q=1}^{Q}$ be the data vector ordered by the time series occurrence in the data set with respect to the observation index. Let $\mathbf{q} = 1, \ldots, \mathbf{Q}$ be the index that represents this ordering. Let $\mathbf{Q_N}$ be the number of non-missing values and let $\mathbf{Q}_{NMISS} = \mathbf{Q} - \mathbf{Q_N}$ be the number of non-missing values in the data vector. Let $\bar{r} = \frac{1}{\mathbf{Q_N}} \sum_{\mathbf{q}=1}^{\mathbf{Q}} \mathbf{r_q}$ be the average value of the data vector with the missing values ignored.

The following example accumulates the observation series, $\mathbf{Z^{(N)}} = \{\mathbf{z}_i\}_{i=1}^{N}$ to the time series, $\mathbf{Y^{(T)}} = \{\mathbf{y}_t\}_{t=1}^{T}$, $\mathbf{y_t} = Accumulate(\mathbf{Z_t^{(T)}})$ for $\mathbf{t} = 1, \ldots, \mathbf{T}$. In this situation, $\mathbf{R} = \mathbf{Z_t^{(T)}}$ and $\mathbf{Q} = \mathbf{N_t^{(T)}}$ for $\mathbf{t} = 1, \ldots, \mathbf{T}$.

Let $\mathbf{a} = Accumulate(\mathbf{R})$ be this accumulated value for this data vector when the following accumulation methods are applied:

None
does not accumulate the vector values.

Sum
accumulates the vector values based on the summation of their values.

$$\mathbf{a} = \sum_{\mathbf{q}=1}^{\mathbf{Q}} \mathbf{r_q}$$

Missing values are ignored in the summation. If $\mathbf{Q_N} = 0$, then $\mathbf{a}$ is set to missing.

Average
accumulates the vector values based on the average of their values.

$$\mathbf{a} = \bar{r} = \frac{1}{\mathbf{Q_N}} \sum_{\mathbf{q}=1}^{\mathbf{Q}} \mathbf{r_q}$$

Missing values are ignored in the summation. If $\mathbf{Q_N} = 0$, then $\mathbf{a}$ is set to missing.

Minimum
accumulates the vector values based on the minimum of their values.

$$\mathbf{a} = min(\{\mathbf{r}_q\}_{q=1}^{Q})$$

Missing values are ignored in the summation. If $\mathbf{Q_N} = 0$, then $\mathbf{a}$ is set to missing.

Maximum
accumulates the vector values based on the maximum of their values.

$$\mathbf{a} = max(\{\mathbf{r}_q\}_{q=1}^{Q})$$

Missing values are ignored in the summation. If $\mathbf{Q_N} = 0$, then $\mathbf{a}$ is set to missing.

Median
accumulates the vector values based on the median of their values.

$$\mathbf{a} = median(\{\mathbf{r}_q\}_{q=1}^{Q})$$

Missing values are ignored in the summation. If $\mathbf{Q_N} = 0$, then $\mathbf{a}$ is set to missing.

Number of Non-Missing Observations
accumulates the vector values based on the number of non-missing values.

$$\mathbf{a} = \mathbf{Q_N}$$

Number of Observations
accumulates the vector values based on the number of values.

$$\mathbf{a} = \mathbf{Q}$$

Number of Missing Observations
accumulates the vector values based on the number of missing values.

$$\mathbf{a} = \mathbf{Q}_{NMISS}$$

First Occurrence
accumulates the vector values based on the first observation in the data.

$$\mathbf{a} = \mathbf{r}_1$$

Last Occurrence
accumulates the vector values based on the last observation in the data.

$$\mathbf{a} = \mathbf{r_Q}$$

Standard Deviation
accumulates the vector values based on their standard deviation.

$$\mathbf{a} = \sqrt{\frac{1}{\mathbf{Q_N}-1} \sum_{\mathbf{q}=1}^{\mathbf{Q}} (\mathbf{r_q} - \bar{r})^2}$$

Missing values are ignored in the summation. If $\mathbf{Q_N} \leq 1$, then **a** is set to missing.

Uncorrected Sum of Squares
accumulates the vector values based on their uncorrected sum of squares.

$$\mathbf{a} = \sum_{\mathbf{q}=1}^{\mathbf{Q}} (\mathbf{r_q})^2$$

Missing values are ignored in the summation. If $\mathbf{Q_N} = 0$, then **a** is set to missing.

Corrected Sum of Squares
accumulates the vector values based on their corrected sum of squares.

$$\mathbf{a} = \sum_{\mathbf{q}=1}^{\mathbf{Q}} (\mathbf{r_q} - \bar{r})^2$$

Missing values are ignored in the summation. If $\mathbf{Q_N} = 0$, then **a** is set to missing.

Working with the Forecasting Settings for the Project

About the Forecasting Settings

All forecast settings are initially set to a default value.

To view the forecasting options for your project, select **Project → Forecasting Settings**.

In the Forecasting Settings dialog box, the forecasting options are divided into the following categories:

- Time ID
- Data Preparation
- Diagnostics
- Models
- Forecasts

To save your changes and exit the Forecasting Settings dialog box, click **OK**. If you have made changes, SAS Forecast Studio automatically reforecasts your data.

NOTE: To reset the options to their default values, click **Reset**. Any changes that you have made are lost.

Time ID Options

When you assign a variable to the time ID role, SAS Forecast Studio tries to detect a time interval and seasonal cycle length automatically. For example, if the time ID variable contains data values that are spaced one month apart, then SAS Forecast Studio uses an interval of month, and a seasonal cycle length corresponding to the default seasonal cycle length.

From the Time ID panel, you can specify the following options for the time ID:

Interval
specifies the interval of the values for the time ID variable. The available options for the interval include the following: **Day**, **Hour**, **Minute**, **Month**, **Quarter**, **Second**, **Semimonth**, **Semiyear**, **Ten-day**, **Week**, **Weekday**, and **Year**. If you select **Weekday** as the interval, then you can select which days are the weekend or inactive days in the week by clicking **Weekend**.

NOTE: If you change the time interval after adding overrides to your project, then the overrides are removed.

Multiplier
specifies the multiplier for the value specified as the interval. For example, if the interval is weekly and you specify 2 as the multiplier, then WEEK2 specifies two-week intervals.

Shift
specifies the offset for the interval. The units of the shift depend on the interval.

The following table shows how the values that you specify for the interval, seasonal cycle length, multiplier, and shift work together.

Interval Name (in SAS code format)	**Units for the Shift**	**Maximum Shift Value**	**Example**
YEARm.s	MONTH	12 x multiplier	YEAR2.7 specifies an interval of biennial years starting in July of even years.
SEMIYEARm.s	MONTH	6 x multiplier	SEMIYEAR3.2 specifies periods of 1 1/2 years starting alternately in July and January.
QTRm.s	MONTH	4 x multiplier	QTR.1 specifies quarterly periods starting in February.
SEMIMONTHm.s	SEMIMONTH	multiplier	
MONTHm.s	MONTH	multiplier	
TENDAYm.s	TENDAY	multiplier	
WEEKm.s	DAY	7 x multiplier	WEEK6.3 specifies six-week intervals starting on Tuesdays.
WEEKDAYm.s	DAY	multiplier x number of weekdays in a week	WEEKDAY3.2 specifies three-weekday intervals with the cycle three-weekday periods aligned to Monday, January 4, 1960.
DAYm.s	DAY	multiplier	
HOURm.s	HOUR	multiplier	HOUR8.7 specifies eight-hour intervals starting at 6:00am, 2:00pm, and 10:00pm.
MINUTE	MINUTE	multiplier	
SECOND	SECOND	multiplier	

Seasonal Cycle Length

specifies the length of a season. This value is populated automatically if SAS Forecast Studio can determine the seasonal cycle length from the time ID variable.

NOTE: 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.

Format

specifies the date format of the values in the time ID variable. This value is populated automatically if SAS Forecast Studio can determine the date format from the time ID variable.

NOTE: The WEEKU*w*., WEEKV*w*., and WEEKW*w*. formats are not supported.

To change the time ID format, click **Edit**. From the Edit Time ID Format dialog box, you can use the format that SAS Forecast Studio selected, select a predefined format, or specify a custom format. Custom formats must be saved in the format library that is defined for the

SAS session that is used by SAS Forecast Studio. For help creating a custom format that can be used in SAS Forecast Studio, contact your system administrator.

NOTE: If you change the time ID format, then SAS Forecast Studio performs a nondestructive diagnosis. SAS Forecast Studio preserves any user-created models. Any system-generated models are re-created.

Data Preparation Options

In the Data Preparation panel, you can specify the following options:

Select how to interpret embedded missing values
specifies how to replace the missing values in the data. 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 non-missing value.
- **Last** - missing values are set to the accumulated last non-missing value.
- **Previous** - missing values are set to the previous accumulated non-missing value. Missing values at the beginning of the accumulated series remain missing.
- **Next** - missing values are set to the next accumulated non-missing values. Missing values at the end of the accumulated series remain missing.

NOTE: When you change the value of this option, then SAS Forecast Studio performs a nondestructive diagnosis. SAS Forecast Studio preserves any user-created models. Any system-generated models are re-created.

Select which leading/trailing missing values to remove
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.

NOTE: When you change the value of this option, then SAS Forecast Studio performs a nondestructive diagnosis. SAS Forecast Studio preserves any user-created models. Any system-generated models are re-created.

Select which leading/trailing zero values to interpret as missing
specifies how beginning or ending zero values (or both) are interpreted in the accumulated time series. You can choose from the following options:

- **None** - the beginning or ending zeros (or both) 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.

NOTE: When you change the value of this option, then SAS Forecast Studio performs a nondestructive diagnosis. SAS Forecast Studio preserves any user-created models. Any system-generated models are re-created.

Ignore data points earlier than the following date
specifies the first date to use in the data preparation. 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.

NOTE: You cannot specify a start date if the time interval for your data is by hour, by minute, or by second. By default, SAS Forecast Studio uses the earliest date in the data set.

Accumulation
opens the Accumulation dialog box. From this dialog box, you can do the following:

- Set the aggregation and accumulation method to the same value for each variable.
- Specify the method for accumulating data in the case of more than one observation per time period.

  The default aggregation method for each dependent, independent, and reporting variable depends on how you forecasted your data.

  – If you forecasted your data hierarchically, then the accumulation method depends on the aggregation method that you selected.
  – If you did not forecast your data hierarchically, then the default accumulation method is **Sum of values**.

  For more information about each of the accumulation methods, see "The Accumulation and Aggregation Methods" on page 83.

  NOTE: If you change the accumulation method for a variable, then SAS Forecast Studio performs a nondestructive diagnosis. SAS Forecast Studio preserves any user-created models. Any system-generated models are re-created.
- Specify the method for aggregating data across levels of the hierarchy. This functionality is available only if you forecasted your data hierarchically.

Diagnostics Options

In the Diagnostics panel, you can specify the following options:

Perform intermittency test
specifies a number greater than 1 that is used to determine whether a time series is intermittent. This option is selected by default, and the default value is 2. To specify a different sensitivity, type a new value in the **Sensitivity** box.

Perform 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. To specify a different sensitivity, type a new value in the **Sensitivity** box.

Minimum number of seasonal cycles for a seasonal model
specifies the minimum number of seasonal cycles for a seasonal model.

For more information, see the MINOBS= option for the HPFENGINE procedure in the *SAS High-Performance Forecasting User's Guide*.

Minimum number of observations for a trend model
specifies minimum number of observations for a trend model.

For more information, see the MINOBS= option for the HPFENGINE procedure in the *SAS High-Performance Forecasting User's Guide*.

Minimum number of observations for a non-mean model
specifies the minimum number of observations for a non-mean model.

For more information, see the MINOBS= option for the HPFENGINE procedure in the *SAS High-Performance Forecasting User's Guide*.

Functional transformation
specifies the transformation to apply to the dependent variable. This transformation applies to the system-generated models. You can select from the following transformations:

- no transformation. This is the default.
- automatic transformation. SAS Forecast Studio uses the model selection criteria to determine whether to apply a logarithmic transformation or no transformation.
- logarithmic transformation
- square-root transformation
- logistic transformation
- Box-Cox transformation. If you select this transformation, then you can specify a parameter value in the **Box-Cox parameter** box. This value should be an integer between -5 and 5. The default value is 0.

Forecast
specifies how to calculate the forecasts for the system-generated models. Forecasts can be based on the mean or median. By default, the mean value is used.

NOTE: If you selected **None** from the **Functional transformation** drop-down list, then this option is not available.

Outlier detection (ARIMA models only)
specifies the settings to use for outlier detection in the ARIMA models. Select the **Detect**

outliers check box to specify whether SAS Forecast Studio should detect outliers in the data when fitting an ARIMA model.

Then specify the following options:

- the significance level for the tests used to detect shifts in the level of the dependent series that are not accounted for by the previously estimated model. For the significance level, you can specify any value between 0 and 1.
- the maximum percentage of series that can be outliers

Models Options

From the Models panel, you can specify the types of models to create when generating model selection lists. You must select at least one model for SAS Forecast Studio to use when generating the model selection lists. By default, SAS Forecast Studio creates an ARIMA model and an exponential smoothing model.

In the Models panel, you can specify the following options:

Generated ARIMA models
includes the ARIMA models that SAS Forecast Studio generates. You can also choose to have SAS Forecast Studio detect outliers in the model. An ARIMA model is created by default.

Generated exponential smoothing models
includes the exponential smoothing models (ESM) that SAS Forecast Studio generates. An exponential smoothing model is created by default.

Generated unobserved components models
includes the unobserved component models (UCM) that SAS Forecast Studio generates.

Models from an external list
enables you to select a model selection list that contains the models that you want to use. If you select this option, then by default SAS Forecast Studio uses the SASHELP.HPFDFLT.TSFSELECT list. For more information about this model selection list, see Appendix D, "Default Model Selection Lists."

To select a different model selection list, click **Browse**. The Select Model Selection List dialog box opens, and you can select the model selection list that you want to use.

NOTE: After SAS Forecast Studio has generated forecasts using the model selection list that you selected, you cannot clear the **Models from an existing list** check box. SAS Forecast Studio will continue to use the external model selection list that you selected.

Only fit generated exponential smoothing models at the lowest level of the hierarchy
specifies whether to fit only the exponential smoothing models that SAS Forecast Studio generates at the lowest level of the hierarchy. You can specify the level by using the **Number of levels** drop-down list. SAS Forecast Studio fits only these exponential smoothing models at the level that you specify and any lower levels. For example, if you select `2` from the **Number of levels** drop-down list, then SAS Forecast Studio fits only exponential smoothing models at the first and second levels of the hierarchy.

Forecasts Options

In the Forecasts panel, you can specify the following options:

Number of periods to forecast (horizon)
specifies the number of periods into the future for which multistep forecasts are made. The larger the horizon value, the larger the prediction error variance at the end of the horizon. By default, the horizon is 12.

Use holdout sample for model selection
specifies the size of the holdout sample to be used for model selection. The holdout sample is a subset of actual time series ending at the last non-missing observation. The default value is 2.

In the **Maximum percentage of series that holdout sample can be** box, specify the size of the holdout sample as a percentage of the length of the time series. The default value is 5%.

For more information about using a holdout sample, see the HOLDOUT= and HOLDOUTPCT= options for the HPFDIAGNOSE procedure in the *SAS High-Performance Forecasting User's Guide*.

Calculate statistics of fit over an out-of-sample range
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.

Selection criterion
specifies the selection criterion for the model. You can select only one value from the drop-down list. For more information about the available selection criterion, see Appendix C, "Statistics of Fit."

Confidence limit
specifies the confidence level for the series. By default, this confidence level is 0.05, which is a 95% confidence limit.

Allow negative forecasts
specifies whether negative forecasts are allowed. If this option is selected, then the series is log transformed. This option is not selected by default.

NOTE: If you clear the **Allow negative forecasts** check box, then the project is reforecast, and any overrides with negative values are removed.

Reforecasting a Series or Project

Reforecast a Series

After you initially create the project, you might change some parameters or create a new model. After these changes, you might want to reforecast a series.

To reforecast the selected series, complete the following steps:

1. Select the series in the hierarchy or filtered table and select **Series** → **Reforecast**. You can also click  in the Model View. The Forecast Series dialog box opens.
2. Specify how to forecast the series. The following options are available:

   - **Diagnose: create new generated models, refit all models, and select a model** - specifies that SAS Forecast Studio should create the forecasts by performing a non-destructive diagnose.
   - **Select: refit all models, and select a model** - specifies that SAS Forecast Studio performs the model selection, estimates the parameters of the selected model, and generates the forecasts. This is the default.

   NOTE: If you choose any option other than **Diagnose: create new generated models, refit all models, and select a model**, then any new series are forecast using the model selection list for the Best Smoothing Model, which ships with SAS Forecast Studio. For more information, see Appendix D, "Default Model Selection Lists."
3. (Optional) Select the **Reset all series to "best fitting model (automatic selection)"** check box.

   When the project is created, SAS Forecast Studio automatically selects the best-fitting model as the forecast model for each series. However, using the **Selected model** drop-down list in the Model View, you can change the forecast model for a series.

   This option is available only when you have selected a model other than the best-fitting model from the **Selected model** drop-down list. If you select this option, then SAS Forecast Studio selects the best-fitting model and ignores any models that you might have selected for a series. For more information about how you select a model, see "Specifying the Forecast Model" on page 195.
4. Click **OK** to reforecast the series.

Reforecast a Project

After you initially create the project, you might change some parameters or create a new model. After these changes, you might want to reforecast the project.

To reforecast all the series in a project, complete the following steps:

1. Select **Project** → **Reforecast** or click  in the toolbar. The Forecast Project dialog box opens.
2. Specify how to forecast the project. The following options are available:

- **Diagnose: create new generated models, refit all models, and select a model** - specifies that SAS Forecast Studio should create the forecasts by performing a non-destructive diagnose.
- **Select: refit all models, and select a model** - specifies that SAS Forecast Studio performs the model selection, estimates the parameters of the selected model, and generates the forecasts. This is the default.
- **Fit: refresh the current forecast model, updating the parameter values** - specifies that SAS Forecast Studio estimates the parameters using the current model. No model selection is performed.
- **Forecast: refresh the current forecast model, using the same parameter values** - specifies that SAS Forecast Studio reforecasts using the current models and parameter estimates. No parameter estimation occurs.

NOTE: If you choose any option other than **Diagnose: create new generated models, refit all models, and select a model**, then any new series are forecast using the model selection list for the Best Smoothing Model, which ships with SAS Forecast Studio. For more information, see Appendix D, "Default Model Selection Lists."

3. (Optional) For the diagnose and select options, select the **Reset all series to "best fitting model (automatic selection)"** check box.

   When the project is created, SAS Forecast Studio automatically selects the best-fitting model as the forecast model for each series. However, using the **Selected model** drop-down list in the Model View, you can change the forecast model for a series.

   If you select this option, then SAS Forecast Studio selects the best-fitting model and ignores any models that you might have selected for a series. For more information about how you select a model, see "Specifying the Forecast Model" on page 195.

4. Select the **Use updated data if available** check box. SAS Forecast Studio determines whether the data source has been updated, and if so, updates the project data and reforecasts the existing series based on the option that you selected.

5. Click **OK** to reforecast the project. After the reforecasting is complete, the Forecast Summary dialog box opens. For more information, see "Reviewing a Summary of the Forecast Run" on page 93.

Reviewing a Summary of the Forecast Run

After SAS Forecast Studio generates the forecasts for the project, the Forecast Summary dialog box appears. You can also open this dialog box by selecting **Project → Forecast Summary**.

From the forecast summary, you can view the following information for each level in the hierarchy or for the single level in the project if you did not forecast your data hierarchically:

- a graph of the distribution of the statistic of fit

- a graph of the different types of model families. This graph shows whether the model included a dependent transformation, a seasonal component, a trend component, inputs (such as independent variables), events, and outliers.
- a graph of the model type. This graph shows how each model type (ARIMA, Exponential Smoothing, Intermittent Demand, and Unobserved Components) fits the series.
- a table of the number of errors. This table shows the number of failed forecasts, the number of series where SAS Forecast Studio used the best-fitting model, the number of warnings in the SAS log, and the number of errors in the logs.

To view more information about a graph or table, double-click the graph or table to open it in a separate window.

Here is an example of the Forecast Summary dialog box:

Figure 8.1 Forecast Summary Dialog Box

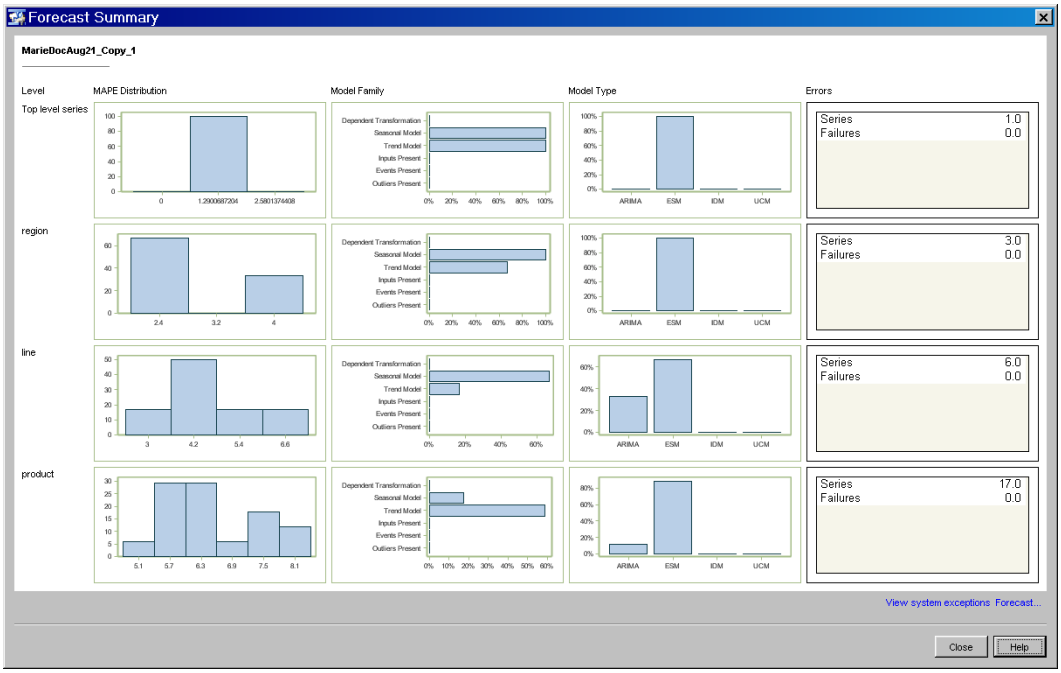

Specifying an Inactive Series

What Is an Inactive Series?

When no forecast is produced for the series, the series is inactive.

When you specify a series as inactive, then the following occurs:

- Any child series must also be inactive. SAS Forecast Studio prompts you to confirm making these child series inactive when you clear the **Active** check box.
- Any overrides that you have specified for this series are removed, and you cannot add new overrides to an inactive series.
- A "Reconciliation is out of date" message appears above the Forecasting View. You can click **Update** to reconcile the hierarchy.
- In the Model View, no forecast model is selected.

NOTE: When SAS Forecast Studio re-creates the model selection list, any inactive series become active again.

Set a Series as Inactive

To specify an inactive series, complete the following steps:

1. In the hierarchy or filtered table, select the series that you want to make inactive.
2. Clear the **Active** check box in either of the following views:
   - in the data table of the Forecasting View
   - in the model selection list of the Model View

   NOTE: If there are child series for the series that you selected, then SAS Forecast Studio prompts you before making those child series inactive. Click **Yes** to continue.

To make the series (and any child series) active again, select the **Active** check box.

Chapter 9

Working with Overrides

Contents

Understanding Overrides

What Is an Override?

An override is a replacement value that you specify in place of a forecasted value. How the override is applied depends on whether you have forecasted your data hierarchically.

- If you have forecasted your data hierarchically, then the override is an adjustment that is done with respect to the reconciled statistical forecast.
- If you have not forecasted your data hierarchically, then the override is an adjustment that is done with respect to the statistical forecast.

You can specify overrides only for future values. Series that have an associated override are identified in the hierarchy with the Override  icon.

Two types of overrides exist in SAS Forecast Studio:

- Locked overrides. A locked override is a user-supplied value for a forecast that is honored when the hierarchy is reconciled. When you specify a locked override, SAS Forecast Studio changes the final forecast to the override value.
- Unlocked overrides. An unlocked override is a user-supplied value for a forecast that acts as a guideline for the final forecast value. SAS Forecast Studio might not honor this override value.

You might want to add a note to explain each override. For more information, see "Add a Note to a Series" on page 103.

When Can You Specify an Override?

You can specify overrides only for future values.

When you are specifying overrides, it is important to remember the following:

- You can enter overrides only for the time ID values in the horizon. The horizon is determined by the value that you specify for the **Number of periods to forecast (horizon)** option in the Forecasting Settings dialog box and the end date for the series. The end date for the series is the largest time ID with a non-missing value for the dependent variable.
- You cannot add overrides when the value for the **Number of periods to forecast (horizon)** option is less than the number of the **Calculate statistics of fit over an out-of-sample range**

option. You cannot edit the time ID values in the horizon from the data table, and the Override Calculator is not available if this condition is met.

If you specified any overrides before this condition was met, then these previous overrides are not removed. The override icons are still visible in the hierarchy. If you make any additional changes to the values of the forecasting options that remove this condition, then these overrides will become visible in the Forecasting View.

- You cannot add a negative override if negative forecasts are not allowed in the project. To allow negative forecasts, you must select the **Allow negative forecasts** check box in the Forecasts panel of the Forecasting Settings dialog box.

For more information about these options, see "Forecasts Options" on page 91.

How Are Overrides Processed?

How SAS Forecast Studio processes overrides depends on whether any override conflicts are detected and how SAS Forecast Studio resolves conflicts between locked overrides. When setting up a hierarchy, you can choose whether the conflicts are resolved before reconciliation or the conflicts are ignored. For more information about the advanced reconciliation settings, see "Understanding the Reconciliation Options" on page 108.

If you require that conflicts be resolved before reconciliation, then SAS Forecast Studio uses the following process:

1. When you add an override, a "Reconciliation is out of date" message appears at the top of the Forecasting View.
2. When you click **Update** to reconcile the hierarchy, SAS Forecast Studio checks for conflicts in the locked overrides.
   - If no conflicts are detected, SAS Forecast Studio performs an override reconciliation. During the reconciliation, one of the following events occurs:
     - The hierarchy is fully reconciled, and the reconciliation message in the Forecasting View disappears.
     - SAS Forecast Studio cannot reconcile the hierarchy. An "Unreconciled nodes in the hierarchy" message appears at the top of the Forecasting View. You can click **Report** to view the Reconciliation Failure Report. For more information, see "About the Reconciliation Failure Report" on page 109.
       NOTE: Unreconciled nodes occur only if you also selected the **Restrict to direction implied by reconciliation method** option as the direction for override reconciliation. If you selected the **No restriction** option instead, then this message does not appear. These options are available from the Advanced Reconciliation Settings dialog box.
   - If conflicts are detected, the Override Conflicts dialog box opens. You view and resolve any conflicts.

After all the conflicts are resolved, SAS Forecast Studio reconciles the hierarchy. During the reconciliation, one of the following events occurs:

- The hierarchy is fully reconciled, and the reconciliation message in the Forecasting View disappears.
- SAS Forecast Studio cannot reconcile the hierarchy. An "Unreconciled nodes in the hierarchy" message appears at the top of the Forecasting View. You can click **Report** to view the Reconciliation Failure Report. For more information, see "About the Reconciliation Failure Report" on page 109.
  NOTE: Unreconciled nodes occur only if you also selected the **Restrict to direction implied by reconciliation method** option as the direction for override reconciliation. If you selected the **No restriction** option instead, then this message does not appear. These options are available from the Advanced Reconciliation Settings dialog box.

If you choose to ignore the conflicts, then SAS Forecast Studio uses the following process:

1. When you add an override, a "Reconciliation is out of date" message appears at the top of the Forecasting View.
2. When you click **Update** to reconcile the hierarchy, SAS Forecast Studio tries to reconcile the hierarchy. During the reconciliation, one of the following events occurs:

   - The hierarchy is fully reconciled, and the reconciliation message in the Forecasting View disappears.
   - SAS Forecast Studio cannot reconcile the hierarchy. An "Unreconciled nodes in the hierarchy" message appears at the top of the Forecasting View. You can click **Report** to view the Reconciliation Failure Report. For more information, see "About the Reconciliation Failure Report" on page 109.

     The Reconciliation Failure Report lists the unreconciled nodes and any conflicts if they were detected but ignored because of the **Ignore conflicts** option. If you also selected the **No restriction** option, then resolving any override conflicts will eliminate the unreconciled nodes. However, if you selected the **Restrict to direction implied by reconciliation method** option, then resolving all the override conflicts might reduce, but not necessarily eliminate, the number of unreconciled nodes.

Adding Overrides

About Adding Overrides

You can add an override in the following ways:

- By typing the override value in the Overrides row of the data table. You can use this method when you want to override a single value. For more information, see "Add Overrides in the Forecasting View" on page 101.
- By using the Override Calculator. This method is good if you want to create overrides for several time periods. For more information, see "Add Overrides by Using the Override Calculator" on page 101.

NOTE: If you clear the **Allow negative forecasts** check box, then the project is reforecast, and any overrides with negative values are removed. For more information, see "Forecasts Options" on page 91.

For each override, you can specify whether the override is locked or unlocked. This locking determines how SAS Forecast Studio treats the override during the reconciliation process. For more information, see "Lock and Unlock Overrides" on page 102.

You might also want to add a note to the series to explain the overrides. For more information, see "Add a Note to a Series" on page 103.

Add Overrides in the Forecasting View

To add an override from the data table in the Forecasting View, complete the following steps:

1. In the Overrides row, double-click in the cell for the time period that you want to override.
2. Type the override value and press ENTER.

After you have added an override, you can edit it by double-clicking in the cell and then typing the new value.

NOTE: If you cannot add an override to a future value, see "When Can You Specify an Override?" on page 98.

Add Overrides by Using the Override Calculator

To create an override by using the Override Calculator, complete the following steps:

1. In the data table in the Forecasting View, select the cells in the Overrides row for the time periods that you want to override, and click  .

   NOTE: If the Override Calculator is not available, see "When Can You Specify an Override?" on page 98.
2. Specify how to calculate the overrides for the selected time periods. You can choose from the following options:

- Select the **Base on existing values** option to base the override values on the specified increase or decrease of the forecast.

  From the drop-down list, select the values that you want to use. You can choose from the following options:

  – **Reconciled forecast** - bases the override values on the specified increase or decrease of the reconciled statistical forecast. This option is available only if you forecasted your data hierarchically.
  – **Statistical forecast** - bases the override values on the specified increase or decrease of the statistical forecast. This option is available only if you did not forecast your data hierarchically or if no reconciliation forecast is available.
  – **Current overrides** - bases the override values on the specified increase or decrease of the current override.
  – **Final forecast** - bases the override values on the specified increase or decrease of the final forecast.

  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 **+** or **-**. In the text box, specify the value; in the second drop-down list, specify whether this value is in units or percent.

  For example, say you want to create an override that is 10% greater than the reconciled statistical forecast. In the first drop-down list, select **+**. In the text box, type `10`, and in the second drop-down list, select **%**. The override values are 10% greater than the reconciled statistical forecast.

- Select **Set to a new value** 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:

  – **Split proportional to the reconciled statistical forecast**
  If there is no reconciled forecast, then this option is **Split proportional to the statistical forecast**.
  – **Even split**
  – **Assign specified value to each period**

3. Click **Apply** to apply the override. Click **OK** to save your changes and close the Override Calculator.

Lock and Unlock Overrides

To lock an individual override, complete the following steps:

1. In the SAS Forecast Studio workspace, open the Forecasting View and scroll to the data table.
2. In the Override Locks row of the data table, select the check boxes for the overrides that you want to lock.

   To unlock an override, clear the check box.

To lock all of the overrides, click  in the toolbar of the data table. To unlock all of the overrides, click  .

If SAS Forecast Studio cannot resolve any override conflicts, then you can also unlock overrides from the Override Conflicts dialog box.

Add a Note to a Series

From the Forecasting View, you can add notes to a series. You might want to add a note when you specify an override for a forecast.

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. In the Forecasting View, expand the Notes section.
2. In the text box, type your note. This note cannot exceed 2,000 characters.

The note is automatically saved with the series when you select a different series in the hierarchy or filtered table.

Series that have an associated note are identified in the hierarchy with a Note  icon. The notes persist when you close the project.

Remove Overrides

You can remove overrides in the following ways:

- To remove individual overrides, in the Overrides row of the data table select the cell(s) of the overrides that you want to delete and click  .
- To remove the overrides for the project, select **Project** → **Remove All Overrides**. In the warning that appears, click **OK** to remove the overrides. Click **Cancel** to keep the current overrides.
- To remove the overrides for the selected series, select **Series** → **Remove All Overrides**. In the warning that appears, click **OK** to remove the overrides. Click **Cancel** to keep the current overrides.

Working with Override Conflicts

What Is an Override Conflict?

An override conflict is a condition that occurs when a locked override for a given time interval violates the limits implied by locked overrides at the same time interval but in lower levels of the hierarchy.

The following figure shows an example of this conflict.

Figure 9.1 Example of an Override Conflict

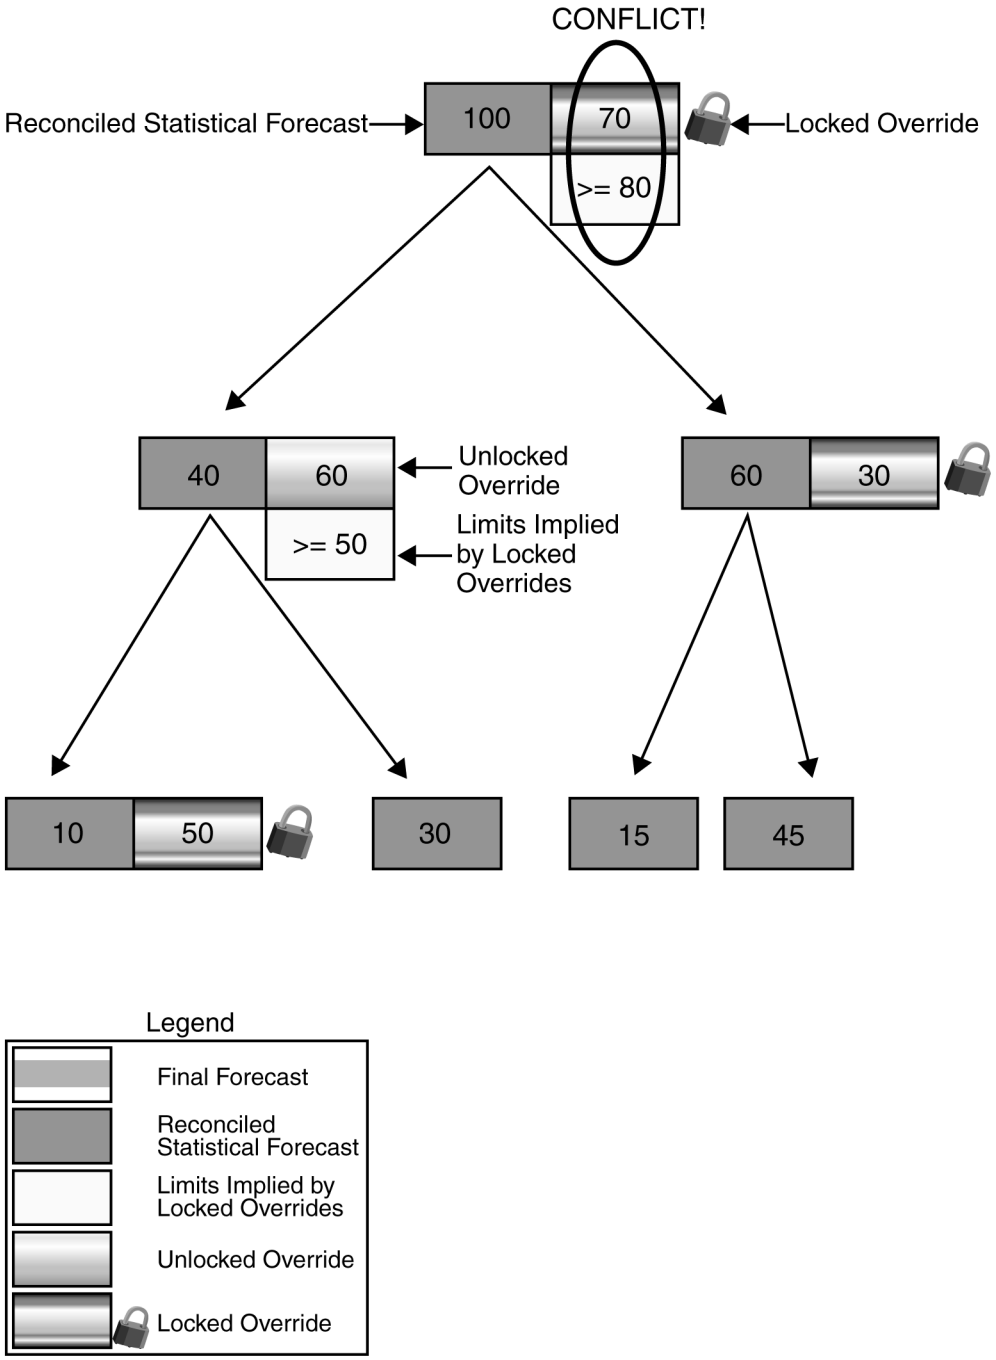

Resolve Override Conflicts

SAS Forecast Studio detects conflicts between locked overrides. When SAS Forecast Studio reconciles the hierarchy, any conflicts appear in the Override Conflicts dialog box.

To resolve conflicts, complete the following steps:

1. From the **Date** pane, select a time period. The right pane of the dialog box shows the conflict for the time period that you selected.
2. To resolve a conflict, you can do the following in the Parent and Children tables:
   - Clear the check boxes for the overrides that can be unlocked.
   - Change the override value in the table.

   NOTE: For more information about when you can specify an override, see "When Can You Specify an Override?" on page 98.
3. (Optional) Repeat steps 1 and 2 for all of the time periods in the **Date** pane.
4. Click **Reconcile**. SAS Forecast Studio attempts to reconcile the hierarchy. After this reconciliation has completed, a message appears. Click **Yes** to close this message.

Understanding Hierarchy Reconciliation

Reconcile the Hierarchy

If you add an override to data that you forecasted hierarchically, then you must reconcile the hierarchy in order for SAS Forecast Studio to calculate the final forecast.

You can reconcile the hierarchy in the following ways:

- Select **Project** → **Reconcile Hierarchy**.
- Click  in the toolbar.
- If the "Reconciliation is out of date" message appears in the Forecasting View, then you can click **Update**.

NOTE: SAS Forecast Studio might not be able to reconcile the hierarchy for a variety of reasons. How SAS Forecast Studio reconciles the hierarchy depends on the values that you set for the advanced reconciliation options. For more information, see "Specify the Advanced Reconciliation Options" on page 107.

Specify the Advanced Reconciliation Options

You might want to specify the reconciliation options for any overrides that you have specified.

To specify the reconciliation options for overrides, complete the following steps:

1. You can specify the reconciliation options in the following ways:
   - If you are creating a project, then you can specify these options from the Configure the Hierarchy step in the New Project Wizard. For more information, see "Configure the Hierarchy" on page 45.
   - If you have already created the project, then you can modify the hierarchy settings by selecting **Project** → **Hierarchy Settings**. The Hierarchy Settings dialog box opens. This dialog box has the same options that are available from the Configure the Hierarchy step in the New Project Wizard.
2. In the New Project Wizard or Hierarchy Settings dialog box, click **Advanced**. The Advanced Reconciliation Settings dialog box opens.
3. Select the **Use confidence interval widths in reconciliation magnitude calculations** check box to specify that the loss function for top-down reconciliation be weighted by the inverse of the variance of the input forecasts. For more information about this option, see the HPFRECONCILE procedure in the *SAS High-Performance Forecasting User's Guide*.
4. Specify how to resolve conflicts between locked overrides. You can choose from the following options:
   - **Require that conflicts be resolved before reconciliation** does not reconcile the hierarchy till all of the override conflicts are resolved.
   - **Ignore conflicts (unresolved conflicts will result in unreconciled nodes)** tries to reconcile the hierarchy without checking for the conflicts. If SAS Forecast Studio finds a conflict for a locked override, then the node that contains that override is not reconciled. This option is useful for reconciling overrides in batch mode.
5. Specify whether to restrict the direction of the reconciliation. You can choose from the following options:
   - **No restriction** reconciles all of the forecasts in the hierarchy regardless of the reconciliation method that you have specified. This option might involve an implicit top-down pass as part of a bottom-up reconciliation method, or it might involve an implicit bottom-up pass as part of a top-down reconciliation method. If you select this option, then the forecasts at the reconciliation level could change.
   - **Restrict to direction implied by reconciliation method (may result in unreconciled nodes)** can result in reconciliation failures in the hierarchy. The forecasts at the reconciliation level do not change unless they are explicitly overridden. Both locked and unlocked overrides at the reconciliation level are honored.

For more information about how these options work together and the results that you will see in SAS Forecast Studio, see "Understanding the Reconciliation Options" on page 108.

Understanding the Reconciliation Options

From the Advanced Reconciliation Settings dialog box, you can specify how you want to resolve conflicts between locked overrides and specify whether to restrict the direction of the reconciliation. Depending on the combination of options that you select, your results will be slightly different. For more information about how to set these options, see "Specify the Advanced Reconciliation Options" on page 107.

The following table explains the possible combinations that you can select and the events that result:

Resolve conflicts	Reconciliation direction	Result
Require that conflicts be resolved before reconciliation	**No restriction**	When you click **Update** in the Forecasting View, SAS Forecast Studio determines whether there are any conflicts. If conflicts exist, the Override Conflicts dialog box opens. In this dialog box, you have the following options: • You can close this dialog box, and the hierarchy remains unreconciled. • You can resolve any conflicts, and SAS Forecast Studio reconciles the hierarchy when you click **Reconcile**.
Ignore conflicts	**No restriction**	When you click **Reconcile** in the Forecasting View, SAS Forecast Studio determines whether there are any conflicts. However, because you selected the **Ignore conflicts** option, the Override Conflicts dialog box does not open if conflicts are detected. Instead the Reconciliation Failure Report opens. This report lists the override conflicts and the reconciliation failures. The relationship between the conflicts and the reconciliation failures is one-to-one. (For every conflict, there should be a reconciliation failure.) You might select these options if you want to reconcile the hierarchy as best as possible, but still allow a few conflicts.

Resolve conflicts	Reconciliation direction	Result
Require that conflicts be resolved before reconciliation	**Restrict direction**	When you click **Update** in the Forecasting View, SAS Forecast Studio determines if there are any conflicts. If conflicts exist, the Override Conflicts dialog box opens. In this dialog box, you have the following options: • You can close this dialog box, and the hierarchy remains unreconciled. • You can resolve any conflicts, and SAS Forecast Studio reconciles the hierarchy when you close this dialog box. However, there could be reconciliation failures even if the override conflicts are resolved. In general, an override conflict implies a reconciliation failure, but failures are not always due to override conflicts. Therefore, resolving override conflicts reduces the number of reconciliation failures, but might not eliminate all of the failures. If additional failures exist, then the Reconciliation Failure Report opens and shows the reconciliation failures.
Ignore conflicts	**Restrict direction**	When you click **Reconcile** in the Forecasting View, SAS Forecast Studio determines whether there are any conflicts. However, because you selected the **Ignore conflicts** option, the Override Conflicts dialog box does not open if conflicts are detected. Instead the Reconciliation Failure Report opens. This report lists the override conflicts and the reconciliation failures. The relationship between the conflicts and the reconciliation failures is not necessarily one-to-one. Additional reconciliation failures might have occurred. You might want to select these options if you want SAS Forecast Studio to reconcile the hierarchy with restrictions but without intervention.

About the Reconciliation Failure Report

If you have unreconciled nodes in the hierarchy, then a warning message appears at the top of the workspace. Click **View** to open the Reconciliation Failure Report.

NOTE: You can also open this report by selecting **Project → Reconciliation Failure Report**.

For more information about the advanced reconciliation options, see "Specify the Advanced Reconciliation Options" on page 107.

The Reconciliation Failures report shows the nodes in the hierarchy that could not be reconciled. The table contains the following information:

- the time period for the unreconciled node
- the name of the unreconciled node
- the value of the override
- the final forecast of the unreconciled node based on the final forecasts of the nodes in the lower levels of the hierarchy

If SAS Forecast Studio detected override conflicts before reconciling the hierarchy, then select **Resolve conflicts** to resolve these conflicts and reduce the number of reconciliation failures.

Examples of Reconciling a Hierarchy with Overrides

Overview of the Examples

The following examples show how SAS Forecast Studio reconciles the hierarchy based on the reconciliation options that you choose.

These examples assume that you have created a hierarchy for Region > Product Category > Product Line. The following table shows the possible values for each level:

Level in the hierarchy	Possible values
Region	Southeast
Product Category	Toys Clothes
Product Line	Trains Dolls Games Shirts Pants

All of the examples include the following assumptions:

- The data in these examples is for a single time period. This period is Dec2008.

- **Forecast Proportions** is selected as the disaggregation method. For more information about disaggregation methods, see "Configure the Hierarchy" on page 45.
- **Sum of Values** is selected as the aggregation method. For more information about aggregation methods, see "Configure the Hierarchy" on page 45.
- The **Allow negative forecasts** option is not checked in the **Forecasts** panel of the Forecasting Settings dialog box. For more information, see "Forecasts Options" on page 91.
- The **Use confidence interval widths in reconciliation magnitude calculations** option is not checked in the Advanced Reconciliation Settings dialog box. For more information, see "Specify the Advanced Reconciliation Options" on page 107.

Example 1: Top-Down Method of Reconciliation with No Conflicts

Purpose of the Example

This example shows the following:

- how SAS Forecast Studio reconciles a hierarchy using the top-down method of reconciliation
- how SAS Forecast Studio honors locked overrides
- how unlocked overrides that are not in the highest level of the hierarchy are used

Setup for the Example

In this example, you specified the following options when you set up the hierarchy:

- **Top Down** as the reconciliation method
- either **No restriction** or **Restrict to direction implied by reconciliation method** as the restriction on the reconciliation direction

For more information about these options, see "Modify the Hierarchy Settings" on page 53.

You also specified the following overrides for the Dec2008 time period:

- 200 as a locked override for the Southeast node in the Region level
- 60 as an unlocked override for the Toys node of the Product Category level
- 80 as a locked override for the Trains node in the Product Line level

NOTE: In addition to these settings, there are several assumptions about the data and option settings that are common to all the examples. To view these assumptions, see "Overview of the Examples" on page 110.

Explanation of the Reconciliation Process

The following figure shows how SAS Forecast Studio reconciles the hierarchy for the Dec2008 time interval:

Figure 9.2 Example 1: SAS Forecast Studio Honors All Locked Overrides

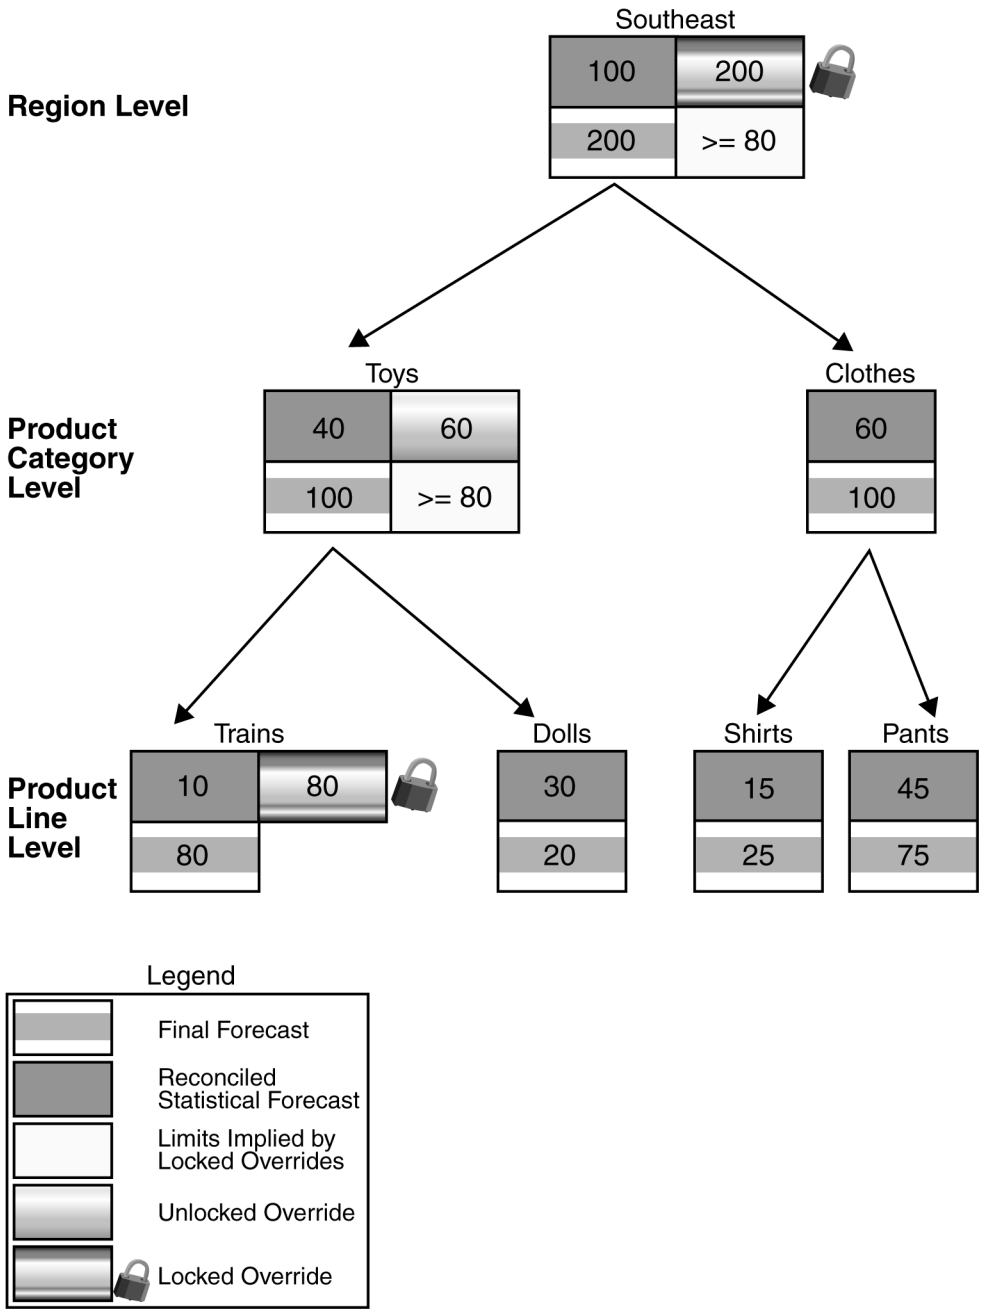

The following table explains the reconciliation process for each level of the hierarchy:

Level in the hierarchy	**Reconciliation process for the nodes in this level**
Region	Because you selected **Top Down** as the reconciliation method, SAS Forecast Studio starts reconciling the hierarchy at the highest level in the hierarchy. In this example, the highest level is the Region level. The reconciled statistical forecast for the Southeast node was 100; however, you specified 200 as a locked override for this level. SAS Forecast Studio honors this locked override, and consequently, the final forecast for the Southeast node is 200.
Product Category	In the Product Category level, there are 2 nodes—the Toys node and the Clothes node. For the Toys node, you specified an unlocked override of 60. SAS Forecast Studio does not honor the value of this unlocked override. Instead, SAS Forecast Studio uses the proportions to determine the final forecasts. The unlocked override that you specified for the Toys node is used instead of the reconciled statistical forecast to determine the final forecast. The unlocked override for the Toys node is 60, and the reconciled statistical forecast for the Clothes node is 60. Because these values are equal, the final forecast for the Southeast node (200) is divided equally between the 2 nodes, so the final forecast for both the Toys and Clothes nodes is 100.
Product Line	In the Product Line level, there are 4 nodes. The 2 nodes below the Toys node are the Trains and Dolls nodes. The 2 nodes below the Clothes node are the Shirts and Pants nodes. For the nodes under the Toys node, you specified 80 as a locked override for the Trains node. SAS Forecast Studio honors this value, so the final forecast for the Trains node is 80. Because the final forecast of the Toys node is 100, the final forecast of the Dolls node is set to 20. In the nodes under the Clothes node, you did not specify any overrides for either node. Therefore, SAS Forecast Studio uses the proportions of the reconciled statistical forecasts to determine the final forecast for each node. In this example, the reconciled statistical forecast for the Shirts node is 25% of the reconciled statistical forecast for the Clothes node, and the reconciled statistical forecast for the Pants node is 75% of the Clothes node. Therefore, the final forecast for the Shirts node is 25% of the final forecast for the Clothes node (or .25 x 100 = 25), and the final forecast for the Pants node is 75% of the final forecast for the Clothes node (or .75 x 100 = 75).

Example 2: Implicit Bottom-Up Method of Reconciliation

Purpose of the Example

This example shows the following:

- how SAS Forecast Studio reconciles a hierarchy using the top-down method of reconciliation
- how using the proportions that were calculated from an unlocked override could violate the limit that is implied by a locked override
- how SAS Forecast Studio uses an implicit bottom-up pass to reconcile the hierarchy

Setup for the Example

In this example, you specified the following options when you set up the hierarchy:

- **Top Down** as the reconciliation method
- **No restriction** as the restriction on the reconciliation direction

For more information about these options, see "Modify the Hierarchy Settings" on page 53.

You also specified the following overrides for the Dec2008 time period:

- 120 as a locked override for the Southeast node in the Region level
- 60 as an unlocked override for the Toys node of the Product Category level
- 80 as a locked override for the Trains node in the Product Line level

NOTE: In addition to these settings, there are several assumptions about the data and option settings that are common to all the examples. To view these assumptions, see "Overview of the Examples" on page 110.

Explanation of the Reconciliation Process

The following figure shows how SAS Forecast Studio reconciles the hierarchy for the Dec2008 time interval:

Figure 9.3 Example 2: Implicit Bottom-Up Method Used to Reconcile Hierarchy

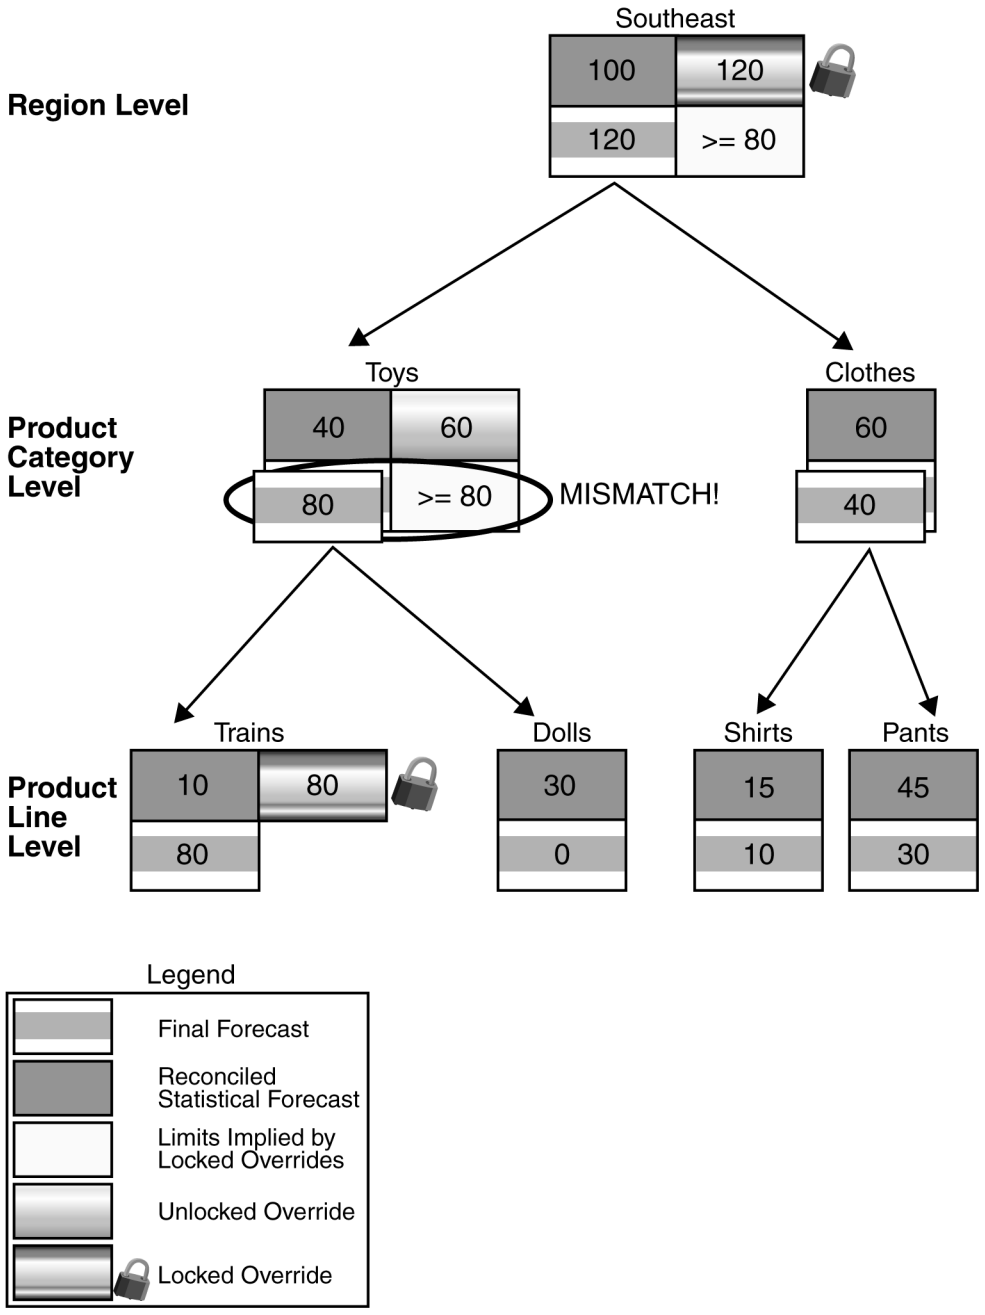

The following table explains the reconciliation process for each level of the hierarchy:

Level in the hierarchy	Reconciliation process for the nodes in this level
Region	Because you selected **Top Down** as the reconciliation method, SAS Forecast Studio starts reconciling the hierarchy at the highest level in the hierarchy. In this example, the highest level is the Region level. The reconciled statistical forecast for the Southeast node is 100; however, you specified 120 as a locked override for this level. SAS Forecast Studio honors this override, and consequently, the final forecast for the Southeast node is 120.
Product Category	In the Product Category level, there are 2 nodes—the Toys node and the Clothes node. For the Toys node, you specified an unlocked override of 60. SAS Forecast Studio does not honor the value of this unlocked override. Instead, SAS Forecast Studio uses the proportions to determine the final forecasts. The unlocked override that you specified for the Toys node is used instead of the reconciled statistical forecast to determine the final forecast. The unlocked override for the Toys node is 60, and the reconciled statistical forecast for the Clothes node is 60. Because these values are equal, the final forecast for the Southeast node (120) is divided equally between the 2 nodes, so the final forecast for both the Toys and Clothes nodes is 60. Now, you have a limit violation. The final forecast of 60 for the Toys node violates the limit implied by the locked override that you specified for the Trains node in the Product Line level. This limit states that the final forecast for the Toys node must be greater than or equal to 80. Therefore, SAS Forecast Studio cannot use the proportion that was implied by the unlocked override to determine the final forecast. Because you selected **No restriction** as the restriction on the reconciliation direction, SAS Forecast Studio sets the final forecast of the Toys node to 80. Because this implied limit was aggregated from a bottom-up approach, this operation is an example of an "implicit" bottom-up method of reconciliation as part of the top-down method of reconciliation. Because the final forecast of the Southeast node is 120, the final forecast of the Clothes node is 40 (or 120 - 80).

Level in the hierarchy	Reconciliation process for the nodes in this level
Product Line	In the Product Line level, there are 4 nodes. The 2 nodes below the Toys node are the Trains and Dolls nodes. The 2 nodes below the Clothes node are the Shirts and Pants nodes. In the branch of the hierarchy below the Toys node, you specified 80 as a locked override for the Trains node. SAS Forecast Studio honors this value, and the final forecast for the Trains node is 80. Because the final forecast for the Toys node is 80, the final forecast for the Dolls node is 0 (or 80 - 80). In the branch of the hierarchy below the Clothes node, you did not specify overrides for either of the nodes in the Product Line level. SAS Forecast Studio uses the proportions of the reconciled statistical forecasts to determine the final forecast for each node. In this example, the reconciled statistical forecast for the Shirts node is 25% of the reconciled statistical forecast for the Clothes node, and the reconciled statistical forecast for the Pants node is 75% of the Clothes node. Therefore, the final forecast for the Shirts node is 25% of the final forecast for the Clothes node (or .25 x 40 = 10). The final forecast for the Pants node is 75% of the final forecast for the Clothes node (or .75 x 40 = 30).

Example 3: Unreconciled Node in the Hierarchy

Purpose of the Example

This example shows the following:

- how SAS Forecast Studio tries to reconcile a hierarchy using the top-down method of reconciliation
- how a node could remain unreconciled because you choose to restrict the direction of the reconciliation process when you set up the hierarchy

Setup for the Example

In this example, you specified the following options when you set up the hierarchy:

- **Top Down** as the reconciliation method
- **Restrict to direction implied by reconciliation method** as the restriction on the reconciliation direction

For more information about these options, see “Modify the Hierarchy Settings” on page 53.

You also specified the following overrides for the Dec2008 time period:

- 120 as a locked override for the Southeast node in the Region level
- 60 as an unlocked override for the Toys node of the Product Category level
- 80 as a locked override for the Trains node in the Product Line level

NOTE: In addition to these settings, there are several assumptions about the data and option settings that are common to all the examples. To view these assumptions, see "Overview of the Examples" on page 110.

Explanation of the Reconciliation Process

The following figure shows how SAS Forecast Studio reconciles the hierarchy for the Dec2008 time interval:

Figure 9.4 Example 3: Unreconciled Node in Hierarchy

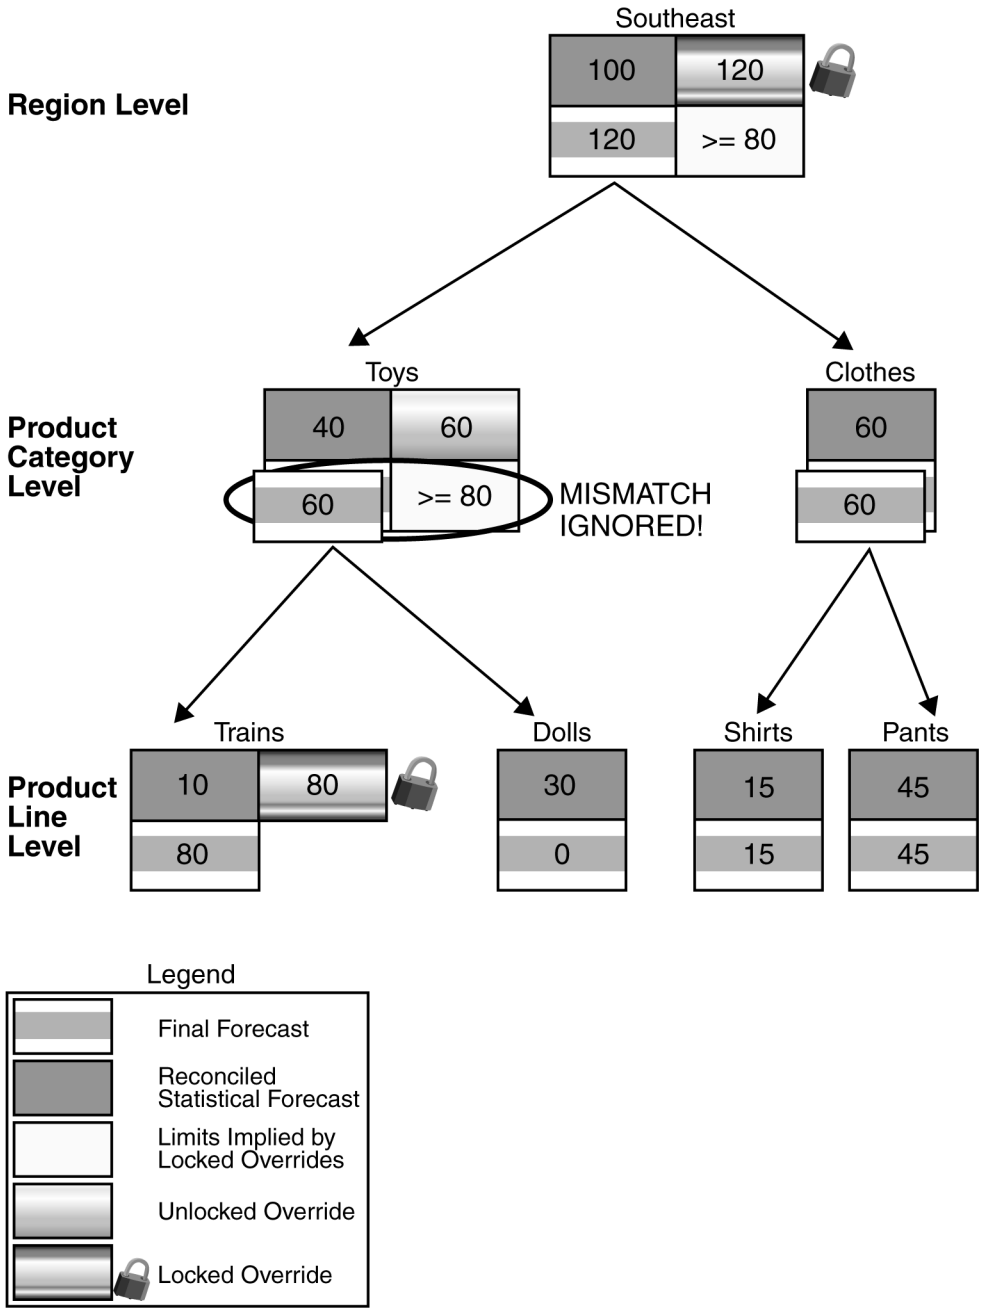

The following table explains the reconciliation process for each level of the hierarchy:

Level in the hierarchy	Reconciliation process for the nodes in this level
Region	Because you selected **Top Down** as the reconciliation method, SAS Forecast Studio starts reconciling the hierarchy at the highest level in the hierarchy. In this example, the highest level is the Region level. The reconciled statistical forecast for the Southeast node is 100; however, you specified 120 as a locked override for this level. SAS Forecast Studio honors this override, and consequently, the final forecast for the Southeast node is 120.
Product Category	In the Product Category level, there are 2 nodes—the Toys node and the Clothes node. For the Toys node, you specified an unlocked override of 60. SAS Forecast Studio does not honor the value of this unlocked override. Instead, SAS Forecast Studio uses the proportions to determine the final forecasts. The unlocked override that you specified for the Toys node is used instead of the reconciled statistical forecast to determine the final forecast. The unlocked override for the Toys node is 60, and the reconciled statistical forecast for the Clothes node is 60. Because these values are equal, the final forecast for the Southeast node (120) is divided equally between the 2 nodes, so the final forecast for both the Toys and Clothes nodes is 60. Now, you have a limit violation. The final forecast of 60 for the Toys node violates the limits implied by the locked override. This limit states that the final forecast for the Toys node must be greater than or equal to 80. However, when you set up the hierarchy, you chose to restrict the direction of the reconciliation, so unlike in the previous example, SAS Forecast Studio does not use the bottom-up method of reconciliation to reconcile the Toys node. Instead, the final forecasts for the Toys and Clothes nodes remain at 60, and the Toys node appears as an unreconciled node in the Reconciliation Failure Report.

Level in the hierarchy	Reconciliation process for the nodes in this level
Product Line	In the Product Line level, there are 4 nodes. The 2 nodes below the Toys node are the Trains and Dolls nodes. The 2 nodes below the Clothes node are the Shirts and Pants nodes. In the branch of the hierarchy below the Toys node, you specified 80 as a locked override for the Trains node. SAS Forecast Studio honors this value, and the final forecast for Trains is 80. Because the final forecast for the Toys node is 60, there is not enough left to allocate to the Dolls node, so the final forecast for the Dolls node is set to 0. In the branch of the hierarchy below the Clothes node, you did not specify overrides for either of the nodes. Therefore, SAS Forecast Studio uses the proportions of the reconciled statistical forecasts to determine the final forecast for each node. In this example, the reconciled statistical forecast for the Shirts node is 25% of the reconciled statistical forecast for the Clothes node, and the reconciled statistical forecast for the Pants node is 75% of the Clothes node. Therefore, the final forecast for the Shirts node is 25% of the final forecast for the Clothes node (or .25 x 60 = 15), and the final forecast for the Pants node is 75% of the final forecast for the Clothes node (or .75 x 60 = 45).

Example 4: Bottom-Up Method of Reconciliation with Locked and Unlocked Overrides in the Lowest Level

Purpose of the Example

This example shows the following:

- how SAS Forecast Studio reconciles a hierarchy using the bottom-up method of reconciliation
- how locked overrides have a direct impact on the final forecasts in the higher levels of the hierarchy
- how an unlocked override has no effect on the final forecast unless it is in the lowest level of the hierarchy

Setup for the Example

In this example, you specified the following options when you set up the hierarchy:

- **Bottom Up** as the reconciliation method

- either **No restriction** or **Restrict to direction implied by reconciliation method** as the restrictions on the reconciliation direction

For more information about these options, see "Modify the Hierarchy Settings" on page 53.

You also specified the following overrides for the Dec2008 time period:

- 60 as an unlocked override for the Toys node of the Product Category level
- 70 as an unlocked override for the Clothes node of the Product Category level
- 80 as a locked override for the Trains node in the Product Line level
- 5 as an unlocked override for the Shirts node in the Product Line level

NOTE: In addition to these settings, there are several assumptions about the data and option settings that are common to all the examples. To view these assumptions, see "Overview of the Examples" on page 110.

Explanation of the Reconciliation Process

The following figure shows how SAS Forecast Studio reconciles the hierarchy for the Dec2008 time interval:

Figure 9.5 Example 4: Bottom-Up Method of Reconciliation with No Conflicts

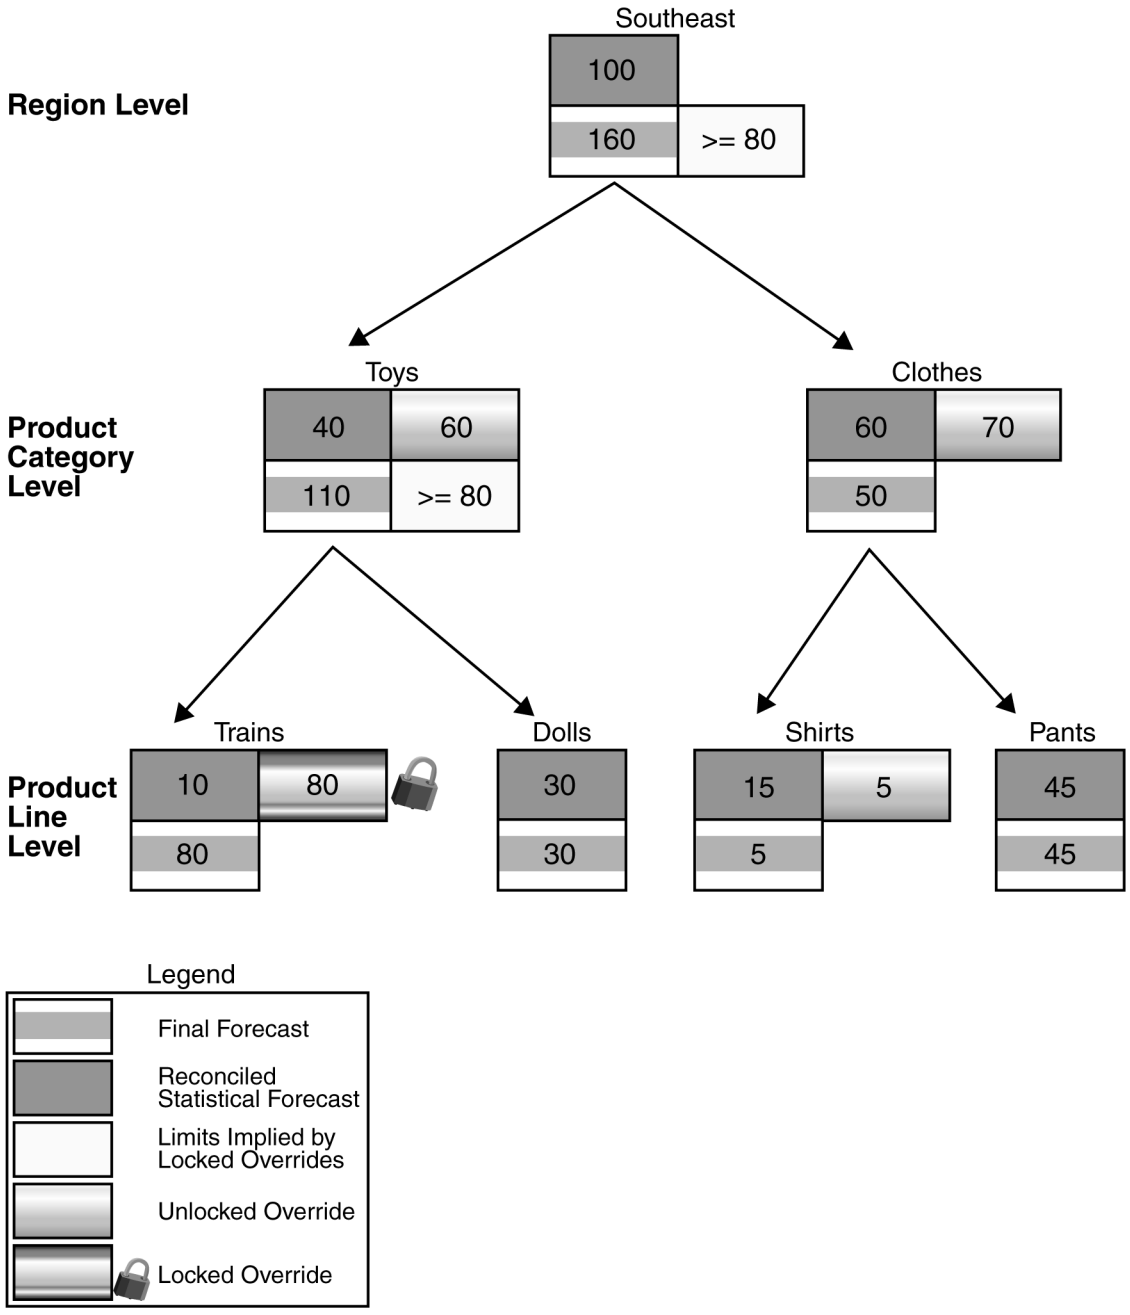

The following table explains the reconciliation process for each level of the hierarchy:

Level in the hierarchy	Reconciliation process for the nodes in this level
Product Line	Because you selected **Bottom Up** as the reconciliation method, SAS Forecast Studio starts reconciling the hierarchy at the lowest level in the hierarchy. In this example, the lowest level is the Product Line level. In the Product Line level, there are 4 nodes. The 2 nodes below the Toys node are the Trains and Dolls nodes. The 2 nodes below the Clothes node are the Shirts and Pants nodes. In the branch of the hierarchy below the Toys node, you specified 80 as a locked override for the Trains node. SAS Forecast Studio honors this value, and the final forecast for Trains is 80. Because you are using the bottom-up method of reconciliation, the final forecast for the Dolls node is equal to the reconciled statistical forecast of 30. In the branch of the hierarchy below the Clothes node, you specified an unlocked override of 5 for the Shirts node. Because you selected the bottom-up method of reconciliation, SAS Forecast Studio honors this unlocked override. Therefore, the final forecast for the Shirts node is 5. You did not specify an override for the Pants node, so the final forecast matches the reconciled statistical forecast of 45.
Product Category	In the Product Category level, there are 2 nodes—the Toys node and the Clothes node. SAS Forecast Studio uses the final forecasts of the nodes at the lower level to determine the final forecast for the nodes at the Product Category level. • For the Toys node, the final forecast for the Trains node (80) is added to the final forecast for the Dolls node (30), and so the final forecast for the Toys node is 110. Note that the unlocked override of 60 for the Toys node has no effect. • For the Clothes node, the final forecast for the Shirts node (5) is added to the final forecast for the Pants node (45), so the final forecast for the Clothes node is 50. Again, SAS Forecast Studio does not honor the unlocked override that you specified for the Clothes node.
Region	The values of the final forecasts for the nodes in the Product Category level determine the final forecast for the Southeast node. The combination of the final forecasts for the Toys node and the Clothes node equals 160, so the final forecast for the Southeast node is 160.

Example 5: Implicit Top-Down Reconciliation Process in the Bottom-Up Method of Reconciliation

Purpose of the Example

This example shows the following:

- how SAS Forecast Studio tries to reconcile a hierarchy using the bottom-up method of reconciliation
- how a mismatch can result at a higher level of the hierarchy, if you specified a locked override at that level
- how SAS Forecast Studio uses an implicit top-down pass to reconcile the hierarchy

Setup for the Example

In this example, you specified the following options when you set up the hierarchy:

- **Bottom Up** as the reconciliation method
- **No restriction** as the restriction on the reconciliation direction

For more information about these options, see "Modify the Hierarchy Settings" on page 53.

You also specified the following overrides for the Dec2008 time period:

- 200 as a locked override for the Southeast node in the Region level
- 60 as an unlocked override for the Toys node of the Product Category level
- 80 as a locked override for the Trains node in the Product Line level

NOTE: In addition to these settings, there are several assumptions about the data and option settings that are common to all the examples. To view these assumptions, see "Overview of the Examples" on page 110.

Explanation of the Reconciliation Process

The following figure shows how SAS Forecast Studio reconciles the hierarchy for the Dec2008 time interval:

Figure 9.6 Example 5: Implicit Top-Down Pass in the Bottom-Up Method of Reconciliation

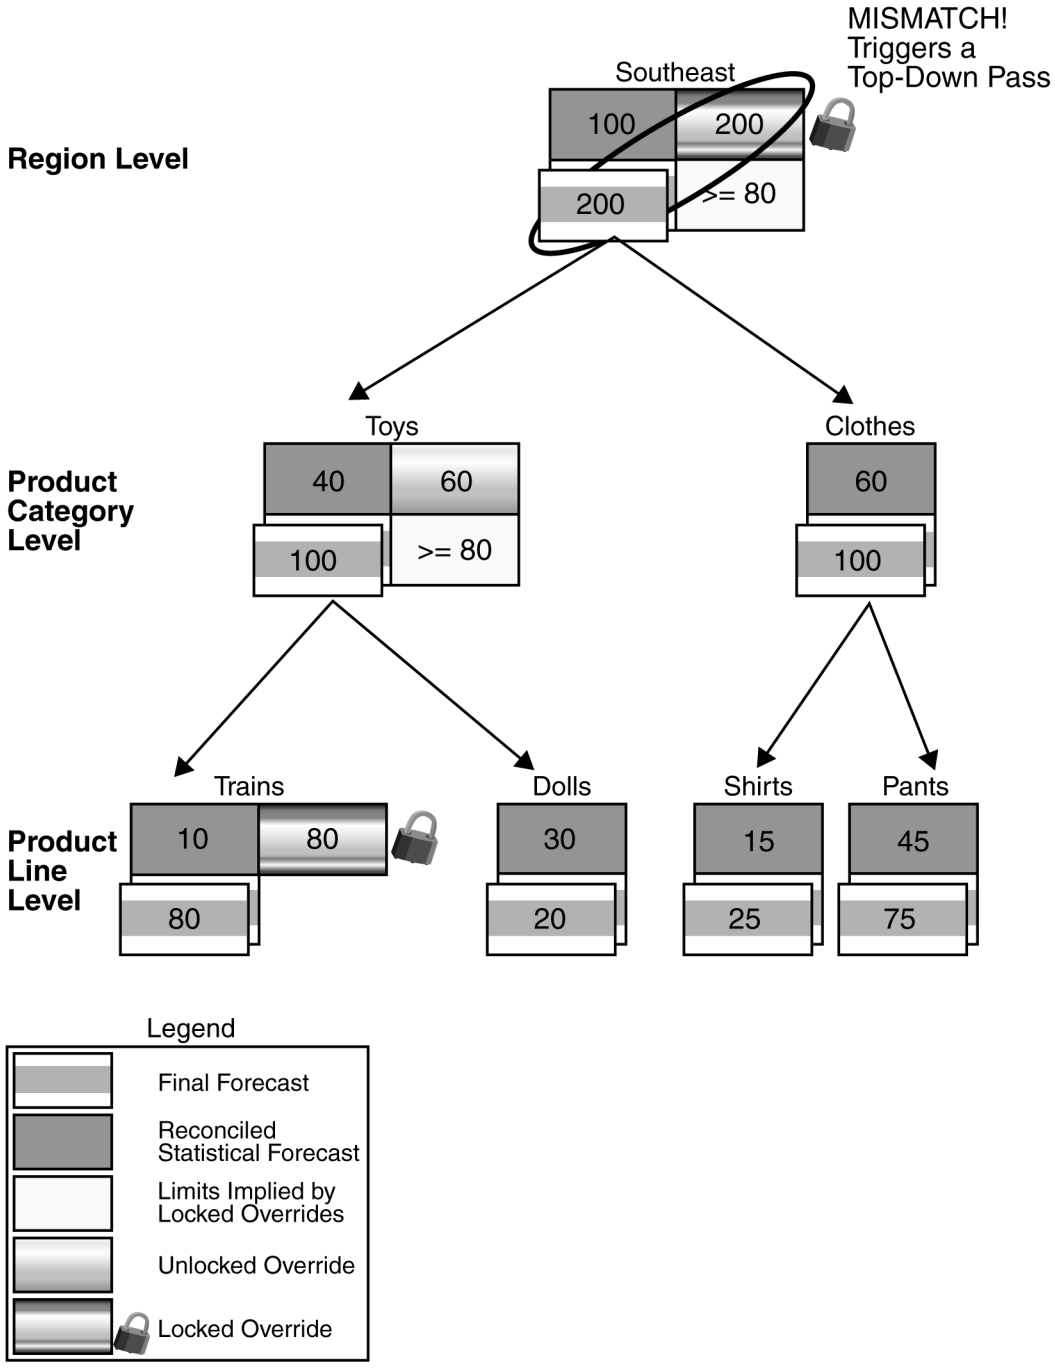

The following table explains the bottom-up reconciliation process for each level of the hierarchy:

Level in the hierarchy	**Reconciliation process for the nodes in this level**
Product Line	Because you selected **Bottom Up** as the reconciliation method, SAS Forecast Studio starts reconciling the hierarchy at the lowest level in the hierarchy. In this example, the lowest level is the Product Line level. In the Product Line level, there are 4 nodes. The 2 nodes below the Toys node are the Trains and Dolls nodes. The 2 nodes below the Clothes node are the Shirts and Pants nodes. In the branch of the hierarchy below the Toys node, you specified 80 as a locked override for the Trains node. SAS Forecast Studio honors this value, and the final forecast for the Trains node is set to 80. Because you are using the bottom-up method of reconciliation, the final forecast for the Dolls node is equal to the reconciled statistical forecast of 30. In the branch of the hierarchy below the Clothes node, the final forecasts equal the reconciled statistical forecasts. Therefore, the final forecast for the Shirts node is 15, and the final forecast for the Pants node is 45.
Product Category	In the Product Category level, there are 2 nodes—the Toys node and the Clothes node. SAS Forecast Studio uses the final forecasts of the nodes at the lower level to determine the final forecast for the nodes at the Product Category level. • For the Toys node, the final forecast for the Trains node (80) is added to the final forecast for the Dolls node (30), and so the final forecast for the Toys node is 110. • For the Clothes node, the final forecast for the Shirts node (15) is added to the final forecast for the Pants node (45), so the final forecast for the Clothes node is 60.
Region	For the Southeast node, you specified a locked override of 200. However, the combination of the final forecasts of the nodes in the Product Category level equal 170. Consequently, there is a mismatch between the final forecast and the locked override that you specified for the Southeast node. Because you selected the **No restriction** option as the restriction to the reconciliation direction, SAS Forecast Studio performs an implicit top-down pass of the hierarchy to reconcile the hierarchy.

The following table explains the implicit top-down reconciliation process for each level of the hierarchy:

Level in the hierarchy	Reconciliation process for the nodes in this level
Region	Because SAS Forecast Studio is now performing an implicit top-down reconciliation, it replaces the final forecast of 170 that was calculated by the bottom-up method of reconciliation with the 200 that you specified as the locked override.
Product Category	You specified an unlocked override of 60 for the Toys node. SAS Forecast Studio uses this unlocked override to determine the proportions during the implicit top-down pass. The unlocked override that you specified for the Toys node is 60, and the final forecast for the Clothes node is 60. Because these nodes are equal, SAS Forecast Studio divides the 200 value from the Southeast node equally between these 2 nodes, so the final forecast for both the Clothes and Toys nodes is 100.
Product Line	In the branch of the hierarchy below the Toys node, you specified a locked override of 80 for the Trains node. SAS Forecast Studio honors this locked override, and the final forecast for the Trains node is set to 80. Because the final forecast for the Toys node is 100, the final forecast for the Dolls node is 100 - 80 = 20. In the branch of the hierarchy below the Clothes node, you did not specify a locked override for either of the nodes. Therefore, SAS Forecast Studio uses the proportions of the reconciled statistical forecasts to determine the final forecast for each node. In this example, the reconciled statistical forecast for the Shirts node is 25% of the reconciled statistical forecast for the Clothes node, and the reconciled statistical forecast for the Pants node is 75% of the Clothes node. Therefore, the final forecast for the Shirts node is 25% of the final forecast for the Clothes node (or .25 x 100 = 25), and the final forecast for the Pants node is 75% of the final forecast for the Clothes node (or .75 x 100 = 75).

Example 6: Implied Limits Have Precedence over an Unlocked Override

Purpose of the Example

This example shows the following:

- how SAS Forecast Studio tries to reconcile a hierarchy using the bottom-up method of reconciliation
- how a mismatch can result at a higher level of the hierarchy if you specified a locked override at that level
- how SAS Forecast Studio uses an implicit top-down pass to reconcile the hierarchy

- how an implied limit at the middle level can have precedence over an unlocked override

Setup for the Example

In this example, you specified the following options when you set up the hierarchy:

- **Bottom Up** as the reconciliation method
- **No restriction** as the restriction on the reconciliation direction

For more information about these options, see "Modify the Hierarchy Settings" on page 53.

You also specified the following overrides for the Dec2008 time period:

- 100 as a locked override for the Southeast node in the Region level
- 60 as an unlocked override for the Toys node of the Product Category level
- 80 as a locked override for the Trains node in the Product Line level

NOTE: In addition to these settings, there are several assumptions about the data and option settings that are common to all the examples. To view these assumptions, see "Overview of the Examples" on page 110.

Explanation of the Reconciliation Process

The following figure shows how SAS Forecast Studio reconciles the hierarchy for the Dec2008 time interval:

Figure 9.7 Example 6: Implied Limits Have Precedence over an Unlocked Override

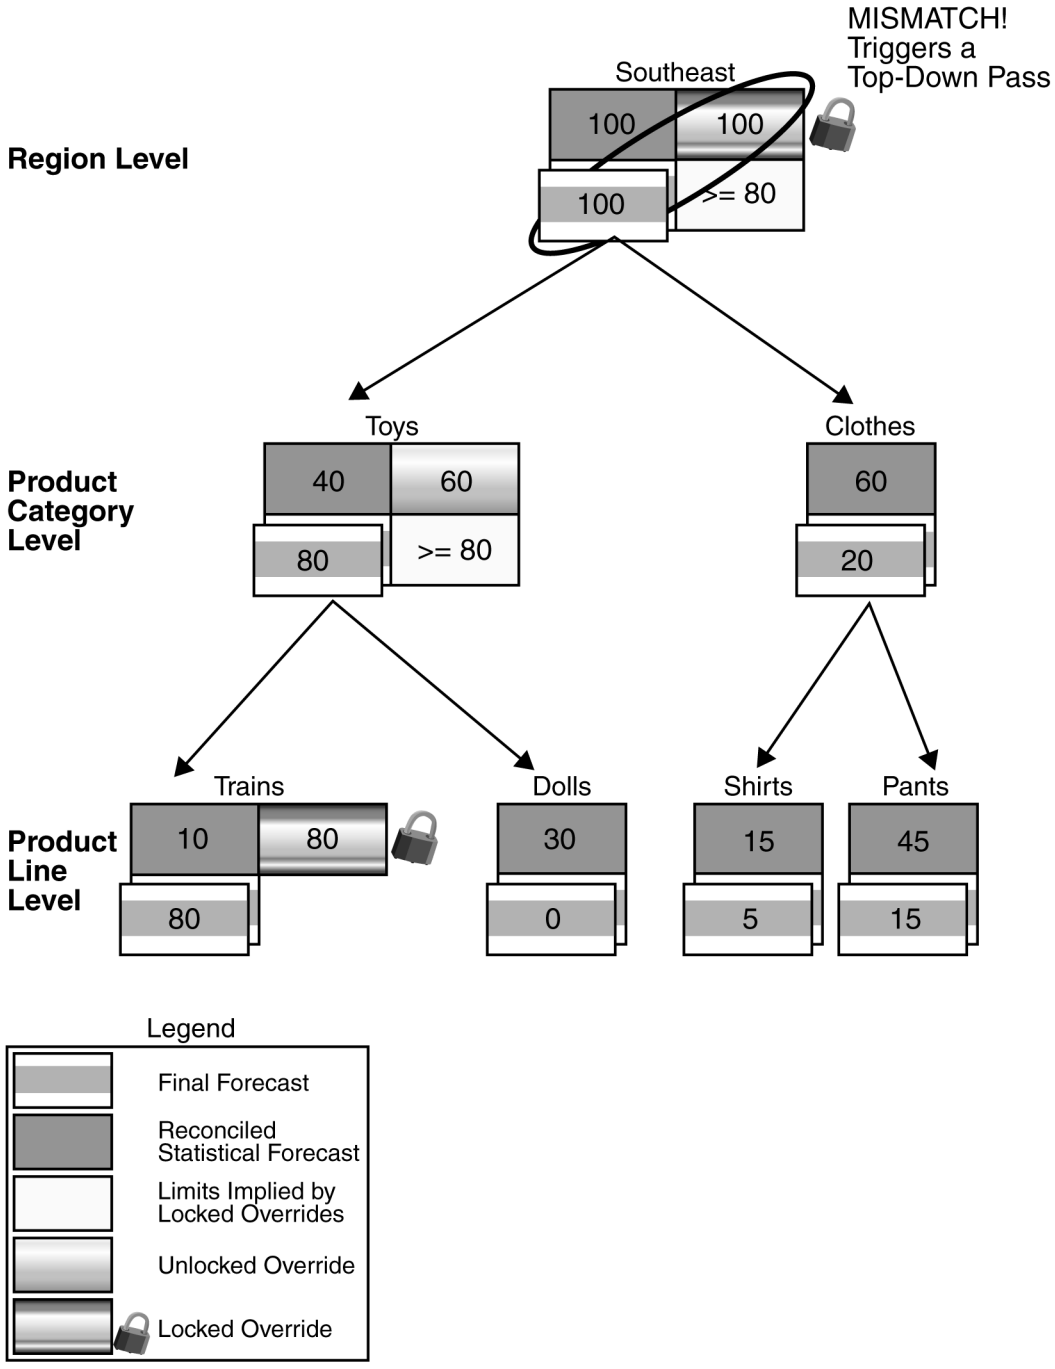

The following table explains the bottom-up reconciliation process for each level of the hierarchy:

Level in the hierarchy	**Reconciliation process for the nodes in this level**
Product Line	Because you selected **Bottom Up** as the reconciliation method, SAS Forecast Studio starts reconciling the hierarchy at the lowest level in the hierarchy. In this example, the lowest level is the Product Line level. In the Product Line level, there are 4 nodes. The 2 nodes below the Toys node are the Trains and Dolls nodes. The 2 nodes below the Clothes node are the Shirts and Pants nodes. In the branch of the hierarchy below the Toys node, you specified 80 as a locked override for the Trains node. SAS Forecast Studio honors this value, and the final forecast for the Trains node is set to 80. Because you are using the bottom-up method of reconciliation, the final forecast for the Dolls node is equal to the reconciled statistical forecast of 30. In the branch of the hierarchy below the Clothes node, the final forecasts equal the reconciled statistical forecasts. Therefore, the final forecast for the Shirts node is 15, and the final forecast for the Pants node is 45.
Product Category	In the Product Category level, there are 2 nodes—the Toys node and the Clothes node. SAS Forecast Studio uses the final forecasts of the nodes at the lower level to determine the final forecast for the nodes at the Product Category level. • For the Toys node, the final forecast for the Trains node (80) is added to the final forecast for the Dolls node (30), and so the final forecast for the Toys node is 110. • For the Clothes node, the final forecast for the Shirts node (15) is added to the final forecast for the Pants node (45), so the final forecast for the Clothes node is 60.
Region	For the Southeast node, you specified a locked override of 100. However, the combination of the final forecasts for the nodes in the Product Category level equals 170, and consequently, there is a mismatch between the final forecast and the locked override that you specified for the Southeast node. Because you selected the **No restriction** option as the restriction on the reconciliation direction, SAS Forecast Studio performs an implicit top-down pass to reconcile the hierarchy.

The following table explains the implicit top-down reconciliation process for each level of the hierarchy:

Level in the hierarchy	Reconciliation process for the nodes in this level
Region	Because SAS Forecast Studio is now performing an implicit top-down reconciliation, it replaces the final forecast of 170 that was calculated by the bottom-up method of reconciliation with the 100 that you specified as the locked override.
Product Category	You specified an unlocked override of 60 for the Toys node. However, you specified a locked override of 80 for the Trains node in the Product Category level. This implies that the value of the Toys node must be greater than or equal to 80. In this case, the implied limit has precedence over the unlocked override in determining the proportions during the implicit top-down pass. The implied limit for the Toys node is 80% of the final forecast for the Southeast node. Therefore, the final forecast for the Toys node is 80 (or .80 x 100 = 80), and the final forecast for the Clothes node is 20 (or .20 x 100 = 20).
Product Line	In the branch of the hierarchy below the Toys node, you specified a locked override of 80 for the Trains node. SAS Forecast Studio honors this locked override, and the final forecast for the Trains node is set to 80. Because the final forecast for the Toys node is 80, the final forecast for the Dolls node is set to 0. In the branch of the hierarchy below the Clothes node, you did not specify a locked override for either of the nodes. Therefore, SAS Forecast Studio uses the proportions of the reconciled statistical forecasts to determine the final forecast for each node. In this example, the reconciled statistical forecast for the Shirts node is 25% of the reconciled statistical forecast for the Clothes node, and the reconciled statistical forecast for the Pants node is 75% of the Clothes node. Therefore, the final forecast for the Shirts node is 25% of the final forecast for the Clothes node (or .25 x 20 = 5), and the final forecast for the Pants node is 75% of the final forecast for the Clothes node (or .75 x 20 = 15).

Example 7: Middle-Out Method of Reconciliation with No Conflicts

Purpose of This Example

This example shows the following:

- how SAS Forecast Studio reconciles a hierarchy using the middle-out method of reconciliation
- how the forecasts at the middle level are unchanged because they do not violate the limits implied by a locked override in a lower level of the hierarchy, and there is no locked override

in a higher level of the hierarchy

Setup for the Example

In this example, you specified the following options when you created the hierarchy:

- **Middle Out** as the reconciliation method
- either **No restriction** or **Restrict to direction implied by reconciliation method** as the restriction on the reconciliation direction

For more information about these options, see "Modify the Hierarchy Settings" on page 53.

You also specified the following overrides for the Dec2008 time period:

- 60 as an unlocked override for the Toys node of the Product Category level
- 20 as a locked override for the Trains node in the Product Line level

NOTE: In addition to these settings, there are several assumptions about the data and option settings that are common to all the examples. To view these assumptions, see "Overview of the Examples" on page 110.

Explanation of the Reconciliation Process

The following figure shows how SAS Forecast Studio reconciles the hierarchy for the Dec2008 time interval:

Figure 9.8 Example 7: Middle-Out Method of Reconciliation with No Conflicts

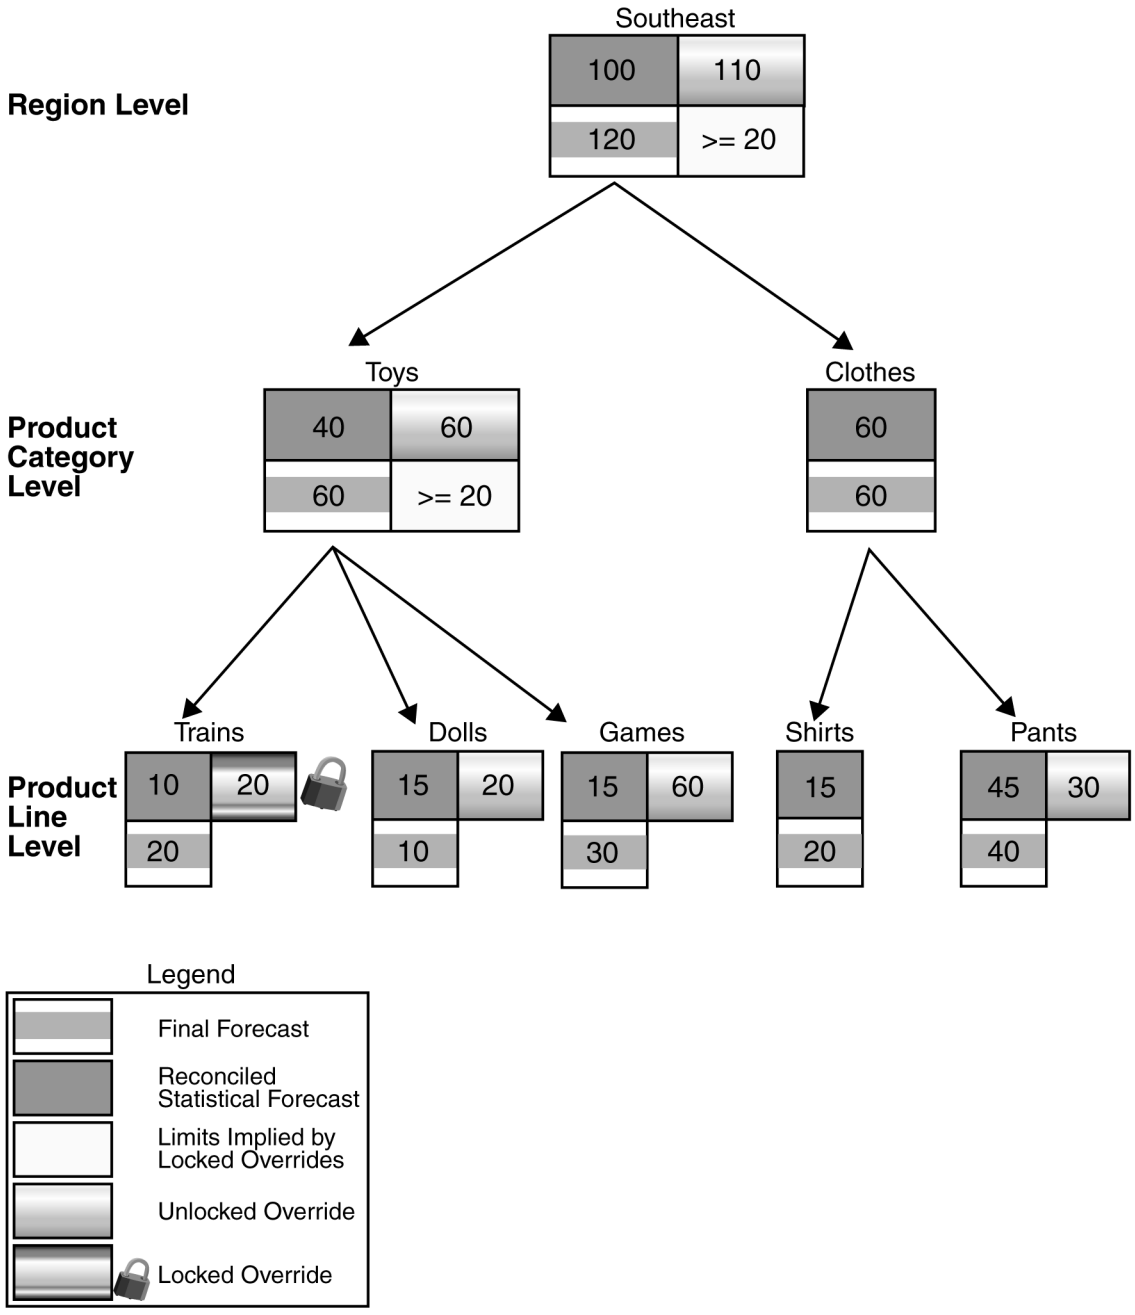

The following table explains the reconciliation process for each level of the hierarchy:

Level in the hierarchy	**Reconciliation process for the nodes in this level**
Product Category	Because you selected **Middle Out** as the reconciliation method, SAS Forecast Studio starts reconciling the hierarchy at the Product Category level. In the Product Category level, there are 2 nodes—the Toys node and the Clothes node. For the Toys node, you specified an unlocked override of 60. Because SAS Forecast Studio is starting the reconciliation at this level, this unlocked override is honored, so the final forecast for the Toys node is 60. For the Clothes node, the final forecast is equal to the reconciled statistical forecast of 60.
Region	SAS Forecast Studio calculates the final forecast for the Southeast node by using the bottom-up method of reconciliation. The values of the final forecasts for the nodes in the Product Category level determine the final forecast for the Southeast node. The combination of the final forecasts for the Toys node and the Clothes node equals 120, so the final forecast for the Southeast node is 120.

Level in the hierarchy	Reconciliation process for the nodes in this level
Product Line	SAS Forecast Studio calculates the final forecast for the nodes in the lowest level in the hierarchy by using the top-down method of reconciliation. In the Product Line level, there are 5 nodes. The 3 nodes below the Toys node are the Trains, Dolls, and Games nodes. The 2 nodes below the Clothes node are the Shirts and Pants nodes. For the branch of the hierarchy below the Toys node, you specified a locked override of 20 for the Trains node. SAS Forecast Studio honors locked overrides, so the final forecast for the Trains node is 20. Because SAS Forecast Studio is using the top-down method of reconciliation for this level of the hierarchy, the unlocked overrides that you specified for the Dolls and Games nodes do not match the final forecasts. However, these unlocked overrides are used to determine the proportions for the final forecasts. The final forecast for the Toys node is 60; however, the final forecast for the Trains node is 20, which means that the combined final forecasts for the Dolls and Games nodes must equal 40. The values of these nodes are calculated by using the following proportions: • For the Dolls node, you specified an unlocked override of 20. The total of the unlocked overrides for the Dolls and Games nodes is 20 + 60 = 80. Because the proportion of the Dolls node is 20/80 = 25% of this total value, the final forecast for the Dolls node is 25% of the final forecast for the Toys node. Therefore, the final forecast for the Dolls node is 40 x .25 = 10. • For the Games node, you specified an unlocked override of 60. Because the proportion of the Games node is 60/80 = 75% of this total value, the final forecast for the Games node is 75% of the final forecast for the Toys node. Therefore, the final forecast for the Games node is 40 x .75 = 30.

Level in the hierarchy	Reconciliation process for the nodes in this level
Product Line (continued)	For the branch of the hierarchy below the Clothes node, you specified an unlocked override of 30 for the Pants node. When calculating the proportions, SAS Forecast Studio uses this unlocked override instead of the reconciled statistical forecast. Using the unlocked override for the Pants node, the total for the Shirts and Pants nodes is 15 + 30 = 45. The values of these nodes are calculated using the following proportions: • Because the Shirt node is 15/45 = 33.3% of this total value, the final forecast for the Shirts node is 33.3% of the final forecast for the Clothes node. Therefore, the final forecast for the Shirts node is 60 x .333 = 20. • Because the Pants node is 30/45 = 66.6% of this total value, the final forecast for the Pants node is 66.6% of the final forecast for the Clothes node. Therefore, the final forecast for the Pants node is 60 x .666 = 40.

Example 8: Middle-Out Method of Reconciliation with an Unreconciled Node

Purpose of This Example

This example shows the following:

- How SAS Forecast Studio reconciles a hierarchy using the middle-out method of reconciliation
- How selecting the **Restrict to direction implied by reconciliation method** option will honor all overrides, including unlocked overrides. As this example shows, this restriction on reconciliation direction could result in unreconciled nodes

Setup for the Example

In this example, you specified the following options when you set up the hierarchy:

- **Middle Out** as the reconciliation method
- **Restrict to direction implied by reconciliation method** as the restriction on the reconciliation direction

For more information about these options, see “Modify the Hierarchy Settings” on page 53.

You also specified the following overrides for the Dec2008 time period:

- 30 as an unlocked override for the Toys node of the Product Category level
- 50 as a locked override for the Trains node in the Product Line level
- 30 as a locked override for the Pants node in the Product Line level

NOTE: In addition to these settings, there are several assumptions about the data and option settings that are common to all the examples. To view these assumptions, see "Overview of the Examples" on page 110.

Explanation of the Reconciliation Process

The following figure shows how SAS Forecast Studio reconciles the hierarchy for the Dec2008 time interval:

Figure 9.9 Example 8: Middle-Out Method of Reconciliation with an Unreconciled Node

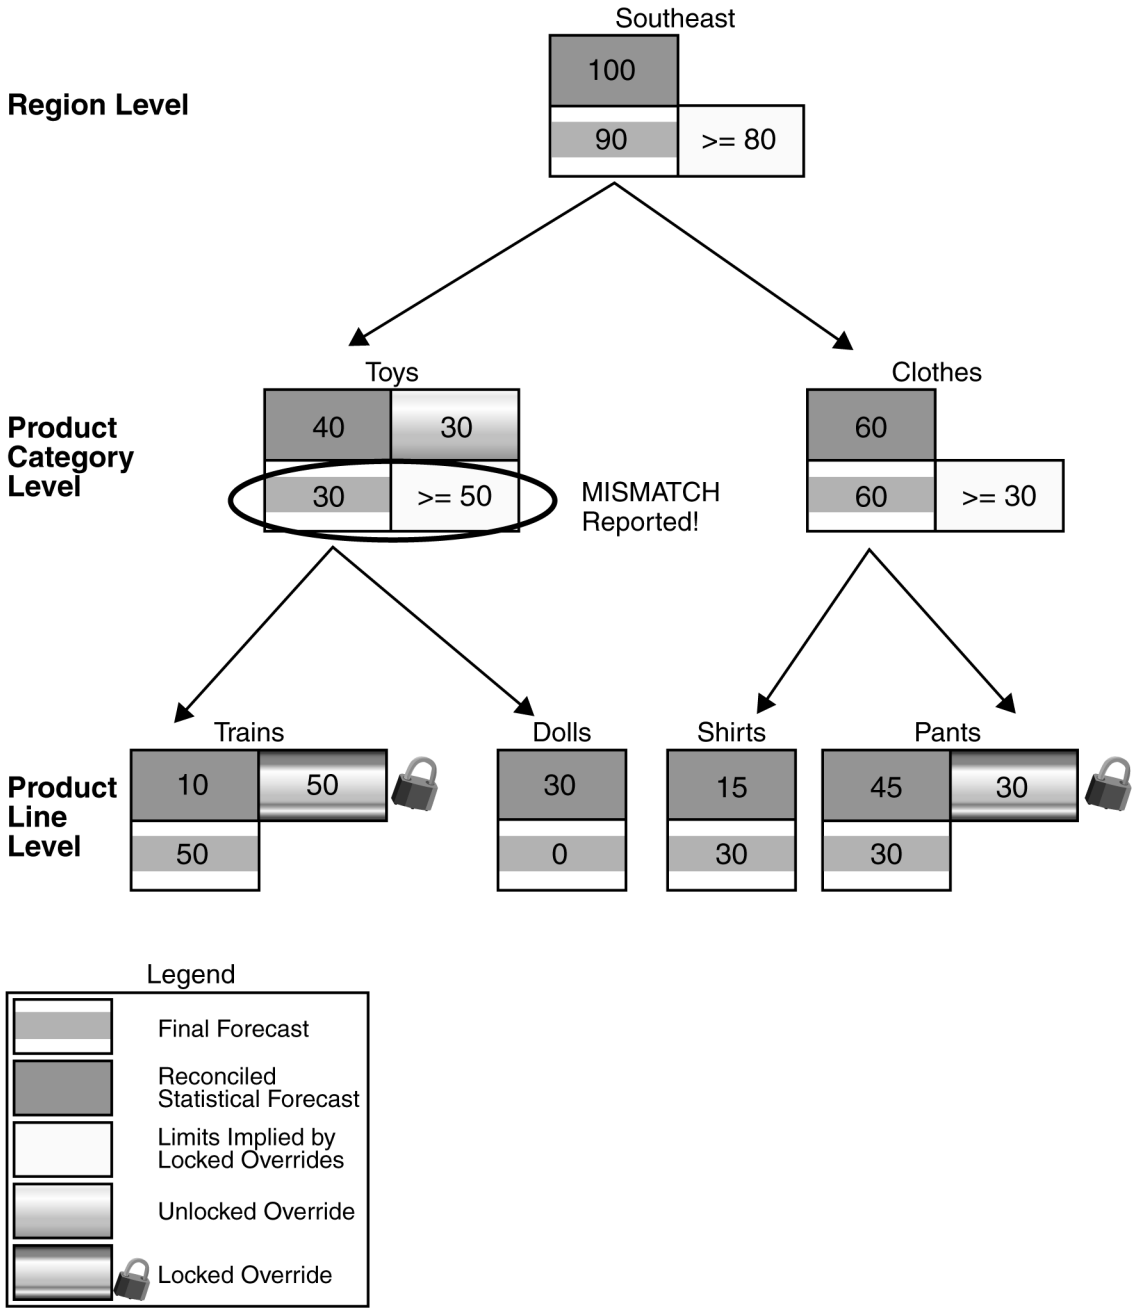

The following table explains the reconciliation process for each level of the hierarchy:

Level in the hierarchy	Reconciliation process for the nodes in this level
Product Category	Because you selected **Middle Out** as the reconciliation method, SAS Forecast Studio starts reconciling the hierarchy at the Product Category level. In the Product Category level, there are 2 nodes—the Toys node and the Clothes node. For the Toys node, you specified an unlocked override of 30. Because you selected the **Restrict to direction implied by reconciliation method** option, SAS Forecast Studio honors this unlocked override, and the final forecast for the Toys node is set to 30. However, the limits that are implied by the locked override in the lowest level indicate that the final forecast must be greater than 50. SAS Forecast Studio leaves this node unreconciled, and this conflict is reported in the Reconciliation Failure Report. For the Clothes node, the final forecast is equal to the reconciled statistical forecast of 60. Because 60 is greater than the limits that are implied by a locked override in the lower level of the hierarchy, there are no unreconciled nodes in this branch of the hierarchy.
Region	SAS Forecast Studio calculates the final forecast for the Southeast node by using the bottom-up method of reconciliation. Because you had a conflict in the Toys node in the Product Category level of the hierarchy, SAS Forecast Studio uses the limit that is implied by the locked override in the Trains node to determine the final forecast for the Southeast node. The final forecasts for the Toys node (30) and the Clothes node (60) are combined to create the final forecast for the Southeast node, so the final forecast for the Southeast node is 90.

Level in the hierarchy	**Reconciliation process for the nodes in this level**
Product Line	SAS Forecast Studio calculates the final forecast for the nodes in the lowest level in the hierarchy by using the top-down method of reconciliation. In the Product Line level, there are 4 nodes. The 2 nodes below the Toys node are the Trains and Dolls nodes. The 2 nodes below the Clothes node are the Shirts and Pants nodes. For the branch of the hierarchy below the Toys node, you specified a locked override of 50 for the Trains node. SAS Forecast Studio honors locked overrides, so the final forecast for the Trains node is 50. Because the final forecast for the Toys node is 30, there is not enough left to allocate to the Dolls node, so the final forecast for the Dolls node is set to 0. For the branch of the hierarchy below the Clothes node, you specified a locked override of 30 for the Pants node. SAS Forecast Studio honors locked overrides, so the final forecast for the Pants node is 30. Because the final forecast for the Clothes node is 60, the final forecast for the Shirts node is 60-30=30.

Chapter 10

Creating User-Defined Models

Contents

Create a New Model Based on a Default Model

SAS Forecast Studio ships with several system-generated models. You might want to copy one of these models and customize it to meet your needs. You can do this from the Model Repository dialog box. For more information, see "Create a New Custom Model Based on a Default Model" on page 190.

Create a User-Defined Model

To create a new model, complete the following steps:

1. Select **Project** → **Model Repository**. The Model Repository dialog box opens. For more information, see "About the Model Repository" on page 188.

   NOTE: You can also create a new model when you add a model to a series. For more information, see "Add a Model to a Series" on page 196.

2. Click the **Custom Models** tab in the Model Repository dialog box.

3. On the **Custom Models** tab, click **New**. The New Model dialog box opens.

4. In the New Model dialog box, select **Create an entirely new model** to create a user-defined model. From the **Model type** drop-down list, you can choose from the following options:

   ARIMA
   creates AutoRegressive Integrated Moving-Average (ARIMA) models. These models can include events and independent variables. For more information, see "Create an ARIMA Model" on page 146.

   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. For more information, see "Create a Subset (Factored) ARIMA Model" on page 150.

   Unobserved Components
   creates unobserved components models. These models can include events and independent variables. For more information, see "Create an Unobserved Components Model" on page 153.

   Exponential smoothing
   uses exponential smoothing models. You can choose from many types of exponential smoothing models. For more information, see "Create an Exponential Smoothing Model" on page 164.

   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; therefore, these models will not work for intermittent time series data. For more information, see "Create an Intermittent Demand Model" on page 167.

   Multiple regression
   creates a multiple linear regression model with autocorrelated errors. You can specify which independent variables and events to include in the model. For more information, see "Create a Multiple Regression Model" on page 170.

   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. For more information, see "Create a Moving Average Model" on page 173.

   Curve fitting
   creates a curve fitting model that enables you to identify trends and relationships in your time series data. For more information, see "Create a Curve Fitting Model" on page 175.

   Random walk
   creates a random walk model. When you create a random walk model, you can specify whether to include the drift and seasonal terms in the model. For more information, see "Create a Random Walk Model" on page 177.

5. Click **OK** to create a new model. The dialog box that opens depends on the options that you selected. After you create a model, it is fitted to every series in the project.

ARIMA Models

Create an ARIMA Model

When you create an ARIMA model, you can specify the autoregressive and moving average polynomials of an ARIMA model. You can also specify one or more input variables.

To create an ARIMA model, complete the following steps:

1. Select **Project** → **Model Repository**. The Model Repository dialog box opens. For more information, see "About the Model Repository" on page 188.

   NOTE: You can also create a new model when you add a model to a series. For more information, see "Add a Model to a Series" on page 196.

2. Click the **Custom Models** tab in the Model Repository dialog box.

3. On the **Custom Models** tab, click **New**. The New Model dialog box opens.

4. In the New Model dialog box, select **Create an entirely new model**.

5. From the **Model type** drop-down list, select **ARIMA** as the model type. The ARIMA Model dialog box opens.

Figure 10.1 ARIMA Model Dialog Box

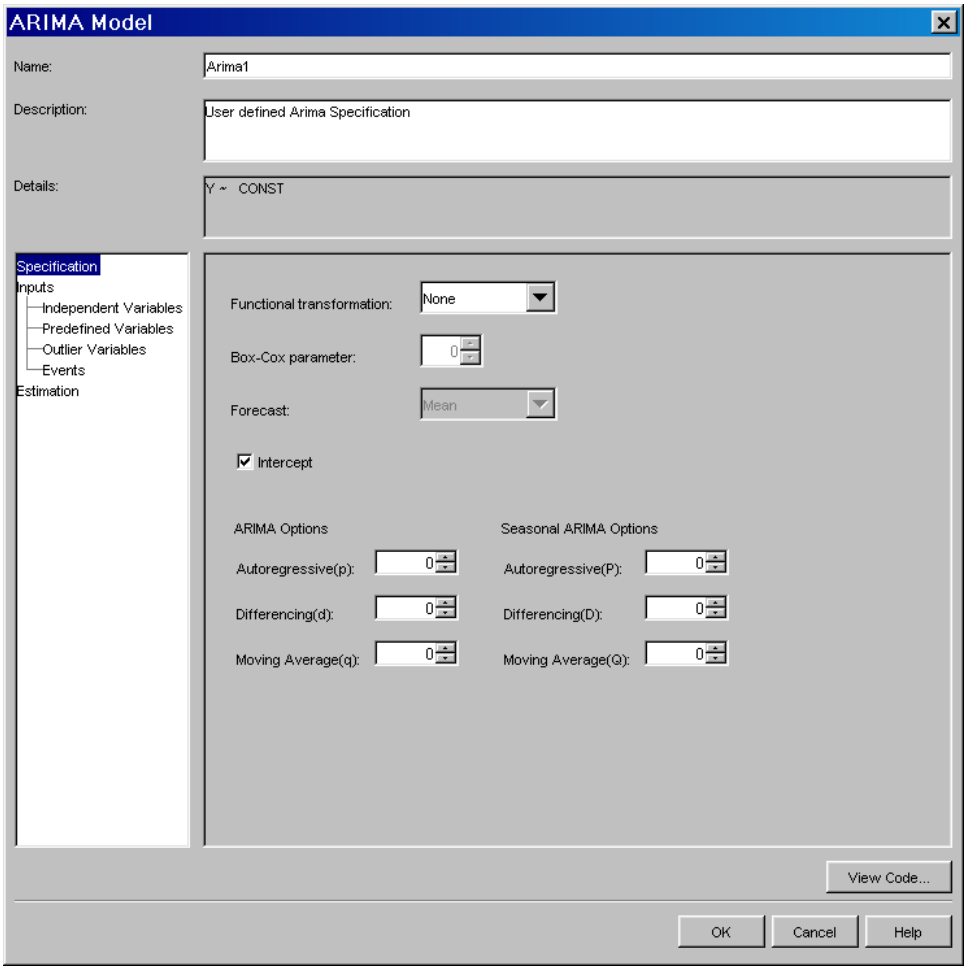

NOTE: The ARIMA Model dialog box also opens when you edit a model. The options that are available when you are creating or editing a model are the same.

6. Specify a unique name for the model. By default, SAS Forecast Studio automatically generates a unique name for the model.

   NOTE: SAS Forecast Studio has a list of reserved model names. For more information, see "Reserved Model Names" on page 179.

7. (Optional) Specify a description for the new model.

   NOTE: The Details pane shows a representation of the model specification. SAS Forecast Studio generates this representation based on the options that you select in the ARIMA Model dialog box. You cannot edit this text.

8. Use the selection pane to specify 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 ARIMA Model dialog box, select **Specification** in the selection pane to specify a basic nonseasonal (ARIMA(p, d, q)) or seasonal (ARIMA(p, d, q)(P, D, Q))s model.

You can specify the following options:

- which functional transformation to apply to the dependent series. You can select from the following transformations:
  - no transformation (which 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 parameter** box.

  NOTE: You can specify a transformation only when the values of the series are positive.
- how to calculate the forecasts. Forecasts can be based on the mean or median. By default, the mean value is used.
- whether an intercept term is included in the model. To include an intercept term, select the **Intercept** check box. By default, the intercept is included.

  NOTE: If you specify a differencing order, then a message appears asking if you want to suppress the intercept. It is recommended that you suppress the intercept when a differencing order is specified.
- the autoregressive (p and P), differencing (d and D), and moving average (q and Q) orders for your model.

  NOTE: The maximum value that you can specify for each order is 13.

  The ARIMA model has nonseasonal and seasonal components. Use the up and down arrows next to each box to specify an integer value for the following orders:
  - the nonseasonal and seasonal autoregressive orders (p and P).

    If you set p = 4, then you are implying autoregressive orders of (1 2 3 4) in the nonseasonal component of the model. If you set p = 4 and P = 3, then you are implying autoregressive orders of (1 2 3 4) in the nonseasonal component of the model and autoregressive orders of (1 2 3) in the seasonal component of the model. In the Details pane, this order appears as follows: P = (1 2 3 4)(1 2 3)s.

    NOTE: The seasonal length of the order is identified by the placeholder "s".

– the nonseasonal and seasonal differencing orders (d and D).

  If you set d = 4, then you are implying differencing orders of (1 1 1 1) in the nonseasonal component of the model. If you set d = 4 and D = 3, then you are implying (1 1 1 1) in the nonseasonal component of the model and (s s s) in the seasonal component of the model. In the Details pane, this order appears as follows: D = (1 1 1 1 s s s).

– the nonseasonal and seasonal moving average orders (q and Q).

  If you set q = 4, then you are implying moving average orders of (1 2 3 4) in the nonseasonal component of the model. If you set q = 4 and Q = 3, then you are implying (1 2 3 4) in the nonseasonal component of the model and (1 2 3) in the seasonal component of the model. In the Details pane, this order appears as follows: Q = (1 2 3 4)(1 2 3)s.

NOTE: 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, and so on) is implied by the Time ID variable in the project.

NOTE: The seasonal options are available only if the data is seasonal.

For more information about the syntax of these orders or for more examples, see "Autoregressive, Differencing, and Moving Average Orders" on page 179.

Inputs in ARIMA Models

By default, no independent variables, predefined variables, or events are used as inputs in the model. Any outliers that were detected are automatically included in the model. You can add and remove inputs from the model.

For more information, see the following topics:

- "Include Independent Variables in User-Defined Models" on page 182
- "Include Predefined Variables in User-Defined Models" on page 182
- "Include Outlier Variables in User-Defined Models" on page 183
- "Include Events in User-Defined Models" on page 183

Estimation for ARIMA Models

For ARIMA and Subset (factored) ARIMA models, you can specify a number of estimation options. For more information, see "Setting the Estimation Options" on page 185.

Subset (Factored) ARIMA Models

Create a Subset (Factored) ARIMA Model

When creating a Subset (factored) ARIMA model, 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 ARIMA Model dialog box is similar to the ARIMA Model dialog box, except that it uses a more general specification of the ARIMA options.

To create a Subset (factored) ARIMA model, complete the following steps:

1. Select **Project** → **Model Repository**. The Model Repository dialog box opens. For more information, see "About the Model Repository" on page 188.

   NOTE: You can also create a new model when you add a model to a series. For more information, see "Add a Model to a Series" on page 196.

2. Click the **Custom Models** tab in the Model Repository dialog box.

3. On the **Custom Models** tab, click **New**. The New Model dialog box opens.

4. In the New Model dialog box, select **Create an entirely new model**.

5. From the **Model type** drop-down list, select **Subset (factored) ARIMA** as the model type. The Subset ARIMA Model dialog box opens.

Figure 10.2 Subset ARIMA Model Dialog Box

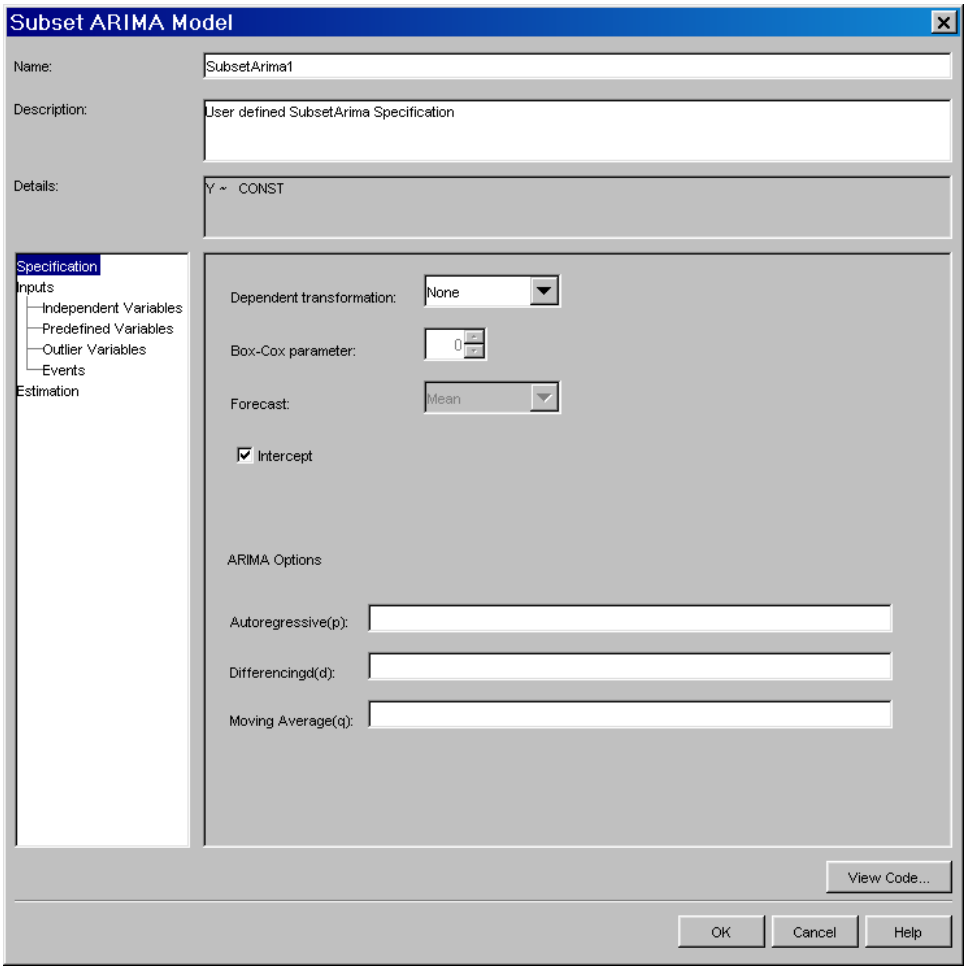

NOTE: The Subset ARIMA Model dialog box also opens when you edit a model. The options that are available when you are creating or editing a model are the same.

6. Specify a unique name for the model. By default, SAS Forecast Studio automatically generates a unique name for the model.

   NOTE: SAS Forecast Studio has a list of reserved model names. For more information, see "Reserved Model Names" on page 179.

7. (Optional) Specify a description for the new model.

   NOTE: The Details pane shows a representation of the model specification. SAS Forecast Studio generates this representation based on the options that you select in the Subset ARIMA Model dialog box. You cannot edit this text.

8. Use the selection pane to specify the following options:

   - specification
   - inputs
   - estimation

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

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

Specification Options for Subset ARIMA Models

In the Subset ARIMA Model dialog box, select **Specification** in the selection pane to access the following options:

- which dependent transformation to apply to the dependent series. You can select from the following transformations:
  - no transformation (which 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 parameter** box.
- how to calculate the forecasts. Forecasts can be based on the mean or median. By default, the mean value is used.
- whether an intercept is included in the model. If you specify a differencing term, then a message appears asking if you want to suppress the intercept.
- the autoregressive (p), differencing (d), and moving average (q) orders for your model.

  For SAS Forecast Studio to interpret these orders correctly, you must specify these options using the correct syntax. For more information about how to specify these orders, see "Autoregressive, Differencing, and Moving Average Orders" on page 179.

Inputs in Subset ARIMA Models

By default, no independent variables, predefined variables, or events are used as inputs in the model. Any outliers that were detected are automatically included in the model. You can add and remove inputs from the model.

For more information, see the following topics:

- "Include Independent Variables in User-Defined Models" on page 182
- "Include Predefined Variables in User-Defined Models" on page 182
- "Include Outlier Variables in User-Defined Models" on page 183
- "Include Events in User-Defined Models" on page 183

Estimation for Subset ARIMA Models

For ARIMA and Subset (factored) ARIMA models, you can specify a number of estimation options. For more information, see "Setting the Estimation Options" on page 185.

Unobserved Components Models

Create an Unobserved Components Model

To create an unobserved components model, complete the following steps:

1. Select **Project** → **Model Repository**. The Model Repository dialog box opens. For more information, see "About the Model Repository" on page 188.

   NOTE: You can also create a new model when you add a model to a series. For more information, see "Add a Model to a Series" on page 196.

2. Click the **Custom Models** tab in the Model Repository dialog box.

3. On the **Custom Models** tab, click **New**. The New Model dialog box opens.

4. In the New Model dialog box, select **Create an entirely new model**.

5. From the **Model type** drop-down list, select **Unobserved Components** as the model type. The Unobserved Components Model dialog box opens.

Figure 10.3 Unobserved Components Model Dialog Box

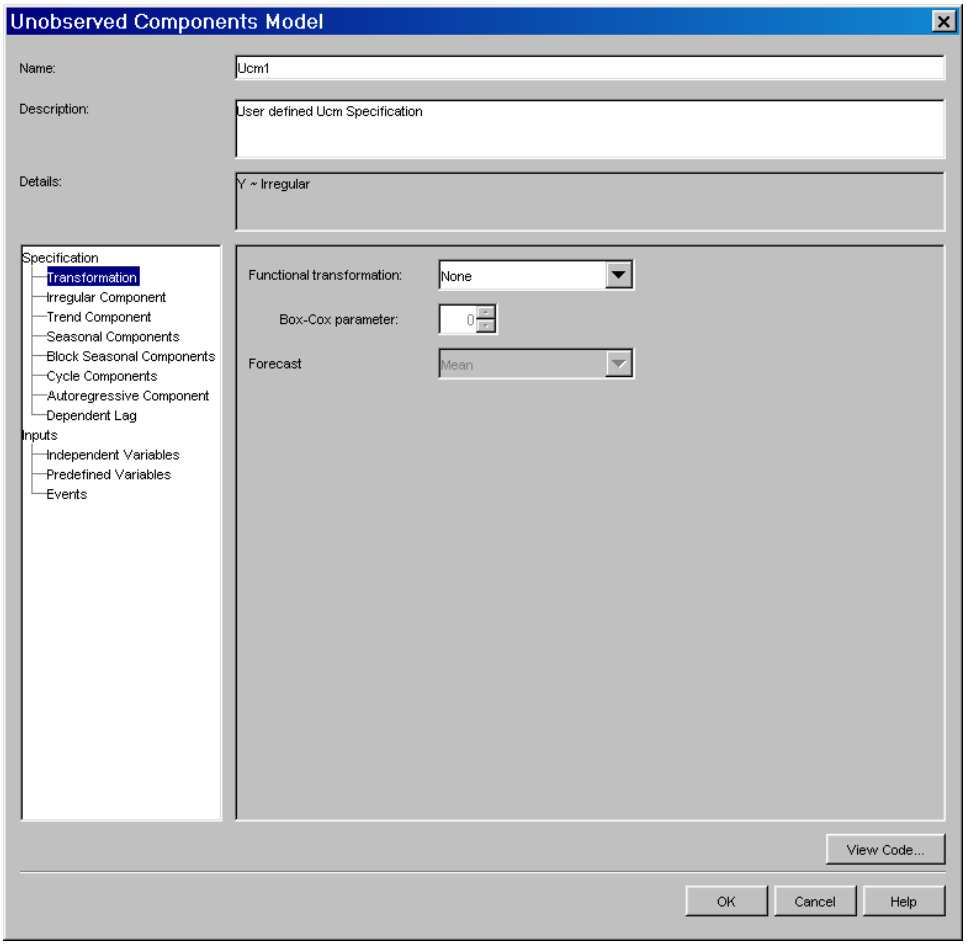

NOTE: The Unobserved Components Model dialog box opens when you edit a model. The options that are available when you are creating or editing a model are the same.

6. Specify a unique name for the model. By default, SAS Forecast Studio automatically generates a unique name for the model.

   NOTE: SAS Forecast Studio has a list of reserved model names. For more information, see "Reserved Model Names" on page 179.

7. (Optional) Specify a description for the new model.

   NOTE: The Details pane shows a representation of the model specification. SAS Forecast Studio generates this representation based on the options that you select in the Unobserved Components Model dialog box. You cannot edit this text.

8. Use the selection pane to specify the following options:

   - transformation
   - irregular component
   - trend component
   - seasonal components

- block-seasonal components
- cycle components
- autoregressive component
- dependent lag
- 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*.

Transformations for Unobserved Components Models

To specify a transformation, complete the following steps:

1. In the Unobserved Components Model dialog box, select **Transformation** in the selection pane.
2. From the **Functional transformation** drop-down list, select the transformation for the model.

   You can select from the following transformations:

   - no transformation (which 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 parameter** box.
3. Specify how to calculate the forecasts. Forecasts can be based on the mean or median. By default, the mean value is used.

Irregular Component for Unobserved Components Models

The irregular component corresponds to the overall random error in the model.

To create an irregular component, complete the following steps:

1. In the Unobserved Components Model dialog box, select **Irregular Component** in the selection pane.
2. Select the **Irregular component** check box.

3. (Optional) 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.

4. If you want the variance to be set at the initial value, select the **Do not perform estimation (noest)** check box. This option is not available unless you specify a value for the initial variance.

Trend Component for Unobserved Components Models

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 trend component, complete the following steps:

1. In the Unobserved Components Model dialog box, select **Trend Component** in the selection pane.

2. Select the **Level** check box to include a level component in the model.

   To specify an initial value for the variance, select the **Define initial variance** check box and specify the initial value.

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

3. (Optional) Select the **Slope** check box to include a slope component in the model.

   To specify an initial value for the variance, select the **Define initial variance** check box and specify the initial value.

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

Seasonal Components for Unobserved Components Models

Working with Seasonal Components

The table in the Seasonal Components pane lists the following properties for each seasonal component:

- the length of the component
- the initial value of the disturbance variance

- whether the disturbance variance should be set at the initial value or if it can be estimated

From this pane, you can complete the following tasks:

- create a seasonal component
- edit a seasonal component
- delete a seasonal component

You can specify a maximum of three seasonal components.

Create a Seasonal Component

To create a seasonal component, complete the following steps:

1. In the Unobserved Components Model dialog box, select **Seasonal Components** in the selection pane.
2. In the Seasonal Components pane, click **New**. The New Seasonal Component dialog box opens.
3. Specify the type of seasonal component. A seasonal component can be one of two types: dummy or trigonometric (which is the default).
4. 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 assign to the Time ID role when you created the project. For more information, see "Assign Data to Roles" on page 44.
5. 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 (noest)** check box. This option is not available unless you specify a value for the initial variance.
7. Click **OK** to create the new seasonal component.

Edit a Seasonal Component

To edit a seasonal component, complete the following steps:

1. In the Unobserved Components Model dialog box, select **Seasonal Components** in the selection pane.
2. In the Seasonal Components pane, select the component from the list and click **Edit**. The Edit Seasonal Component dialog box opens.

3. Specify the new values for the options that you want to change. These options are the same that are available when you create the seasonal component. For more information, see "Create a Seasonal Component" on page 157.
4. Click **OK** to save your changes.

Delete a Seasonal Component

To delete a seasonal component, complete the following steps:

1. In the Unobserved Components Model dialog box, select **Seasonal Components** in the selection pane.
2. In the Seasonal Components pane, select the component that you want to delete and click **Delete**. You are prompted to confirm this deletion. Click **Yes** to delete the seasonal component.

Block-Seasonal Components for Unobserved Components Models

Working with Block-Seasonal Components

The table in the Block Seasonal Components pane lists the following properties for each component:

- the type of component
- the block size
- the number of blocks
- the offset if the first measurement of the series is not at the start of a block
- the initial value for the disturbance variance
- whether the disturbance variance should be set at the initial value or if it can be estimated

From this pane, you can complete the following tasks:

- add a block-seasonal component
- edit a block-seasonal component
- delete a block-seasonal component

For a model, you can specify a maximum of three block-seasonal components.

Create a Block-Seasonal Component

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

1. In the Unobserved Components Model dialog box, select **Block Seasonal Components** in the selection pane.
2. In the Block Seasonal Components pane, click **New**. The New Blockseasonal Component dialog box opens.
3. Specify the type of block-seasonal component. A block-seasonal component can be one of two types: dummy (which is the default) or trigonometric.
4. Specify the block size. The block size can be any integer larger than or equal to 2.
5. Specify the number of blocks. The number of blocks can be any integer larger than or equal to 2.

   NOTE: The block-seasonal component requires the values for the block size and the number of blocks. For example, you can specify the block-seasonal component for an hourly series. For this block-seasonal component, you specify a block size of 24 and the number of blocks as 7. The resulting block-seasonal component is periodic with a period of 168 hours, which is equal to the number of hours in a week. The periodic pattern changes from day to day within a given week, but the pattern is constant within a given day.
6. If the first measurement of the series is not at the start of a block, then select the **Define offset** check box. You can specify the location of the first measurement within the block. This value must be between 1 and the block size.
7. To supply an initial value for the variance, select the **Define initial variance** check box. Specify the initial value for the disturbance variance.
8. If you want the disturbance variance to be set at the initial value, select the **Do not perform estimation (noest)** check box. This option is not available unless you specify a value for the initial variance.

Edit a Block-Seasonal Component

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

1. In the Unobserved Components Model dialog box, select **Block Seasonal Components** in the selection pane.
2. In the Block Seasonal Components pane, select the component from the list and click **Edit**. The Edit Blockseasonal Component dialog box opens.
3. Specify the new values for the options that you want to change. These options are the same that are available when you create the block-seasonal component. For more information, see "Create a Block-Seasonal Component" on page 159.
4. Click **OK** to save your changes.

Delete a Block-Seasonal Component

To delete a block-seasonal component, complete the following steps:

1. In the Unobserved Components Model dialog box, select **Block Seasonal Components** in the selection pane.
2. In the Block Seasonal Components pane, select the component that you want to delete and click **Delete**.

Cycle Components for Unobserved Components Models

Working with Cycle Components

The table in the Cycle Components pane lists the following properties for each cycle component:

- the initial value of the period.
- the initial value of the damping factor.
- the initial value of the disturbance variance.
- whether to use the initial value for the period, damping factor, or disturbance variance, or if this value can be estimated. These values appear in the **Noest** columns.

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

From this pane, you can complete the following tasks:

- add a cycle component
- edit a cycle component
- delete a cycle component

Create a Cycle Component

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

To specify a cycle component, complete the following steps:

1. In the Unobserved Components Model dialog box, select **Cycle Components** in the selection pane.
2. In the Cycle Components pane of the Unobserved Components Model dialog box, click **New**. The New/Edit Cycle Component dialog box opens.

3. 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. By default, the initial value is 3.

4. 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). By default, the initial value is 0.01.

5. To specify an initial value for the disturbance variance parameter that SAS Forecast Studio uses 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. By default, the initial value is 0.

For the period, damping, and variance values, you can select the **Do not perform estimation (noest)** 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 period, damping, or variance.

Edit a Cycle Component

To edit a cycle component, complete the following steps:

1. In the Unobserved Components Model dialog box, select **Cycle Components** in the selection pane.

2. In the Cycle Components pane of the Unobserved Components Model dialog box, select the component from the list and click **Edit**. The New/Edit Cycle Component dialog box opens.

3. Specify the new values for the options that you want to change. These options are the same that are available when you create the cycle component. For more information, see "Create a Cycle Component" on page 160.

4. Click **OK** to save your changes.

Delete a Cycle Component

To delete a cycle component, complete the following steps:

1. In the Unobserved Components Model dialog box, select **Cycle Components** in the selection pane.

2. In the Cycle Components pane of the Unobserved Components Model dialog box, select the component that you want to delete and click **Delete**.

Autoregressive Components for Unobserved Components Models

To specify an autoregressive component, complete the following steps:

1. In the Unobserved Components Model dialog box, select **Autoregressive Component** in the selection pane.
2. Select the **Autoregressive component** check box.
3. (Optional) 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.
4. (Optional) 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 (noest)** 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 Unobserved Components Models

For unobserved components models, you can specify the forecast variable lags to be included as predictors in the model.

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

1. In the Unobserved Components Model dialog box, select **Dependent Lag** in the selection pane.
2. Select the **Lag dependent variable** check box.
3. Specify the lag values. By default, the number of lags is 0.

   To specify a lag, use the following syntax:

   - For a dependent lag with degrees from 1 to an integer value, use *integer* as the syntax for the **Lags** option. This syntax creates a lag of 1, . . . ,*integer*.
   - For the lag of a specific integer, use (*integer*) as the syntax for the **Lags** option. This syntax creates a lag of (*integer*).
   - To include a multiplier, use (*integer, . . . , integer*)*m* as the syntax for the **Lags** option. This syntax creates a lag of (*integer*m, . . . , integer*m*).
   - For a factored lag, use (*integer . . . integer*)...(*integer, . . . , integer*) as the syntax for the **Lags** option. This syntax creates a lag of (*integer, . . . , integer*)...(*integer, . . . , integer*).

The following table shows some examples of this syntax:

Model description	Value for lags option	Result
Lag from 1 to an integer value	5	(1 2 3 4 5)
Lag for a specific integer Lag with multiplier	(5) (1 2 3)(2 7)12	(5) (1 2 3)(24 84)
Factored lag	(1 2 3)(2 7)	(1 2 3)(2 7)

4. 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.

5. To set the parameter values to the initial value, select the **Do not perform estimation (noest)** check box. This option fixes the values of the component parameters to those specified. This option is not available unless you specify an initial value for phi.

Inputs in Unobserved Components Models

By default, no independent variables, predefined variables, or events are used as inputs in the model. You can add and remove inputs from the model.

For more information, see the following topics:

- "Include Independent Variables in User-Defined Models" on page 182
- "Include Predefined Variables in User-Defined Models" on page 182
- "Include Events in User-Defined Models" on page 183

Exponential Smoothing Models

Overview of Exponential Smoothing Models

Exponential smoothing is a forecasting technique that uses exponentially declining weights to produce a weighted moving average of time series values.

Using SAS Forecast Studio, you can create many types of exponential smoothing models. For more information about the Exponential Smoothing model, see the HPFESMSPEC procedure in the *SAS High-Performance Forecasting User's Guide*.

Create an Exponential Smoothing Model

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

1. Select **Project** → **Model Repository**. The Model Repository dialog box opens. For more information, see "About the Model Repository" on page 188.

   NOTE: You can also create a new model when you add a model to a series. For more information, see "Add a Model to a Series" on page 196.

2. Click the **Custom Models** tab in the Model Repository dialog box.

3. On the **Custom Models** tab, click **New**. The New Model dialog box opens.

4. In the New Model dialog box, select **Create an entirely new model**.

5. From the **Model type** drop-down list, select **Exponential smoothing** as the model type. The Exponential Smoothing Model dialog box opens.

Figure 10.4 Exponential Smoothing Model Dialog Box

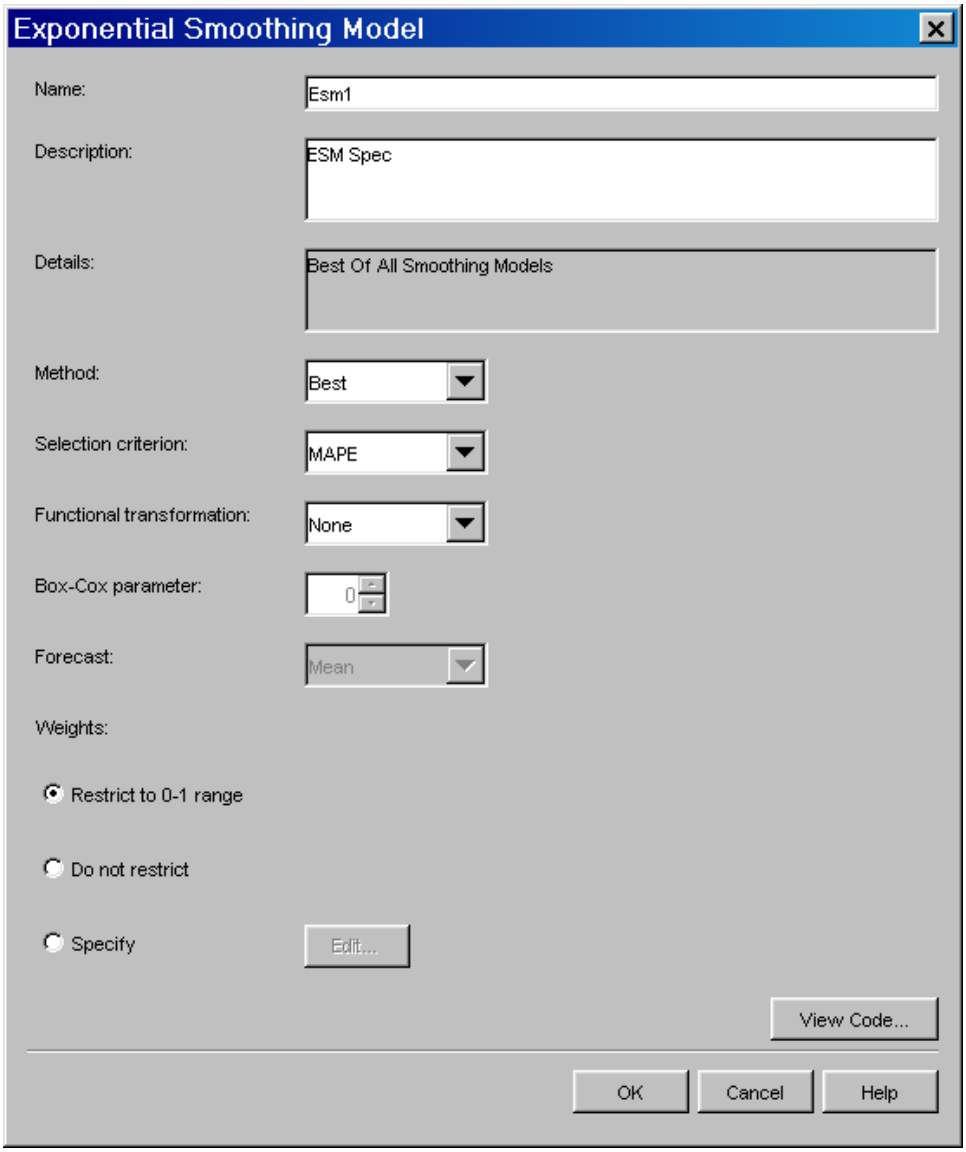

NOTE: The Exponential Smoothing Model dialog box also opens when you edit a model. The options that are available when you are creating or editing a model are the same.

6. Specify a unique name for the model. By default, SAS Forecast Studio automatically generates a unique name for the model.

   NOTE: SAS Forecast Studio has a list of reserved model names. For more information, see "Reserved Model Names" on page 179.

7. (Optional) Specify a description for the new model.

   NOTE: The Details pane shows a representation of the model specification. SAS Forecast Studio generates this representation based on the options that you select in the Exponential Smoothing Model dialog box. You cannot edit this text.

8. From the **Method** drop-down list, select the forecasting method to be used to forecast the time series. You can choose from the following methods:

   - **Simple** - Simple (Single) Exponential Smoothing
   - **Double (Brown)** - Double (Brown) Exponential Smoothing model
   - **Linear (Holt)** - Linear (Holt) Exponential Smoothing model
   - **Damp Trend** - Damped Trend Exponential Smoothing model
   - **Seasonal** - Seasonal Smoothing model
   - **Winters** - Winters (Multiplicative) model
   - **Additive Winters** - Winters (Additive) model
   - **BestN** - Best Candidate Nonseasonal Smoothing model
   - **Bests** - Best Candidate Seasonal Smoothing model
   - **Best** - Best Candidate Smoothing model

   NOTE: The **Winters** option is not available if the series that you selected contains negative values.

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

9. Select the selection criterion to use. The default selection criterion is the project default.

   NOTE: This option is available only when either of the following is true:

   - You selected **Best**, **BestN**, or **Bests** as the forecasting method.
   - In the **Functional transformation** drop-down list, you select **Auto**, which means that the transformation for the dependent series is automatically chosen based on the model selection criteria.

10. From the **Functional transformation** drop-down list, select the dependent series transformation to be applied to the time series. You can select from the following transformations:

    - no transformation (which 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 parameter** box.
    - An automatically chosen transformation based on the model selection criteria.

    NOTE: You can specify a transformation only when the values of the series are positive.

11. Specify how to calculate the forecasts. Forecasts can be based on the mean or median. By default, the mean value is used.

12. Specify whether to restrict all weights to the 0 to 1 range. You can also specify the smoothing weights. For more information, see "Smoothing Weights" on page 186.

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

Intermittent Demand Models

Create an Intermittent Demand Model

To create an intermittent demand model, complete the following steps:

1. Select **Project** → **Model Repository**. The Model Repository dialog box opens. For more information, see "About the Model Repository" on page 188.

   NOTE: You can also create a new model when you add a model to a series. For more information, see "Add a Model to a Series" on page 196.

2. Click the **Custom Models** tab in the Model Repository dialog box.

3. On the **Custom Models** tab, click **New**. The New Model dialog box opens.

4. In the New Model dialog box, select **Create an entirely new model**.

5. From the **Model type** drop-down list, select **Intermittent demand** as the model type. The Intermittent Demand Model dialog box opens.

Figure 10.5 Intermittent Demand Model Dialog Box

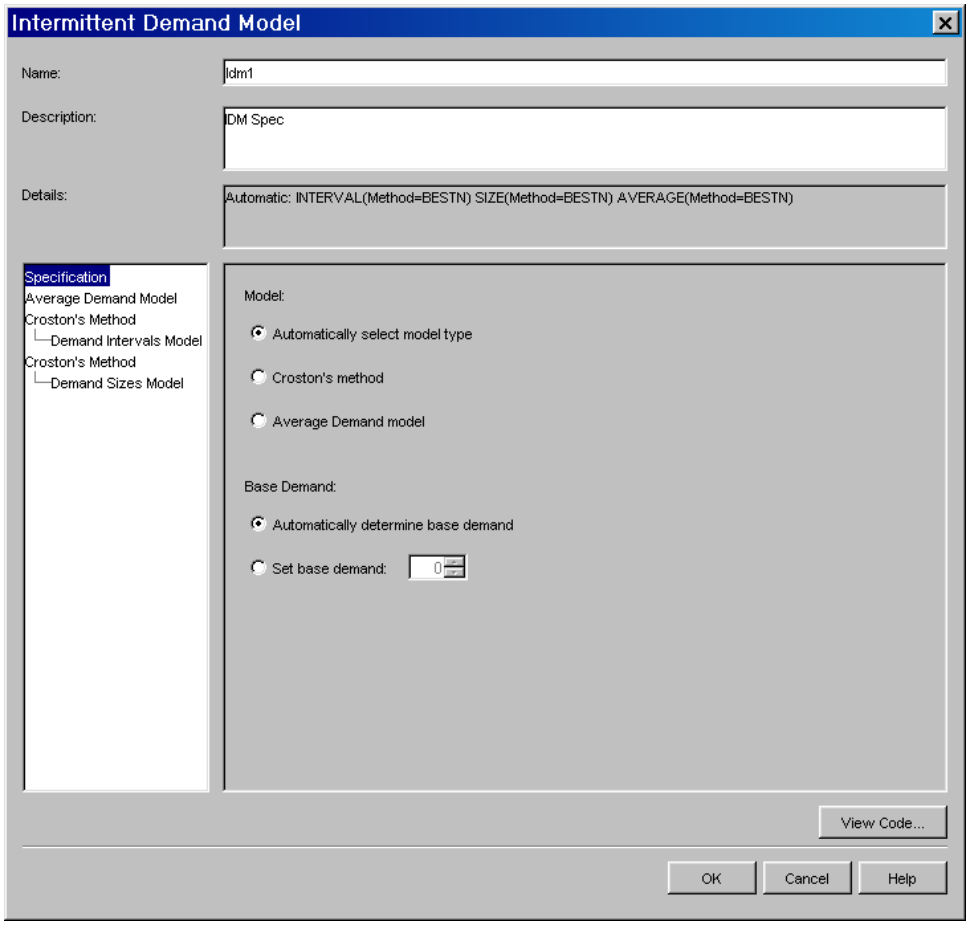

NOTE: The Intermittent Demand Model dialog box opens when you edit a model. The options that are available when you are creating or editing a model are the same.

6. Specify a unique name for the model. By default, SAS Forecast Studio automatically generates a unique name for the model.

   NOTE: SAS Forecast Studio has a list of reserved model names. For more information, see "Reserved Model Names" on page 179.

7. (Optional) Specify a description for the new model.

   NOTE: The Details pane shows a representation of the model specification. SAS Forecast Studio generates this representation based on the options that you select in the Intermittent Demand Model dialog box. You cannot edit this text.

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

   - specification
   - average demand model, demand intervals model, or demand sizes model

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

Specification Options for Intermittent Demand Models

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

- the model type. You can choose from the following options:
  - **Automatically select model type** - specifies that SAS Forecast Studio should automatically determine the model type.
  - **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.

Model Options for Intermittent Demand Models

For an Intermittent Demand model, you can create an average demand model, a Croston's method demand intervals model, and a Croston's method demand sizes model.

For the model that you want to create, select **Average Demand**, **Demand Intervals Model**, or **Demand Sizes Model** in the selection pane to access the following options:

- the forecasting method to use to forecast the time series. You can choose from the following methods:
  - **BestN** - Best Candidate Nonseasonal Smoothing model
  - **Simple** - Simple (Single) Exponential Smoothing (which is the default)
  - **Double (Brown)** - Double (Brown) Exponential Smoothing model
  - **Linear (Holt)** - Linear (Holt) Exponential Smoothing model
  - **Damp Trend** - Damped Trend Exponential Smoothing model

  For more information about these models, see the "Forecasting Process Details" chapter in the *SAS High-Performance Forecasting User's Guide*.
- the selection criterion to use. The default selection criterion is the project default.

  NOTE: This option is available only when either of the following is true:

  - You selected **BestN** as the forecasting method.
  - In the **Functional transformation** drop-down list, you select **Auto**, which means that the transformation for the dependent series is automatically chosen based on the model selection criteria.

- the functional transformation to apply to the time series. You can select from the following transformations:
  - no transformation (which 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 parameter** box.
  - an automatically chosen transformation based on the model selection criteria.
- 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. You can also choose to specify the smoothing weights. For more information, see "Smoothing Weights" on page 186.

Multiple Regression Models

Create a Multiple Regression Model

To create a multiple regression model, complete the following steps:

1. Select **Project** → **Model Repository**. The Model Repository dialog box opens. For more information, see "About the Model Repository" on page 188.

   NOTE: You can also create a new model when you add a model to a series. For more information, see "Add a Model to a Series" on page 196.
2. Click the **Custom Models** tab in the Model Repository dialog box.
3. On the **Custom Models** tab, click **New**. The New Model dialog box opens.
4. In the New Model dialog box, select **Create an entirely new model**.
5. From the **Model type** drop-down list, select **Multiple Regression** as the model type. The Multiple Regression Model dialog box opens.

Figure 10.6 Multiple Regression Model Dialog Box

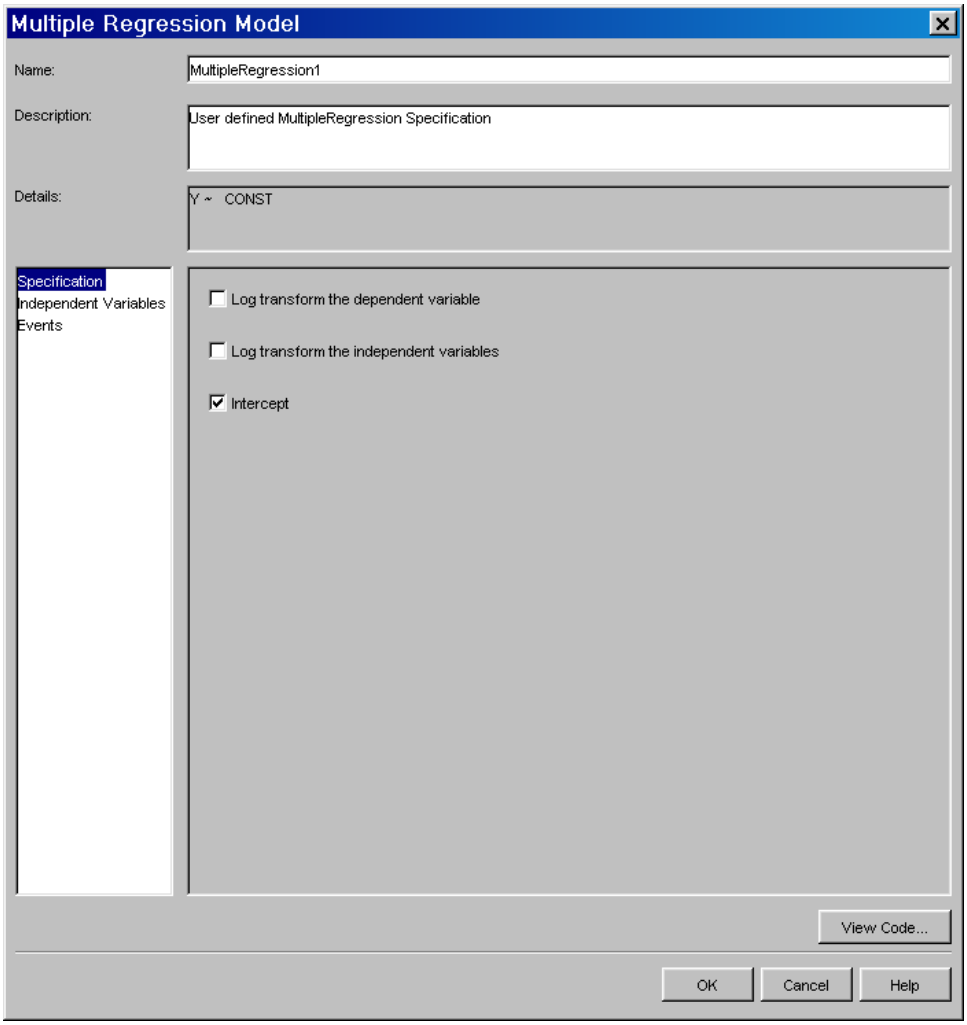

NOTE: The Multiple Regression Model dialog box also opens when you edit a model. The options that are available when you are creating or editing a model are the same.

6. Specify a unique name for the model. By default, SAS Forecast Studio automatically generates a unique name for the model.

   NOTE: SAS Forecast Studio has a list of reserved model names. For more information, see "Reserved Model Names" on page 179.

7. (Optional) Specify a description for the new model.

   NOTE: The Details pane shows a representation of the model specification. SAS Forecast Studio generates this representation based on the options that you select in the Multiple Regression Model dialog box. You cannot edit this text.

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

   - specification
   - independent variables

- events

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

Specification Options for Multiple Regression Models

In the Multiple Regression Model dialog box, select **Specification** in the selection pane to access the following options:

- **Log transform the dependent variable** - specifies whether to log transform the dependent variable.
- **Log transform the independent variables** - specifies whether to log transform the independent variables.
- **Intercept** - specifies whether to include the intercept in the model. By default, the intercept is included.

Independent Variables for Multiple Regression Models

By default, none of the variables that you assign to the Independent variables role are inputs in the model. However, using the Multiple Regression Model dialog box, you can specify to include these independent variables in the model. For more information, see "Include Independent Variables in User-Defined Models" on page 182.

Events for Multiple Regression Models

By default, none of the events in the project are inputs in the model. However, using the Multiple Regression Model dialog box, you can specify to include these events in the model. You can also manage the events for this model and specify the input series options for an event. For more information, see "Include Events in User-Defined Models" on page 183.

Moving Average Models

Overview of Moving Average Models

Using SAS Forecast Studio, you can create a moving average model. The formula for the moving average model with width k is $y_t = [y_{(t-1)} + \ldots + y_{(t-k)}]/k + error$.

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

Create a Moving Average Model

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

1. Select **Project** → **Model Repository**. The Model Repository dialog box opens. For more information, see "About the Model Repository" on page 188.

   NOTE: You can also create a new model when you add a model to a series. For more information, see "Add a Model to a Series" on page 196.

2. Click the **Custom Models** tab in the Model Repository dialog box.

3. On the **Custom Models** tab, click **New**. The New Model dialog box opens.

4. In the New Model dialog box, select **Create an entirely new model**.

5. From the **Model type** drop-down list, select **Moving Average** as the model type. The Moving Average Model dialog box opens.

Figure 10.7 Moving Average Model Dialog Box

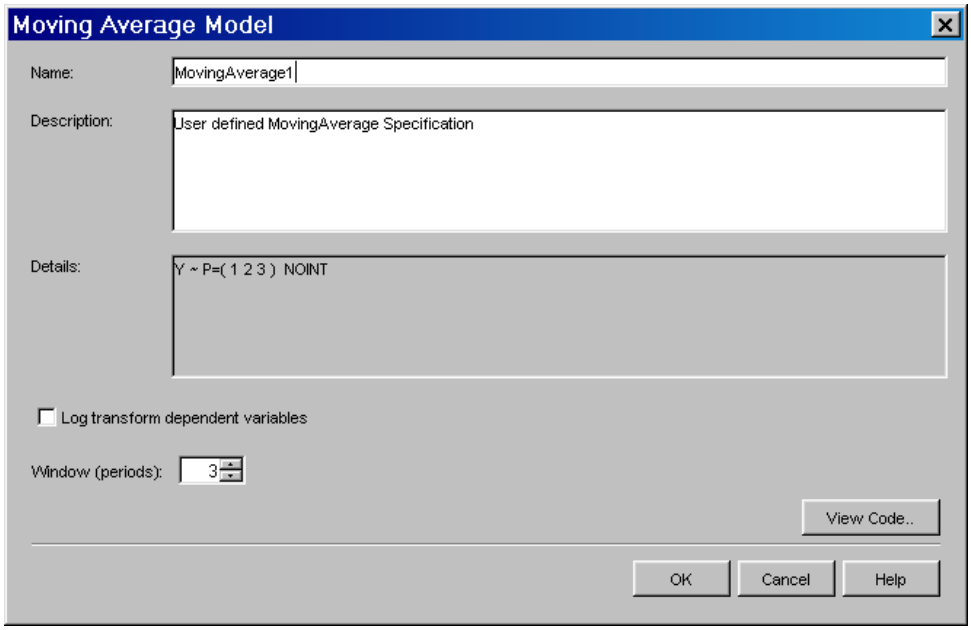

NOTE: The Moving Average Model dialog box also opens when you edit a model. The options that are available when you are creating or editing a model are the same.

6. Specify a unique name for the model. By default, SAS Forecast Studio automatically generates a unique name for the model.

   NOTE: SAS Forecast Studio has a list of reserved model names. For more information, see "Reserved Model Names" on page 179.

7. (Optional) Specify a description for the new model.

   NOTE: The Details pane shows a representation of the model specification. SAS Forecast Studio generates this representation based on the options that you select in the Moving Average Model dialog box. You cannot edit this text.

8. Select the **Log transform dependent variables** check box to log transform the dependent variable.

9. 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**.

Curve Fitting Models

Overview of Curve Fitting Models

Curve fitting models enable you to identify trends and relationships in your time series data. Using SAS Forecast Studio, you can create a curve fitting model with a linear or quadratic trend.

Create a Curve Fitting Model

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

1. Select **Project** → **Model Repository**. The Model Repository dialog box opens. For more information, see "About the Model Repository" on page 188.

   NOTE: You can also create a new model when you add a model to a series. For more information, see "Add a Model to a Series" on page 196.

2. Click the **Custom Models** tab in the Model Repository dialog box.

3. On the **Custom Models** tab, click **New**. The New Model dialog box opens.

4. In the New Model dialog box, select **Create an entirely new model**.

5. From the **Model type** drop-down list, select **Curve Fitting** as the model type. The Curve Fitting Model dialog box opens.

Figure 10.8 Curve Fitting Model Dialog Box

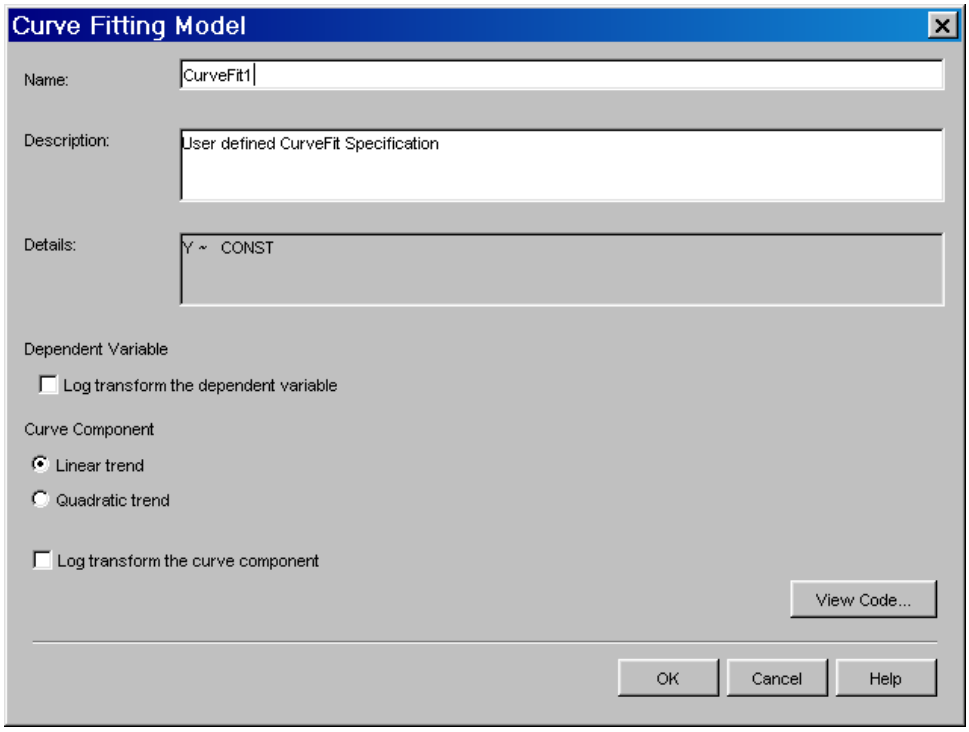

NOTE: The Curve Fitting Model dialog box also opens when you edit a model. The options that are available when you are creating or editing a model are the same.

6. Specify a unique name for the model. By default, SAS Forecast Studio automatically generates a unique name for the model.

   NOTE: SAS Forecast Studio has a list of reserved model names. For more information, see "Reserved Model Names" on page 179.

7. (Optional) Specify a description for the new model.

   NOTE: The Details pane shows a representation of the model specification. SAS Forecast Studio generates this representation based on the options that you select in the Curve Fitting Model dialog box. You cannot edit this text.

8. Select the **Log transform the dependent variable** check box if you want to log transform the dependent variable.

9. Select the curve component for the model. 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 $X_t = (t - c)^2$.

10. Select the **Log transform the curve component** check box if you want to log transform the curve component.

11. Click **OK** to save the model.

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

Random Walk Models

Overview of Random Walk Models

Using SAS Forecast Studio, 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

Create a Random Walk Model

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

1. Select **Project** → **Model Repository**. The Model Repository dialog box opens. For more information, see "About the Model Repository" on page 188.

   NOTE: You can also create a new model when you add a model to a series. For more information, see "Add a Model to a Series" on page 196.

2. Click the **Custom Models** tab in the Model Repository dialog box.

3. On the **Custom Models** tab, click **New**. The New Model dialog box opens.

4. In the New Model dialog box, select **Create an entirely new model**.

5. From the **Model type** drop-down list, select **Random walk** as the model type. The Random Walk Model dialog box opens.

Figure 10.9 Random Walk Model Dialog Box

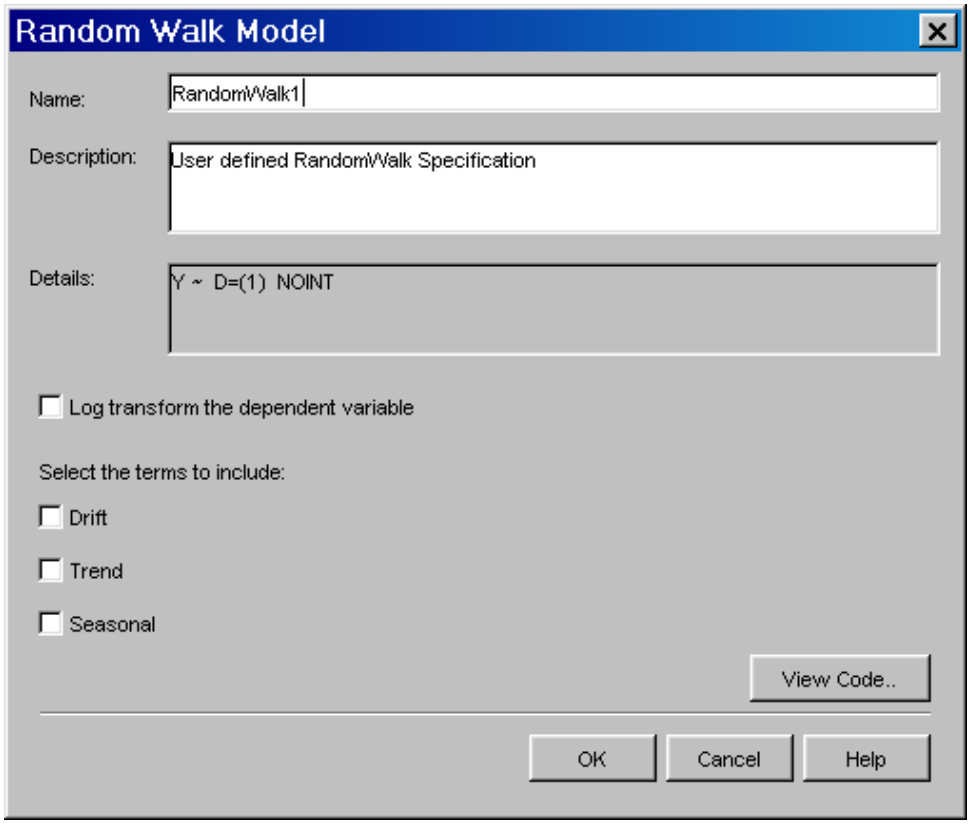

NOTE: The Random Walk Model dialog box also opens when you edit a model. The options that are available when you are creating or editing a model are the same.

6. Specify a unique name for the model. By default, SAS Forecast Studio automatically generates a unique name for the model.

   NOTE: SAS Forecast Studio has a list of reserved model names. For more information, see "Reserved Model Names" on page 179.

7. (Optional) Specify a description for the new model.

   NOTE: The Details pane shows a representation of the model specification. SAS Forecast Studio generates this representation based on the options that you select in the Random Walk Model dialog box. You cannot edit this text.

8. Select the **Log transform the dependent variable** check box to log transform the dependent variable. When you select this option, the model uses the data for the log transformed series.

9. Select the terms to include in the model. You can choose from the following options:

   - **Drift** - includes a drift term in the model.
   - **Trend** - includes a trend term in the model.
   - **Seasonal** - includes a seasonal term in the model.

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

Common Model Options

Reserved Model Names

SAS Forecast Studio uses the following keywords in the names of the automatically generated models:

- LEAF
- TOP
- HP

These words are reserved for SAS Forecast Studio models. If the name of the model that you are creating contains these keywords, then a warning message appears, and you are prompted to select a different name.

NOTE: If you are copying a system-generated model that contains one of these keywords, then the copied name can contain one of these keywords. For example, if you copy a model named LEAF_*n*, then SAS Forecast Studio names the copied model LEAF_Copy*n*, where *n* is the lowest available integer value. Because this name is unique, SAS Forecast Studio allows this name. For more information about copying a system-generated model, see "Create a New Custom Model Based on a Default Model" on page 190.

Autoregressive, Differencing, and Moving Average Orders

What Models Use Autoregressive, Differencing, and Moving Average Orders?

For ARIMA and Subset (factored) ARIMA models, you can specify the autoregressive, differencing, and moving average orders for the model. How you specify these orders differs slightly between the models. For the ARIMA model, you can specify integer values for the orders that make up the nonseasonal and seasonal components. For Subset (factored) ARIMA models, you have more flexibility. SAS Forecast Studio uses the syntax from PROC HPFARIMASPEC to specify this order. For more information about the underlying code, see the HPFARIMASPEC procedure in the *SAS High-Performance Forecasting User's Guide*.

For ARIMA, Subset (factored) ARIMA, and Unobserved Components models, you can also specify the differencing order for inputs, such as events.

Autoregressive Orders

The **Autoregressive (p)** option specifies the autoregressive part of the model.

To specify an autoregressive order, use the following syntax:

- For an autoregressive model with degrees from 1 to an integer value, use p = *integer*. This syntax creates a model of P = (1, . . . ,*integer*).
- For an autoregressive model with parameters at specified lags, use p = (*integer*). This syntax creates a model of P = (*integer*).
- For an autoregressive model with a seasonal component, use p = (*integer, . . . , integer*)(*integer, . . . , integer*)s, where *s* is a placeholder for the seasonal component. This placeholder value is supplied later. You can also specify an integer value for this seasonal component. This syntax creates a model of P = (*integer, . . . , integer*)(*integer, . . . , integer*)s.
- For an autoregressive model with multiple factors, use p = (*integer, . . . , integer*)...(*integer, . . . , integer*). This sytanx creates a model of P = (*integer, . . . , integer*)...(*integer, . . . , integer*).

The following table shows some examples of this syntax:

Model description	Value of p	Result in Details pane
Autoregressive model with degrees from 1 to an integer value	p = 5	P = (1 2 3 4 5)
Autoregressive model with parameters at specified lags	p = (5)	P = (5)
Autoregressive model with a seasonal component with "s" as a placeholder	p = (1 2 3)(2 7)s	P = (1 2 3)(2 7)s
Autoregressive model with a seasonal component of 12	p = (1 2 3)(2 7)12	P = (1 2 3)(24 84)
Autoregressive model with three factors	p = (1 2)(1)(1)s	P = (1 2)(1)(1)s

Differencing Orders

The **Differencing (d)** option specifies the differencing orders for the dependent series.

For differencing orders, you must enclose the value or values that you specify in parentheses. Unlike the autoregressive and moving average orders, you must include all orders that you specify in a single list, and you can specify an integer value multiple times in the same list. You can use placeholders (such as "s") to specify seasonal differencing.

To specify a differencing order, use the following syntax:

- For a single differencing order, use D = (*integer*). This syntax creates a model of D = (*integer*).
- For multiple differencing orders, use d = (*integer, . . . , integer*). This syntax creates a model of D = (*integer . . . integer*).
- For a differencing order with a seasonal component, use d = (*integer, . . . , s*). This syntax creates a model of D = (*integer, . . . , s*), where *s* is a placeholder for the seasonal order. This placeholder value is supplied later.

The following table includes examples of this syntax:

Model description	Value of d	Result in Details pane
Single differencing order	d = 5	Incorrect syntax; no order specified
Single differencing order	d = (5)	D = (5)
Multiple differencing orders	d = (1 1 3 5)	D = (1 1 3 5)
Differencing order with a seasonal component	d = (1 s s)	D = (1 s s)

Moving Average Orders

The **Moving average (q)** option specifies the moving average part of the model. The syntax of the moving average order is identical to the syntax of the autoregressive order.

To specify a moving average order, use the following syntax:

- For a moving average model with degrees from 1 to an integer value, use q = *integer*. This syntax creates a model of Q = (1, . . . ,*integer*).
- For a moving average model with parameters at specified lags, use q = (*integer*). This syntax creates a model of Q = (*integer*).
- For a moving average model with a seasonal component, use q = (*integer* . . . *integer*)(*integer* . . . *integer*)s, where *s* is a placeholder for the seasonal component. This placeholder value is supplied later. You can also specify an integer value for this seasonal component. This syntax creates a model of Q = (*integer, . . . , integer*)(*integer, . . . , integer*)s.
- For a moving average model with multiple factors, use q = (*integer, . . . , integer*)...(*integer, . . . , integer*). This syntax creates a model of Q = (*integer, . . . , integer*)...(*integer, . . . , integer*).

The following table shows some examples of this syntax:

Model description	Value of q	Result in Details pane
Moving average model with degrees from 1 to an integer value	q = 5	Q = (1 2 3 4 5)
Moving average model with parameters at specified lags	q = (5)	Q = (5)
Moving average model with a seasonal component with "s" as a placeholder	q = (1 2 3)(2 7)s	Q = (1 2 3)(2 7)s
Moving average model with a seasonal component of 12	q = (1 2 3)(2 7)12	Q = (1 2 3)(24 84)
Moving average model with three factors	q = (1 2)(1)(1)s	Q = (1 2)(1)(1)s

Example: Specifying the Orders for the Airline Model

Using the syntax that is available in SAS Forecast Studio, you can create the familiar Airline model, ARIMA(0,1,1)(0,1,1)s, by specifying the differencing and moving average orders in the Subset (factored) ARIMA model. You can specify this model in any of the following ways:

Model description	Value of d	Value of q	Result
Season length is specified as a wildcard	d = (1 s)	q = (1)(1)s	D = (1 s) Q = (1)(1)s
Monthly seasonality is specified	d = (1 12)	q = (1)(1)12	D = (1 12) Q = (1)(12)
Alterative way to specify monthly seasonality	d = (1 12)	q = (1)(12)	D = (1 12) Q = (1)(12)

Including Independent Variables, Predefined Variables, Events, and Outliers

Include Independent Variables in User-Defined Models

By default, none of the variables that you assign to the Independent variables role are inputs in the model. You assign independent variables from the New Project Wizard when you create the project or when you use the Variables dialog box. For more information, see "Modify Variable Assignments" on page 49.

In the ARIMA, Subset (factored) ARIMA, and Unobserved Components models, you can choose to include these independent variables in the model.

To include an independent variable in one of these models, complete the following steps:

1. In the selection pane of the ARIMA Model, Subset ARIMA Model, or Unobserved Components Model dialog box, select **Independent Variables**.
2. For each independent variable that you want to include in the model, select the check box in the **Include in Model** column.

For ARIMA, Subset ARIMA, and Unobserved Component models, you can also specify the transfer function options if you include only one independent variable in the model. To specify the transfer function options, click **Transfer Function**. The Transfer Function dialog box opens. For more information, see "Specifying the Transfer Function Options" on page 184.

NOTE: You cannot specify the transfer function options for an independent variable in a multiple regression model.

Include Predefined Variables in User-Defined Models

By default, none of the predefined variables from the input data set are inputs in the model.

In the ARIMA, Subset (factored) ARIMA, and Unobserved Components models, you can choose to include these predefined variables in the model.

To include a predefined variable in the model, complete the following steps:

1. In the selection pane of the ARIMA Model, Subset ARIMA Model, or Unobserved Components Model dialog box, select **Predefined Variables**.
2. For each predefined variable that you want to include in the model, select the check box in the **Include in Model** column.

Include Outlier Variables in User-Defined Models

By default, any outliers that were automatically detected during the creation of the model are inputs in the model. In the ARIMA or Subset (factored) ARIMA models, you can choose to delete these outliers from the model.

To delete an outlier, complete the following steps:

1. In the selection pane of the ARIMA Model or Subset ARIMA Model dialog box, select **Outlier Variables**.
2. Select the outlier that you want to delete and click **Delete**.

NOTE: You cannot add an outlier back to the model after you have deleted it. The next time that you open the Outlier Variables pane the outlier that you deleted does not appear.

Include Events in User-Defined Models

By default, none of the events in the project are inputs in the model.

In the ARIMA, Subset (factored) ARIMA, Unobserved Components, and Multiple Regression models, you can choose to include these events in the model.

To include an event in the model, complete the following steps:

1. In the selection pane of the ARIMA Model, Subset ARIMA Model, Unobserved Components Model, or Multiple Regression Model dialog box, select **Events**.
2. For each event variable that you want to include in the model, select the check box in the **Include in Model** column.

On the Events pane, you can also perform the following tasks:

- To create a new event, click **New**. The Event Properties dialog box opens. For more information, see "Create an Event" on page 73.
- If you included only one event in the model, then you can specify a transfer function by clicking **Transfer Function**. The Transfer Function dialog box opens. For more information, see "Specifying the Transfer Function Options" on page 184.

Specifying the Transfer Function 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 Transfer Function dialog box.

From this dialog box, you can specify the following information:

Functional transformation
specifies the transformation to be applied to the time series. You can select from the following transformations:

- no transformation (which 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 parameter** box.

NOTE: The transformation is applicable only if the input variable is positive.

Lagging order
specifies the delay (or lag) for the input series.

Differencing orders
specifies the differencing orders for the input series. How you specify the differencing order depends on the options that are available.

- If the **Simple differencing** and **Seasonal differencing** options are available, then you can specify the integer value to use for these differencing orders.
- If the **Differencing orders** option is available, then you must specify the syntax for the differencing order. For more information about the syntax, see "Autoregressive, Differencing, and Moving Average Orders" on page 179.

NOTE: For ARIMA and Subset (factored) ARIMA models, the differencing order for new events is the same as the differencing order that is specified for the model. You can view the differencing order for the model in the Details pane of the ARIMA Model dialog box or the Subset ARIMA Model dialog box.

Numerator and Denominator polynomials
If you are creating an ARIMA or subset ARIMA model, specify whether you want an ordinary or dynamic regression.

NOTE: These options do not apply to Unobserved Components models and are not available.

How you specify these polynomials depends on the options that are available.

- If the **Numerator factors** and **Denominator factors** options are available, then you can specify the integer value to use for the simple and seasonal factors.

- If the **Numerator polynomial** and **Denominator polynomial** options are available, then you must use the following syntax to specify the numerator (NUM) and denominator (DEN) polynomial of the transfer function.

  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 autoregressive (p) option in "Autoregressive Orders" on page 179.

After you have finished specifying these options, click **OK** to continue creating or editing the model.

Setting the Estimation Options

In the ARIMA and Subset (factored) ARIMA models, you can specify the estimation options to use.

To specify the estimation options, complete the following steps:

1. In the ARIMA Model or Subset ARIMA Model dialog box, select **Estimations** in the selection pane.
2. In the Estimations pane, the following options are available:

   - **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.

Smoothing Weights

For exponential smoothing and intermittent demand models, you can specify the smoothing weights to use.

To specify a smoothing weight, complete the following steps:

1. In the Exponential Smoothing Model or Intermittent Demand Model dialog box, select **Specify** as the weight and click **Edit**. The Edit Weights dialog box opens.
2. 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.
3. (Optional) Specify the following general settings:
   - **Do not perform estimation (noest)** - specifies that 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.
   - **Restrict weights to stable values** - specifies to restrict the weights to stable values.

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

Chapter 11

Working with Models

Contents

Managing Models

About the Model Repository

To view all of the models that are available for the current project, select **Project→ Model Repository**.

The Model Repository dialog box shows all of the models that ship with SAS Forecast Studio and any user-defined models. These models are stored in separate model repositories. For more information, see "How Models Are Stored" on page 189.

NOTE: The list of models shown in the Model Repository dialog box is different from the model selection list in the Model View that shows only the models for the current series. For more information, see "Using the Model Selection List" on page 191.

Figure 11.1 Model Repository Dialog Box

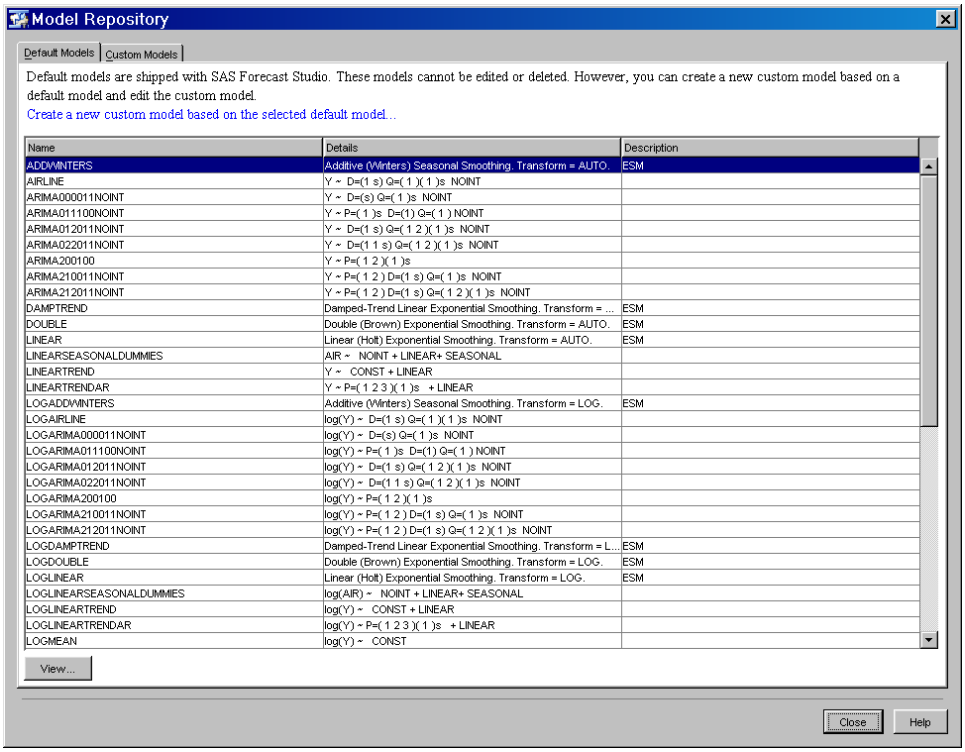

The Model Repository dialog box contains the following tabs:

Default Models
: lists the models that ship with SAS Forecast Studio. You cannot edit or delete these models. However, you can copy a default model and customize it to meet your needs. You can also view the model specification for the selected model. For more information, see "Working with Default Models" on page 189.

Custom Models
: lists the models that you create in SAS Forecast Studio. When you create a model for a particular series, that model is automatically added to this tab. After a model is added, you can apply it to other series in the project.

  Because these are user-defined models, you can create new models, edit an existing model, copy a model, or delete a model from this tab. For more information, see "Working with Custom Models" on page 190.

How Models Are Stored

The following table describes each of the model repositories that are used by SAS Forecast Studio:

Model Repository	Description
project level model repository	This model repository is created when you create the project. Any model specifications that you define are saved in projectmodrep.sas7bcat that is in the SAS Forecast Studio home directory. This is the only model repository that you can modify. The model specifications in this repository take precedence over model specifications in the other model repositories.
HPFDFLT model repository	This model repository is installed with SAS High-Performance Forecasting. You cannot modify the content in this repository.
model repositories for each level in the hierarchy	SAS Forecast Studio automatically creates a model repository for each level in the hierarchy or one model repository if your project does not have a hierarchy. These repositories contain the system-generated model specifications. You cannot modify the content in this model repository. SAS Forecast Studio identifies these models by using specific model names. You cannot use these model names for any user-created models. For more information, see "Reserved Model Names" on page 179.

Working with Default Models

View the Model Specification for a Default Model

To view the model specification for a default model, complete the following steps:

1. Select **Project** → **Model Repository**. The Model Repository dialog box opens.
2. Click the **Default Models** tab.
3. Select the default model from the table and click **View**. A read-only version of the model specification dialog box for that model opens. For more information about the options that are available for each model type, see Chapter 10, "Creating User-Defined Models."
4. When you have finished reviewing the options in the model specification dialog box, click **Cancel** to close this dialog box.

Create a New Custom Model Based on a Default Model

To create a new custom model based on a default model, complete the following steps:

1. Select **Project** → **Model Repository**. The Model Repository dialog box opens.
2. Click the **Default Models** tab.
3. Click **Create a new custom model based on a default model**. The Select Default Model dialog box opens.
4. Select the default model that you want to use as the basis for the custom model and click **OK**. The model specification dialog box for that model opens. By default, the name of the model is *Existing-model-name*COPY*n*, where *n* is the lowest available integer.
5. Edit the model options and click **OK**. For more information about the options that are available for each model, see Chapter 10, "Creating User-Defined Models."

SAS Forecast Studio saves the new model and displays it on the **Custom Models** tab in the Model Repository dialog box.

Working with Custom Models

From the **Custom Models** tab in the Model Repository 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 User-Defined Model" on page 144.
- 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 197.

- 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.

  You can then edit the model by using the model specification dialog box that opens. For more information, see "Edit a Model" on page 197.

- To delete a model, select the model from the table and click **Delete**. You are prompted to confirm this deletion. To delete the model, click **Yes**.

Using the Model Selection List

What Is the Model Selection List?

The model selection list is shown in the Model View. The model selection list enables you to compare the statistics of fit for different models and to select the model that you want to use for the current series. To view or save the code for the model selection list, click  .

The model selection list displays the generated, external, and custom models that SAS Forecast Studio selected as the models for generating forecasts. You cannot edit the model specification for generated, default, or external models, so these models are listed in the model selection list as read-only.

By default, the best-fitting model appears in bold. It is the forecast model that is displayed in the forecast plot and data table in the Forecasting View. If you select a different model from the **Selected model** drop-down list, then the selected model appears in bold.

The following is an example of a model selection list:

Figure 11.2 Model Selection List in Model View

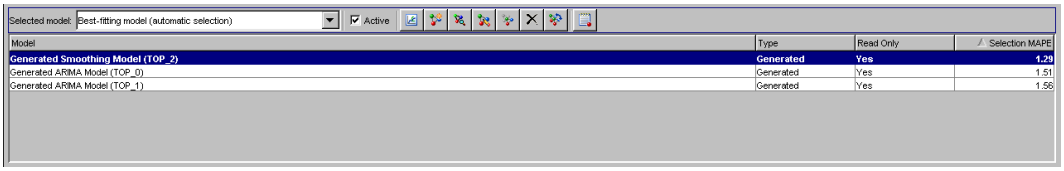

How SAS Forecast Studio Fits Models

When SAS Forecast Studio diagnoses a project, SAS Forecast Studio attempts to fit all the models in the model selection list. SAS Forecast Studio runs a series of diagnostics to determine the characteristics of the data (such as seasonality or intermittency), and it avoids models that are inappropriate for the data.

- If the diagnostics determine that a series is intermittent, then you cannot use continuous time series models, such as ARIMA, exponential smoothing, or unobserved components models.
- If the diagnostics determine that a series is continuous, then you cannot use an intermittent time series model, such as the intermittent demand model.

If SAS Forecast Studio fails to fit a model, then no statistics of fit are displayed for that model. Instead, the last column in the model selection list says "Failed." If none of the models in the model selection list fit, then SAS Forecast Studio attempts to use the model selection list that ships with SAS Forecast Studio. This default model selection list is made up of exponential smoothing models.

However, the status of a model can change depending on the other models in the model selection list. If you change a model selection list (for example, you add, edit, or delete a model from the list or you reforecast a series) and you have selected the **Perform seasonality test** option in the Forecasting Settings dialog box, then the seasonality test that is used to re-create the model selection list is more rigorous than the test that is used when SAS Forecast Studio diagnoses and forecasts the entire project. The more rigorous test could result in models that previously fit appearing as failed in the model selection list. For more information about the **Perform seasonality test** option, see "Diagnostics Options" on page 88.

Working with the Model Selection List

Using the model selection list in the Model View, you can complete the following tasks:

- select a model as the forecast model
- set a series as inactive
- reforecast the series
- add a model to a series
- edit a model
- copy a model
- delete a model
- compare models

Analyzing Models

View the Model Plots

You might want to view a plot of the model to gain a better understanding of the model.

To view a model plot, complete the following steps:

1. Select **View** → **Model View** to open the Model View.
2. From the model selection list, select a model.
3. Click **Plots**.
4. From the drop-down list, select from the following options:

   - **Forecast** - displays a plot of the model and the forecast that covers the time range of the actual data and the forecast horizon.
   - **Forecast Only** - displays the forecast in the forecast horizon only.
   - **Model** - displays a plot of the model with predictions and confidence limits that cover the time range of the actual data.
   - **Model and Forecast** - displays a model plot that includes the generated forecasts in the forecast horizon.
   - **Model Components** - displays plots of each component that is used by the model.
   - **Prediction Errors** - displays a needle plot and histogram of the prediction errors.
   - **Autocorrelation Function** - displays the autocorrelation function and standardized autocorrelation function of the prediction errors.
   - **Partial Autocorrelation Function** - displays the partial autocorrelation function and standardized partial autocorrelation function of the prediction errors.
   - **Inverse Autocorrelation Function** - displays the inverse autocorrelation function and standardized inverse autocorrelation function of the prediction errors.
   - **White Noise Probability** - displays the white noise probability plot and log white noise probability plot of the prediction errors.

   NOTE: Not all plots might be available for the current model. If a plot contains content that does not apply to the current model, then SAS Forecast Studio does not open this plot when you select the plot option from the drop-down list.

NOTE: In SAS Forecast Studio, you can copy a graph and save it to an external file. To export a graphic, right-click the graphic and select **Copy** from the pop-up menu. Then paste the graphic in the image editor of your choice.

View the Model Tables

You might want to view the statistics for the model to gain a better understanding of the model.

To view a model table, complete the following steps:

1. Select **View** → **Model View** to open the Model View.
2. From the model selection list, select a model.
3. Click **Tables**.
4. From the drop-down list, select from the following options:
   - **Model Parameter Estimates** - displays estimates, standard errors, and significance tests for each of the model parameters.
   - **Forecasts** - displays forecasted values, standard errors, and confidence limits for the time range of the forecast horizon.
   - **Statistics of Fit** - displays all available statistics of fit for the forecasting model.
   - **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.
   - **Test for Unbiasedness** - displays all the tests for unbiasedness for the forecasting model.

   NOTE: Not all tables might be available for the current model. If a table contains content that does not apply to the current model, then SAS Forecast Studio does not open this table when you select the table option from the drop-down list.

The model table opens in the Model 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).
- **Show Obs # Column** - includes a column with an observation number in the table.
- **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.

Specifying the Forecast Model

What Is the Forecast Model?

The forecast model is the model that SAS Forecast Studio uses to generate a numerical prediction of the future values for the time series. By default, SAS Forecast Studio chooses the best-fitting model in the list as the forecast model. This model appears as **Best-fitting model (automatic selection)** in the **Selected model** drop-down list. The best-fitting model can change over time as the performance changes for the other models in the list.

However, you can select the model to use as the forecast model by using the **Selected model** drop-down list in the Model View. When you reforecast the project, SAS Forecast Studio determines the forecast model based on the current forecast model and how you specified to reforecast the data.

- If you select an automatically generated model as the forecast model and then reforecast the project to re-create the automatically generated models, then the selection in the **Selected model** drop-down list reverts to the **Best-fitting model (automatic selection)** option.
- You change the forecast model for one or more series. You then reforecast the project using the **Refit all models, and select a forecast model** option. SAS Forecast Studio keeps the forecast model that you selected rather than selecting the best model from all the fitted models.
- You change the forecast model for an individual series. You then reforecast only that series using the **Refit all models, and select a forecast model** option. SAS Forecast Studio prompts you to choose between the following:
  - keep the forecast model that you selected
  - change the forecast model to the best model from all the fitted models

For more information about the reforecasting options, see "Reforecasting a Series or Project" on page 91.

Select the Forecast Model

To select a forecast model, complete the following steps:

1. Select **View** → **Model View** to open the Model View.
2. From the **Selected model** drop-down list, select the model that you want to be the forecast model. If you want SAS Forecast Studio to always use the best-fitting model as the forecast model, then select the **Best-fitting model (automatic selection)** option.

This drop-down list is not available if you have set the current series to inactive. For more information, see “Specifying an Inactive Series” on page 94.

NOTE: If you select a generated model, then the Selected Model dialog box opens. Because generated models are replaced when a series is diagnosed, the selected model does not persist when the series is diagnosed. You must specify whether to copy the generated model and use the copy or whether to always use the generated model.

When you change the forecast model for a series that is part of a hierarchy, a "Reconciliation is out of date" warning appears. Click **Update** to reconcile the hierarchy. For more information about how SAS Forecast Studio reconciles the hierarchy, see “Understanding Hierarchy Reconciliation” on page 106.

Add a Model to a Series

If you are working in the model selection list, then SAS Forecast Studio adds the model to the current series.

To add a model to the series that is currently being displayed in the Model View, complete the following steps:

1. Click  in the model selection list in the Model View. The Add Model dialog box opens.
2. Specify how to create the new model that you are adding to the series. You can choose from the following options:
   - **Fit an existing model without editing it** - enables you to select a default model or a custom model to add to the series. For more information, see “About the Model Repository” on page 188.
   - **Create and fit a new model by editing an existing model** - enables you to edit a default model or a custom model before adding it to the current series.
   - **Create and fit an entirely new model** - enables you to create an entirely new model. For more information, see “Create a User-Defined Model” on page 144.
3. Click **OK**.

View a Model

To view the information for a model in a report format, complete the following steps:

1. Select **View** → **Model View** to open the Model View.

2. Select the model from the model selection list.
3. Click  . A dialog box with a read-only view of the model specification opens.

You can also view the model specification by double-clicking the name of the model in the model selection list.

Import Models

You might want to use a model that is external to SAS Forecast Studio. You can do this by importing models from an external list.

To import models from an external list, complete the following steps:

1. Select **Project** → **Forecasting Settings**. The Forecasting Settings dialog box opens.
2. In the selection pane, select **Models**.
3. Select the **Try models from an existing model list** check box. By default SAS Forecast Studio uses the SASHELP.HPFDFLT.TSFSELECT list. For more information about this model selection list, see Appendix D, "Default Model Selection Lists."
4. To select a different model selection list, click **Browse**. The Select Model Selection List dialog box opens, and you can select the model selection list that you want to use.

   Click **OK**.
5. Click **OK** in the Forecasting Settings dialog box to save your changes.

SAS Forecast Studio reforecasts the data and includes the imported models.

Edit a Model

About Editing Models

In SAS Forecast Studio, you can edit a model in the following ways:

- from the model selection list in the Model View
- from the model repository

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 refit the model to all of the series in the project, complete the following steps:

1. Select **Project** → **Reforecast**. The Reforecast Project dialog box opens.
2. Select **Refresh the current forecast model, updating the parameter values** and click **OK**.

Edit a Model from the Model Selection List

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

1. Select **View** → **Model View** to open the Model View. The model selection list appears at the top of this view. For more information, see "Using the Model Selection List" on page 191.
2. Select the model that you want to edit and click  . The model specification dialog box for that model opens.
3. In the model specification dialog box, specify the options that you want to edit and click **OK**. For information about the options that are available for a model, see "Additional Information about the Model Specifications" on page 199.

NOTE: If you cannot edit some of the models in the model selection list, then a warning appears above the model list, and the edit button is disabled for those models. You cannot edit the following models:

- default models that shipped with SAS Forecast Studio. For more information, see "Working with Default Models" on page 189.
- models that SAS Forecast Studio created using the HPFDIAGNOSE procedure. You select the models that SAS Forecast Studio generates from the Models pane in the Forecasting Settings dialog box. For more information, see "Models Options" on page 90.
- imported models. For more information, see "Import Models" on page 197.

Edit a Model from the Model Repository

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

1. Select the series in the hierarchy or filtered table and select **Project → Model Repository**. The Model Repository dialog box opens. For more information about this dialog box, see "About the Model Repository" on page 188.
2. Click the **Custom Models** tab.
3. Select the model that you want to edit from the list and click **Edit**. The model specification dialog box for that model opens.
4. In the model specification dialog box, specify the options that you want to edit and click **OK**. For information about the options that are available for a model, see "Additional Information about the Model Specifications" on page 199.

Additional Information about the Model Specifications

For more information about each of the model specification dialog boxes that you use to edit the models, see the following topics:

- "Create an ARIMA Model" on page 146
- "Create a Curve Fitting Model" on page 175
- "Create an Exponential Smoothing Model" on page 164
- "Create an Intermittent Demand Model" on page 167
- "Create a Moving Average Model" on page 173
- "Create a Multiple Regression Model" on page 170
- "Create a Random Walk Model" on page 177
- "Create a Subset (Factored) ARIMA Model" on page 150
- "Create an Unobserved Components Model" on page 153

Copy Models

To copy a model, complete the following steps:

1. Select **View** → **Model View** to open the Model View.
2. In the model selection list, select the model that you want to copy and click  .
3. Edit the model specification. By default, the name of the copied model is *model-name*Copy*n* where *model-name* is the original model name and *n* is the lowest integer value available. This name appears in the model selection list unless you edit it. For more information, see "Edit a Model" on page 197.

Remove Models from the Model Selection List

When you delete a model from the model selection list, the model is not deleted from the model repository, but it no longer appears in the model selection list.

To delete a model from the model selection list, select the model name and click  . You are prompted to confirm this removal.

NOTE: If you delete all the models in the model selection list, then SAS Forecast Studio automatically fits an exponential smoothing model so that a forecast can be produced for the series. This exponential smoothing model is automatically removed when you add another model to the model selection list.

Compare Models

If you have created several models for a series, then you might want to compare the statistics of fit for the models that were fitted to that series. You can use the Compare Models dialog box to compare these statistics of fit. Through this comparison, you can determine which model that you want to use.

To compare models, complete the following steps:

1. Select **View** → **Model View** to open the Model View.
2. To compare all of the models for the series, select **Series** → **Compare Models** or click  . The Compare Models dialog box opens.
3. Compare the statistics of fit for each model in the series. For each statistic of fit, the cell for the best-fitting model is shaded in green. For more information about the available statistics, see Appendix C, "Statistics of Fit".
4. Click **Close** after you have finished this comparison.

Figure 11.3 Compare Models Dialog Box

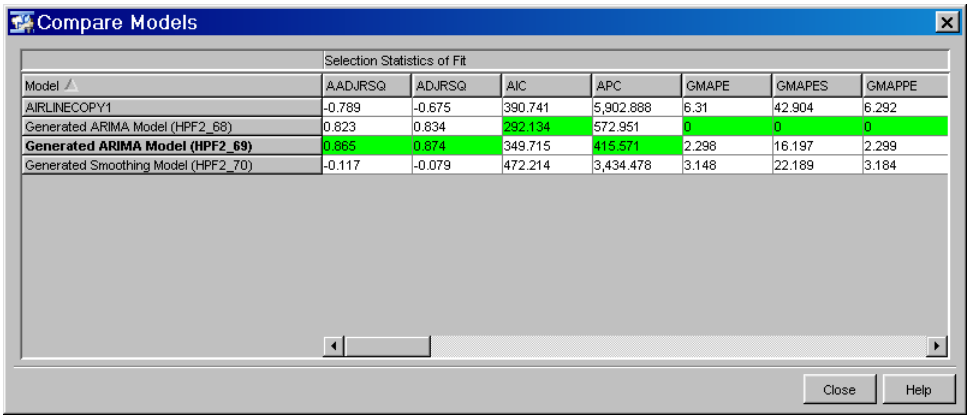

Chapter 12

Generating Reports

Contents

Understanding Reports

Overview of Reports

You can use reports to share 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.

One advantage to using stored processes to create reports is that you can include parameters in the stored process code. Including parameters in your code increases your ability to reuse code and also enables you to customize the report at run time. The author of the stored process can also specify whether parameters are required or whether you can use the default values. If a parameter is required, then you must specify a value before SAS Forecast Studio will run the report.

Sample Reports

SAS Forecast Studio ships with some sample reports. These reports are visible from the Reports and Stored Processes dialog box. For a description of each of the sample reports, see Appendix E, "Sample Reports in SAS Forecast Studio."

In addition to these sample reports, you might want to create your own reports. For more information, see "Create a Report" on page 204.

NOTE: Some sample reports, included for illustration purposes, use SAS procedures that are not directly licensed with SAS Forecast Server. If your site has licenses for these additional SAS procedures, then you can run these sample reports. If your site does not have licenses for these procedures, then these reports will not run. For more information about what is available at your site, contact your site administrator.

Run a Report

To run a report for the selected series, complete the following steps:

1. Select **Tools** → **Reports and Stored Processes**. The Reports and Stored Processes dialog box opens.
2. From the tree, select the report that you want to run. When you select a report from the list, the Information pane displays a brief description (if one has been provided) and the date that the report was created.
3. Click **Run**. If the stored process code that generates the report includes any parameters that need to be customized at run time, the Report Parameters dialog box opens. If no parameters are included, then this dialog box does not appear.

   After you have specified values for any required parameters, the report opens in a new Web browser window.

   NOTE: Some reports require a reporting variable or independent variable to run. If the report you select requires one of these variables and the current project does not have either of these variables defined, then an error message appears and the report does not run. For information about assigning an independent or reporting variable, see "Modify Variable Assignments" on page 49.

Create a Report

Although SAS Forecast Studio ships with sample reports, you can create your own reports. To create a report, you must write a stored process and register it to SAS Management Console by using the BI Manager.

For more information about how to write a stored process and make it available in SAS Forecast Studio, see "Additional Administration Tasks" in the *SAS Forecast Server Administrator's Guide*.

Part III

Advanced Topics

Chapter 13

Working in Batch Mode

Contents

Running Projects in Batch Mode

Saving the Project Code to a Batch File

Using SAS Forecast Studio, you can save the project code as a series of SAS files. You can use these files to create batch jobs for forecasting.

NOTE: In order for this code to execute properly in batch mode, you must declare the library references that the code uses to run. You can declare these library references in your `autoexec.sas` file. For help declaring library references, contact the SAS administrator at your site.

There are two versions of each SAS file. The files differ in whether or not they include the code for importing data.

- If the filename contains DO_NOT_IMPORT_DATA, then the code for importing new data is not included in this batch file. When this code is run, only the data that was available during the project creation is used in the calculations.
- If the filename contains IMPORT_DATA, then the code for importing new data is included in this batch file.

  This data specification code does the following:

  1. creates a copy of the original data that is sorted by the BY variables
  2. creates XML files that describe the hierarchy

3. accumulates the data for each level

When this code is run, any new data that is available is used in the calculation.

Descriptions of the Project Code in Each Batch File

The following list describes the project code that is contained in each batch file:

CREATE_PROJECT code files

The CREATE_PROJECT_DO_NOT_IMPORT_DATA and CREATE_PROJECT_IMPORT_DATA files contain the SAS code to re-create the project. This code includes all of the data preparation and forecasting steps.

These files include the following project code:

- Library declarations that create librefs, which reference existing directories inside a project. Library declarations include definitions of macro variables that describe the project location, the data specification location, and the event repository location.
- Catalog declarations for model repositories that are used by the project.
- Macro variable declarations that describe procedure settings.
- PROC HPFDIAGNOSE statements for each level in the project. These are "destructive," which means that SAS Forecast Studio rediagnoses all series in the project and re-creates the model repositories.

  PROC HPFENGINE statements for each level in the project are set to TASK=SELECT. The HPFENGINE procedure selects and fits the models for each series and creates the forecasts.
- Calls to the HPFRECON macro that performs reconciliation for a hierarchical project.
- Calls to macros that reconcile overrides if any overrides have been specified.
- A statement that clears catalog and filename declarations.

DIAGNOSE_DESTRUCTIVE code files

The DIAGNOSE_DESTRUCTIVE_DO_NOT_IMPORT_DATA and DIAGNOSE_DESTRUCTIVE_IMPORT_DATA files contain the SAS code to rediagnose the project. This code does not include the data preparation steps.

These files contain the following project code:

- Library declarations that create librefs, which reference existing directories inside a project. Library declarations include definitions of macro variables that describe the project location, the data specification location, and the event repository location.
- Catalog declarations for model repositories that are used by the project.
- Macro variable declarations that describe procedure settings.
- PROC HPFDIAGNOSE statements for each level in the project. These are "destructive," which means that SAS Forecast Studio rediagnoses all series in the project and re-creates the model repositories.

PROC HPFENGINE statements for each level in the project are set to TASK=SELECT. The HPFENGINE procedure selects and fits the models for each series and creates the forecasts.

- Calls to the HPFRECON macro that performs reconciliation for a hierarchical project.
- Calls to macros that reconcile overrides if any overrides have been specified.
- A statement that clears catalog and filename declarations.

DIAGNOSE_NON_DESTRUCTIVE_PROJECT code files

The DIAGNOSE_NON_DESTRUCTIVE_PROJECT_DO_NOT_IMPORT_DATA and DIAGNOSE_NON_DESTRUCTIVE_PROJECT_IMPORT_DATA files contain the SAS code to rediagnose the project. This code does not include the data preparation steps.

These files include the following project code:

- Library declarations that create librefs, which reference existing directories inside a project. Library declarations include definitions of macro variables that describe the project location, the data specification location, and the event repository location.
- Catalog declarations for model repositories that are used by the project.
- Macro variable declarations that describe procedure settings.
- PROC HPFDIAGNOSE statements for each level in the project are not destructive. These statements use the INEST option of the HPFDIAGNOSE procedure to preserve all customizations that are created by user-defined models.

  PROC HPFENGINE statements for each level in the project are set to TASK=SELECT. The HPFENGINE procedure selects and fits the models for each series and creates the forecasts.
- Calls to the HPFRECON macro that performs reconciliation for a hierarchical project.
- Calls to macros that reconcile overrides if any overrides have been specified.
- A statement that clears catalog and filename declarations.

SELECT_MODELS code files

The SELECT_MODELS_DO_NOT_IMPORT_DATA and SELECT_MODELS_IMPORT_DATA code files contain the SAS code to perform model selection, to estimate the parameters of the selected model, and to produce forecasts. This code does not include the data preparation, diagnose, fit, and forecasting steps.

These files include the following project code:

- Library declarations that create librefs, which reference existing directories inside a project. Library declarations include definitions of macro variables that describe the project location, the data specification location, and the event repository location.
- Catalog declarations for model repositories that are used by the project.
- Macro variable declarations that describe procedure settings.
- PROC HPFENGINE statements (TASK=SELECT) for each level in the project. The HPFENGINE procedure selects and fits the models for each series and creates the forecasts.
- A statement that replaces missing values for stochastic independent variables.

- Calls to the HPFRECON macro that performs reconciliation for a hierarchical project.
- Calls to macros that reconcile overrides if any overrides have been specified.
- A statement that clears catalog and filename declarations.

FIT_MODELS code files

The FIT_MODELS_DO_NOT_IMPORT_DATA and FIT_MODELS_IMPORT_DATA code files contain the SAS code to estimate parameters using the model that you specified and then to create a forecast. No model selection is performed. This code does not include the data preparation and diagnose steps.

These files include the following project code:

- Library declarations that create librefs, which reference existing directories inside a project. Library declarations include definitions of macro variables that describe the project location, the data specification location, and the event repository location.
- Catalog declarations for model repositories that are used by the project.
- Macro variable declarations that describe procedure settings.
- PROC HPFENGINE statements (TASK= FIT) for each level in the project. The HPFENGINE procedure refits selected models for each series and creates forecasts.
- A statement that replaces missing values for stochastic independent variables.
- Calls to the HPFRECON macro that performs reconciliation for a hierarchical project.
- Calls to macros that reconcile overrides if any overrides have been specified.
- A statement that clears catalog and filename declarations.

FORECAST_MODELS code files

The FORECAST_MODELS_DO_NOT_IMPORT_DATA and FORECAST_MODELS_IMPORT_DATA contains the SAS code to reforecast the model and parameters. When refitting the model parameters, SAS Forecast Studio uses the estimate of the previous parameter as a starting point for reestimation. This code does not include the data preparation, diagnose, and fit model steps.

These files include the following project code:

- Library declarations that create librefs, which reference existing directories inside a project. Library declarations include definitions of macro variables that describe the project location, the data specification location, and the event repository location.
- Catalog declarations for model repositories that are used by the project.
- Macro variable declarations that describe procedure settings.
- PROC HPFENGINE statements (TASK= FORECAST) for each level in the project. The HPFENGINE procedure re-creates forecasts for each series by using selected and fitted models.
- A statement that replaces missing values for stochastic independent variables.
- Calls to the HPFRECON macro that performs reconciliation for a hierarchical project.
- Calls to macros that reconcile overrides if any overrides have been specified.
- A statement that clears catalog and filename declarations.

About the SAS Forecast Server Macros

Summary of SAS Forecast Server Macros

The following SAS macros are provided with SAS Forecast Server software. A SAS macro is a program that generates SAS statements. Macros make it easy to produce and execute complex SAS programs that would be time-consuming to write yourself.

Table 13.1 shows the SAS macros that are included in the SAS Forecast Server software.

Table 13.1 SAS Forecast Server Macros

Macro Name	Description
%FSCOPY	Copies a SAS Forecast Studio project to a new destination. **NOTE:** Both servers must run the same version or higher of SAS Forecast Studio.
%FSCREATE	Creates a new SAS Forecast Studio project in batch mode.
%FSDELARC	Deletes an archived SAS Forecast Studio project.
%FSDELPRJ	Deletes an existing SAS Forecast Studio project.
%FSEXP14	Exports all SAS Forecast Studio 1.4 projects to archived files. This macro must be run on the SAS Forecast Server 1.4 middle tier machine.
%FSEXPALL	Exports all SAS Forecast Studio projects to archive files.
%FSEXPORT	Exports a single SAS Forecast Studio project to an archive file.
%FSGETPRJ	Retrieves the metadata about the SAS Forecast Studio projects.
%FSIMPALL	Imports all SAS Forecast Studio projects from archived files.
%FSIMPORT	Imports a SAS Forecast Studio project from an archived file.
%FSLOAD	Opens an existing SAS Forecast Studio project, and loads global macro variables that describe the project.
%FSMIGPRJ	Migrates an existing SAS Forecast Studio project to the current version of SAS Forecast Server.
%FSMIGALL	Migrates all existing SAS Forecast Studio projects to the current version of SAS Forecast Server.
%FSMIGPRJ	Migrates an existing SAS Forecast Studio project to the current version of SAS Forecast Server.
%FSMOVE	Moves a SAS Forecast Studio project to a new destination. **NOTE:** Both servers must run the same version or higher of SAS Forecast Server.
%FSREN	Renames a single SAS Forecast Studio project.
%FSRUNPRJ	Opens an existing SAS Forecast Studio project, and runs the project at a given stage.
%FSSETCRB	Assigns the creator of a SAS Forecast Studio project.
%FSSETLOC	Changes the default location for SAS Forecast Studio projects.
%FSSETPUB	Determines whether public access to a SAS Forecast Studio project should be enabled.

Table 13.1 SAS Forecast Server Macros (continued)

Macro Name	Description
%FSUNREG	Unregisters an existing SAS Forecast Studio project from the metadata server. **NOTE:** Only an adminstrative user can deploy this macro.

Additional Information

These macros are available automatically for you to use in your SAS programs.

For more information, see the following resources:

- *SAS Forecast Server Administrator's Guide* describes the syntax for each of the macros and provides examples of how to use the macros.
- *SAS Macro Language: Reference* describes the SAS macro facility.

Chapter 14

Working with SAS Forecast Studio Tasks in the SAS Add-In for Microsoft Office

Contents

Overview of the SAS Forecast Studio Tasks

What Are the SAS Forecast Studio Tasks?

If the SAS Enterprise Business Intelligence Server is also installed at your site, you can use the SAS Add-In for Microsoft Office to run the following tasks:

- the Forecast Studio Create Project task to create a SAS Forecast Studio project from an Excel or SAS data source
- the Forecast Studio Open Project task to open the results from a selected series in an existing SAS Forecast Studio project
- the Forecast Studio Submit Overrides tasks to submit overrides for the forecast data in an existing SAS Forecast Studio project

What Is the SAS Add-In for Microsoft Office?

The SAS Add-In for Microsoft Office extends the functionality of Microsoft Excel, Microsoft Word, and Microsoft PowerPoint by enabling you to access SAS analytics and SAS reporting functionality without any SAS programming experience. The SAS add-in is designed for users who are familiar with these Microsoft Office programs but who might be new to SAS.

When the SAS add-in is installed on your computer, a SAS menu and the SAS Analysis Tools toolbar are automatically integrated into Excel, Word, and PowerPoint.

Figure 14.1 Location of SAS Menu in Microsoft Excel

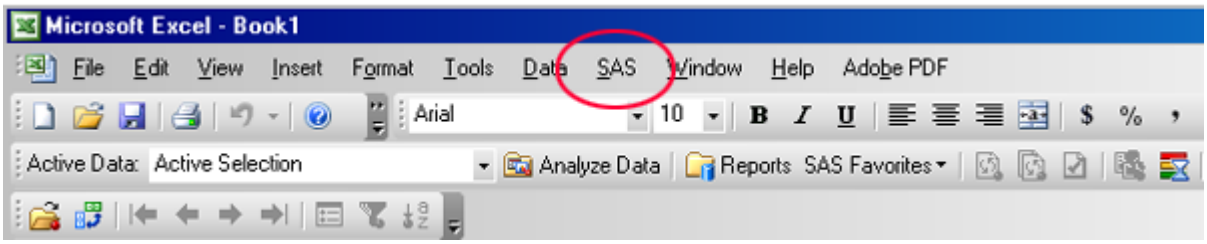

The SAS add-in includes approximately 80 SAS tasks that enable you to perform a variety of analyses. The Forecast Studio Create Project task, the Forecast Studio Open Project task, and the Forecast Studio Submit Overrides task are available in the Analyze Data window. In the SAS Tasks folder, click **Time Series** to view these tasks.

Figure 14.2 Forecasting Tasks in Analyze Data Window

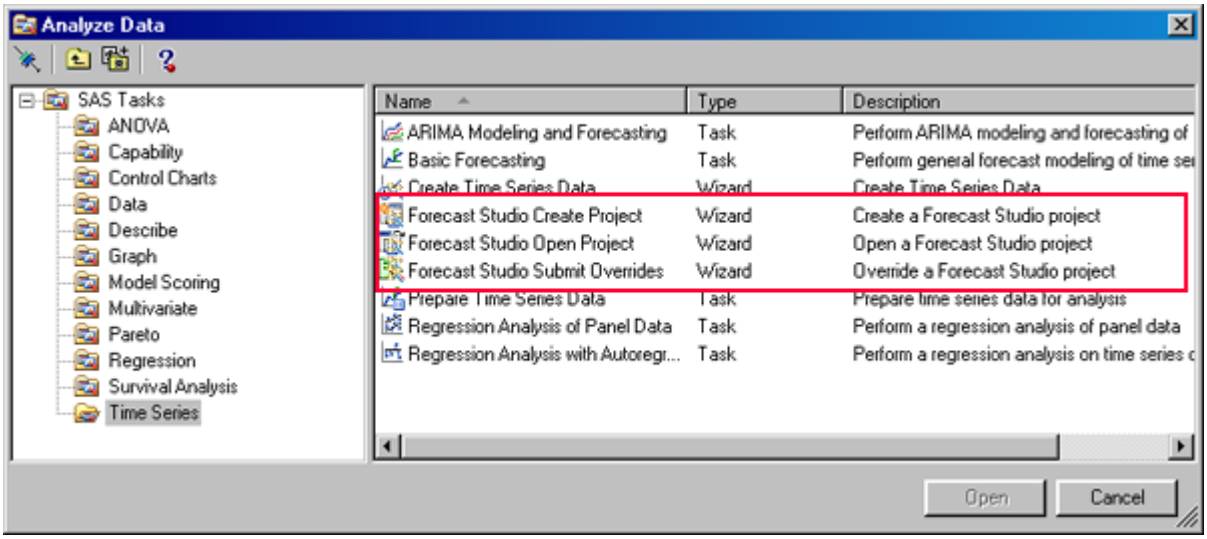

Prerequisites for Using the SAS Forecast Studio Tasks

In order to use the SAS Forecast Studio tasks in the SAS Add-In for Microsoft Office, your site administrator must have completed the following steps:

- installed SAS Forecast Server 2.1.
- installed the SAS Add-In 2.1 for Microsoft Office, which is a product in the SAS Enterprise Business Intelligence Server bundle.

  NOTE: The Forecast Studio Submit Overrides task is available only if you download it from the Software Downloads page at `http://www.sas.com/download`. Click the **SAS Add-In for Microsoft Office** link to download the task.
- configured SAS Forecast Studio Server to use the SAS Add-In for Microsoft Office. For more information, see the *SAS Forecast Server Administrator's Guide*.

The Forecast Studio Create Project Task

The Forecast Studio Create Project task enables a SAS add-in user to create a SAS Forecast Studio project by using an Excel or SAS data source. Using this task, a SAS add-in user can specify the forecasting variables, choose whether to forecast the data hierarchically, and specify the forecast horizon for the new SAS Forecast Studio project.

The Forecast Studio Open Project Task

The Forecast Studio Open Project task enables a SAS add-in user to select a series from an existing SAS Forecast Studio project and specify the result types to open in Microsoft Excel, Word, or PowerPoint.

Here are some examples of the results that a SAS add-in user can open in Microsoft Excel, Word, or PowerPoint:

- the forecasting plot
- the data table
- all of the models that have been defined for the current project
- the estimates, standard errors, and significance tests for each of the model parameters
- tables for each of the components that are used by the model

- all the available statistics of fit for the forecasting model
- the forecast values of the dependent variable for the time range of the forecast horizon

After opening these results, a SAS add-in user can perform additional analyses using the other SAS tasks that ship with the SAS Add-In for Microsoft Office.

The Forecast Studio Overrides Task

The Forecast Studio Submit Overrides task enables a SAS add-in user to submit overrides for the forecast data in an existing SAS Forecast Studio project. After the SAS add-in user submits an override, SAS Forecast Server reconciles the project. How the override is applied depends on whether the data in the project is forecasted hierarchically.

- If the data in the project is forecasted hierarchically, then the override is an adjustment that is done with respect to the reconciled statistical forecast.
- If the data in the project is not forecasted hierarchically, then the override is an adjustment that is done with respect to the statistical forecast.

If SAS Forecast Server detects an override conflict when reconciling the project, then the SAS add-in user sees an error message stating that the project cannot be reconciled and that the project must be opened in SAS Forecast Studio to reconcile the conflicts.

When the project is opened in SAS Forecast Studio, the "Reconciliation is out of date" warning appears at the top of the workspace. Click **Update** to open the Override Conflicts dialog box to see all of the overrides that were submitted. To resolve the override conflicts, the SAS Forecast Studio user can modify the overrides by using the Override Conflicts dialog box, or the SAS Forecast Studio user can close this dialog box and update overrides by using the data table.

Additional Information

For more information about how to use each of these tasks, see the Help for the SAS Add-In for Microsoft Office.

Part IV

Appendixes

Appendix A

Reserved Variable Names in SAS Forecast Studio

Contents

Overview of the Reserved Variables Names

For each project, SAS Forecast Studio creates several output data sets. The variable names in your input data set cannot match any of the variable names in these output data sets. The variable names in your input data set also cannot start with an underscore. If you try to assign a variable to a role and the variable name matches either of these conditions, then an error message appears.

For more information about the output data sets that SAS Forecast Studio creates, see Appendix B, "Output Data Sets in the Project Directory."

Reserved Variable Names

The following table lists alphabetically the variables that are used by SAS Forecast Studio.

Variable Name	Description
_VariableName	Any variable name that begins with an underscore
ACTUAL	Dependent series value
ACTUAL	Dependent series value
COMP	Name of the component
COMPONENT	Model component
CROSS	Cross-variable name
DEPTRANS	Dependent transformation indicator
DSVAR	Data set variable mapping
EST	Parameter estimate
EVENTS	Variable that indicates the presence of events

Variable Name	Description
FACTOR	Model factor
INPUTS	variable that indicates the presence of inputs
LABEL	Parameter or statistic label
LAG	Lag of input
LEADt	*t*-step ahead forecasts
LOWER	Lower confidence limit
MODE	Mode of decomposition
MODEL	Name of model
MODELTYPE	Model type
MODELVAR	Model variable mapping
NAME	Variable name
OUTLIERS	Variable that indicates the presence of outliers
PARM	Parameter name
PREDICT	Component forecast
PVALUE	Parameter estimate *p*-value
REGION	Forecast region
SEASON	Seasonal index
SEASONAL	Variable that indicates the presence of a seasonal model
SELECT	Name of selection list
SHIFT	Shift
STAT	Statistic name
STATUS	Success or 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
TREND	Variable that indicates the presence of a trend model
UPPER	Upper confidence limit
AADJRSQ	Amemiya's adjusted R-square
ACF	Autocorrelations
ACF2STD	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
AICC	Finite sample corrected Akaike Information Criterion
APC	Amemiya's Prediction Criterion
AVG	Average value
CC	Cycle component
CCF	Cross-correlations
CCF2STD	Cross-correlations beyond two standard errors
CCFNORM	Normalized cross-correlations

Variable Name	Description
CCFLPROB	Cross-correlation log probabilities
CCFPROB	Cross-correlation probabilities
CCFSTD	Cross-correlation standard errors
CCOV	Cross-covariances
CSS	Corrected sum of squares
DFE	Degrees of freedom error
ERROR	Prediction errors
GMAPE	Geometric mean percent error
GMAPPE	Geometric mean predictive percent error
GMAPES	Geometric mean absolute error percent of standard deviation
GMASPE	Geometric mean symmetric percent error
GMRAE	Geometric mean relative absolute error
IACF	Inverse autocorrelations
IACF2STD	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
LAG*h*	Correlation or cross-correlation statistics for lag *h*
LOWER	Lower confidence limits
MAE	Mean absolute error
MAPE	Mean absolute percent error
MAPPE	Symmetric mean absolute predictive percent error
MASE	Mean absolute scaled error
MAXAPES	Maximum absolute error percent of standard deviation
MAXERR	Maximum error
MAXIMUM	Maximum value
MAXPE	Maximum percent error
MAXPPE	Maximum predictive percent error
MAXRE	Maximum relative error
MAXSPE	Maximum symmetric percent error
MDAPE	Median percent error
MDAPES	Median absolute error percent of standard deviation
MDAPPE	Median predictive percent error
MDASPE	Median symmetric percent error
MDRAE	Median relative absolute error
ME	Mean error
MEAN	Mean value
MEDIAN	Median value
MINAPES	Minimum absolute error percent of standard deviation
MINERR	Minimum error
MINIMUM	Minimum value
MINPE	Minimum percent error
MINPPE	Minimum predictive percent error
MINRE	Minimum relative error

Variable Name	Description
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	Number of nonmissing observations or number of variance products
NAME	Variable name
NMISS	Number of missing observations
NMISSA	Number of missing actuals
NMISSP	Number of missing predicted
NPARMS	Number of parameters
NOBS	Number of observations
ORIGINAL	Original series index
PACF	Partial autocorrelations
PACF2STD	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 interval *t*
PREDICT	Predicted values
RANGE	Difference between the maximum and minimum values
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*
SMAPE	Symmetric mean absolute percent error
SSE	Sum of squares error
SST	Corrected total sum of squares
STD	Prediction standard errors
STDDEV	Standard deviation
SUM	Summation value
TC	Trend component
TCC	Trend-cycle component
TCS	Trend-cycle-seasonal component
TSS	Total sum of squares
UMSE	Unbiased mean square error
URMSE	Unbiased root mean square error
UPPER	Upper confidence limits

Variable Name	Description
USS	Uncorrected sum of squares
WN	White noise test statistics
WNLPROB	White noise test log probabilities
WNPROB	White noise test probabilities

Appendix B

Output Data Sets in the Project Directory

Contents

Overview of Output Data Sets

SAS Forecast Studio creates several output data sets. In general, these data sets contain the variables that you assign as BY variables. These output data sets are created for each level of the hierarchy and are saved in the following locations:

Windows operating environments
!ROOT\SAS\ForecastStudio\Projects*project-name***hierarchy**

UNIX operating environments
!ROOT/SAS/ForecastStudio/Projects/*project-name***/hierarchy**

Descriptions of Output Data Sets

From the project directory, you can view the following output data sets. If the forecast fails for one series, then the forecasts in the horizon appear as .F in the data. If the series is at or above the reconciliation level, then all the parent nodes have missing reconciled forecasts, and these missing values appear as .F in the data. If the series is at or below the reconciliation level, then all the children nodes have missing reconciled forecasts, and these missing values appear as .F in the data.

FINALFOR= data set
The FINALFOR= data set contains the observations that have a time ID value in the forecast horizon. If the RECFOR= data set is created for a level in the hierarchy, then the FINALFOR= data set is a subset of the RECFOR= data set; otherwise, the FINALFOR= data set is a subset of the OUTFOR= data set.

OUT= data set
The OUT= data set contains the variables that you assign as the time ID, BY variables, and

dependent variables. The values for the time ID are based on the interval of the time series. The values of the dependent variables are accumulated, and missing values are interpreted using the project options that you set. For more information about these options, see "Working with the Forecasting Settings for the Project" on page 84.

If any of the forecasting steps fail for a particular variable, the variable values are extended by missing values.

OUTCOMPONENT= data set

The OUTCOMPONENT= data set contains the variables that you assign as BY variables. This data set can also contain the variables listed below:

- _NAME_ - variable name or the name of the dependent variable
- _COMP_ - the name of the component
- _TIMEID_ - time ID values or the time ID variables
- _ACTUAL_ - the value of the dependent series
- _PREDICT_ - the forecast value of the component
- _LOWER_ - lower confidence limit
- _UPPER_ - upper confidence limit
- _STD_ - prediction standard error

OUTEST= data set

The OUTEST= data set contains the variables that you assign as BY variables. This data set can also contain the variables listed below if a forecasting model has been selected to forecast the actual time series. The following variables contain data that is related to the parameter estimation step:

- _NAME_ - variable name or the name of the dependent variable
- _SELECT_ - the name of the model selection list
- _MODEL_ - the name of the forecasting model
- _MODELVAR_ - the model variable that is used for mapping
- _DSVAR_ - the data set variable that is used for mapping
- _VARTYPE_ - the type of the variable to be forecast. In this case, the value is DEPENDENT.
- _TRANSFORM_ - the transformation for the model that was applied
- _COMPONENT_ - the model component. Examples of model components include autoregressive (AR), moving average (MA), and trend.
- _COMPMODEL_ - the model portion of an intermittent demand component
- _FACTOR_ - the model factor
- _LAG_ - the lag for the input
- _SHIFT_ - the shift
- _PARM_ - the parameter name
- _LABEL_ - the parameter label
- _EST_ - the parameter estimate

- _STDERR_ - standard error of the parameter estimate
- _TVALUE_ - the *t*-values of the parameter estimates
- _PVALUE_ - the *p*-values of the parameter estimates
- _STATUS_ - indicates the success or failure of estimating the parameter

OUTFOR= data set

The OUTFOR= data set contains the variables that you assign as BY variables. This data set can also contain the variables listed below:

- _NAME_ - variable name or the name of the dependent variable
- _TIMEID_ - time ID values or the time ID variables
- ACTUAL - values of the dependent series
- PREDICT - predicted values
- STD - prediction standard errors
- LOWER - lower confidence limits
- UPPER - upper confidence limits
- ERROR - prediction errors

If the forecasting step fails for a particular variable, no observations are recorded. If you specify to transform the time series, then the values in the variables listed above are the inverse transform forecasts. When you transform the time series, you can specify whether the median or mean forecasts are recorded.

OUTOVRD= data set

The OUTOVRD= data set contains any overrides that you specify in SAS Forecast Studio. This data set contains the variables that you assign as BY variables and the following variables:

- _NAME_ - variable name or the name of the dependent variable
- _TIMEID_ - time ID values or the time ID variables
- OVERRIDE - the values of the overrides that you specify
- OLOCK - an indicator variable that specifies whether the override is locked. A value of 0 means the override is locked, and a value of 1 means an unlocked override.

This data set is created only if you have specified overrides. If you did not specify any overrides, then SAS Forecast Studio does not create this data set.

OUTSTAT= data set

The OUTSTAT= data set contains the variables that you assign as BY variables. This data set can also contain the variables listed below. The following variables contain observations related to the statistics-of-fit step:

- _NAME_ - variable name or the name of the dependent variable
- _REGION_ - the region in which the statistics are calculated. Statistics that are calculated in the fit region are indicated by FIT. Statistics that are calculated in the forecast region are indicated by FORECAST.

- DFE - degrees of freedom error
- N - number of observations
- NOBS - number of observations that were used in calculating the statistics
- NMISSA - number of missing actuals
- NMISSP - number of missing predicted values
- NPARMS - number of parameters
- TSS - total sum of squares
- SST - corrected total sum of squares
- SSE - sum of squares error
- MSE - mean square error
- UMSE - unbiased mean square error
- RMSE - root mean square error
- URMSE - unbiased root mean square error
- MAPE - mean absolute percent error
- MAE - mean absolute error
- MASE - mean absolute scaled error
- RSQUARE - R-square
- ADJRSQ - adjusted R-square
- AADJRSQ - Amemiya's adjusted R-square
- RWRSQ - random walk R-square
- AIC - Akaike information criterion
- AICC - finite sample corrected AIC
- SBC - Schwarz Bayesian information criterion
- APC - Amemiya's prediction criterion
- MAXERR - maximum error
- MINERR - minimum error
- MINPE - minimum percent error
- MAXPE - maximum percent error
- ME - mean error
- MPE - mean percent error
- MDAPE - median absolute percent error
- GMAPE - geometric mean absolute percent error
- MINPPE - minimum predictive percent error
- MAXPPE - maximum predictive percent error
- MSPPE - mean predictive percent error
- MAPPE - symmetric mean absolute predictive percent error
- MDAPPE - median absolute predictive percent error
- GMAPPE - geometric mean absolute predictive percent error

- MINSPE - minimum symmetric percent error
- MAXSPE - maximum symmetric percent error
- MSPE - mean symmetric percent error
- SMAPE - symmetric mean absolute percent error
- MDASPE - mean absolute symmetric percent error
- GMASPE - geometric mean absolute symmetric percent error
- MINRE - minimum relative error
- MAXRE - maximum relative error
- MRE - mean relative error
- MRAE - mean relative absolute error
- MDRAE - median relative absolute error
- GMRAE - geometric mean relative absolute error
- MAPES - mean absolute error percent of standard deviation
- MDAPES - median absolute error percent of standard deviation
- GMAPES - geometric mean absolute error percent of standard deviation

If the statistics-of-fit step fails for a particular variable, no observations are recorded.

OUTSTATSELECT= data set

The OUTSTATSELECT= data set contains the same variables as the OUTSTAT= data set with the addition of the following:

- _MODEL_ - the name of the forecasting model
- _SELECT_ - the name of the model selection list
- _SELECTED_ - whether or not the model was chosen to forecast the dependent series

OUTSUM= data set

The OUTSUM= data set contains the variables that you assign as BY variables. The OUTSUM= data set records the summary statistics for each dependent variable. If you select a model for the series, then the forecast values can include the prediction values, the upper confidence limits, or the lower confidence limits, depending on the options that you specify. If you do not select a model, then the forecast values are set to missing.

The following variables that are related to summary statistics can also appear in this data set. These variables depend on the options that you set for accumulation and for the interpretation of missing values.

- _NAME_ - variable name or the name of the dependent variable
- _STATUS_ - forecasting status. Nonzero values imply that no forecast was generated for the series.
- NOBS - number of observations
- N - number of non-missing observations or number of observations
- NMISS - number of missing observations
- MIN - minimum value

- MAX - maximum value
- MEAN - mean value
- STDDEV - standard deviation

The following variables that are related to forecast summation are based on the values that you specify for the lead and for starting the summation:

- PREDICT - forecast summation of the predicted values
- STD - forecast summation of the prediction standard errors
- LOWER - forecast summation of the lower confidence limits
- UPPER - forecast summation of the upper confidence limits

Variance-related computations are computed only when no transformation is specified for the time series.

Your data set can also contain the _LEAD_ variable if you have a multistep forecast. A multistep forecast ranges from 1 to the lead number that you specify. This forecast can also contain the lower confidence limits, the upper confidence limits, or the predicted values.

If the forecast step fails for a particular variable, the variables related to forecasting are set to missing. The OUTSUM= data set contains both a summary of the (accumulated) time series and optionally its forecasts for all series.

RECFOR= data set

The RECFOR= data set contains the values of the reconciled forecasts. This data set can also contain the variables listed below. The following variables contain data that is related to the forecasting step:

- _NAME_ - variable name or the name of the dependent variable
- _TIMEID_ - time ID values or the time ID variables
- PREDICT - predicted values
- STD - prediction standard errors
- LOWER - lower confidence limits
- UPPER - upper confidence limits
- ERROR - prediction errors
- _RECONSTATUS_ - the reconciliation status. This variable contains a code that specifies whether the reconciliation is successful or not. A corresponding message is also printed to the log.

  The _RECONSTATUS_ variable can have the following values:

0	Success.
500	A locked equality constraint has been imposed.
1000	ID value out of the range with respect to the START= and END= interval.
2000	Insufficient data to reconcile.
3000	Reconciliation failed for the predicted value. This implies that it also failed for the confidence limits and standard error.

4000 Reconciliation failed for the standard error.

5000 Reconciliation failed for the confidence limits.

6000 The constrained optimization problem is infeasible.

7000 The option DISAGGREGATION=PROPORTION has been changed to DISAGGREGATION=DIFFERENCE for this observation due to discordant signs in the input.

8000 The option STDMETHOD= provided by the user has been changed for this observation.

9000 The option CLMETHOD= provided by the user has been changed for this observation.

10000 The standard error hit the limits imposed by the STDDIFBD= option.

11000 Multiple warnings have been printed to the log for this observation.

NOTE: The RECFOR= data set is not created if you chose not to forecast your data hierarchically or when the reconciliation functionality is disabled. If you chose to forecast the data hierarchically and reconciliation functionality is enabled, then SAS Forecast Studio does not create the RECFOR= data set for the reconciliation level in the hierarchy. In this case, the reconciled forecasts are identical to the forecasts generated by the statistical model, so the content of the RECFOR= data set will be the same as the OUTFOR= data set.

RECSTAT= data set

The RECSTAT= data set contains the values of the reconciled forecasts. This data set also contains the same variables as the OUTSTAT= data set.

If the statistics-of-fit step fails for a particular variable, no observations are recorded. If you specified to transform the time series, the values in the variables listed above are computed based on the inverse transform forecasts. When you transform the time series, you can specify whether the median or mean forecasts are recorded.

NOTE: The RECSTAT= data set is not created if you chose not to forecast your data hierarchically or when the reconciliation functionality is disabled. If you chose to forecast the data hierarchically and reconciliation functionality is enabled, then SAS Forecast Studio does not create the RECSTAT= data set for the reconciliation level in the hierarchy. In this case, the reconciled forecasts are identical to the forecasts generated by the statistical model, so the content of the RECSTAT= data set will be the same as the OUTSTAT= data set.

RECSUM= data set

The RECSUM= data set contains the variables that you assigned to the BY variables role. The RECSUM= data set records the summary statistics for each dependent variable. If you selected a model for the series, then the reconciled forecasts can include the prediction values, the upper confidence limits, or the lower confidence limits, depending on the options that you specify. If you did not select a model, then the reconciled forecasts are set to missing.

The following variables that are related to summary statistics can also appear in this data set. These variables depend on the options that you set for accumulation and for the interpretation of missing values.

- _NAME_ - variable name or the name of the dependent variable
- _STATUS_ - forecasting status. Nonzero values imply that no forecast was generated for the series.

- NOBS - number of observations
- N - number of non-missing observations or number of observations
- NMISS - number of missing observations
- MIN - minimum value
- MAX - maximum value
- MEAN - mean value
- STDDEV - standard deviation

The following variables that are related to forecast summation are based on the values that you specify for the lead and for starting the summation:

- PREDICT - forecast summation of the predicted values
- STD - forecast summation of the prediction standard errors
- LOWER - forecast summation of the lower confidence limits
- UPPER - forecast summation of the upper confidence limits

Variance-related computations are computed only when no transformation is specified for the time series.

Your data set can also contain the _LEAD_ variable if you have a multistep forecast. A multistep forecast ranges from one to the lead number that you specified. This reconciled forecast can also contain the lower confidence limits, the upper confidence limits, or the predicted values.

If the forecast step fails for a particular variable, the variables related to forecasting are set to missing. The RECSUM= data set contains both a summary of the (accumulated) time series and optionally its forecasts for all series.

NOTE: The RECSUM= data set is not created if you chose not to forecast your data hierarchically or when the reconciliation functionality is disabled. If you chose to forecast the data hierarchically and reconciliation functionality is enabled, then SAS Forecast Studio does not create the RECSUM= data set for the reconciliation level in the hierarchy. In this case, the reconciled forecasts are identical to the forecasts generated by the statistical model, so the content of the RECSUM= data set will be the same as the OUTSUM= data set.

Appendix C

Statistics of Fit

Contents

What Are Statistics of Fit?

The statistics of fit are statistical values that are used to evaluate how well a forecasting model performs by comparing the actual data to the predictions. For a given forecast model that has been fitted to the time series data, the model should be checked or evaluated to see how well it fits or forecasts the data. Commonly used statistics of fit are Root Mean Square Error (RMSE), Mean Absolute Percent Error (MAPE), Akaike Information Criteria (AIC), and many others. The statistics of fit can be computed from the model residuals or the prediction errors.

When the full range of data is used to both fit and evaluate the model, this is referred to as in-sample evaluation. When the most recent data is excluded for parameter estimation (holdout) and this holdout sample is used for evaluation, this is referred to as holdout sample evaluation. Holdout sample analysis is similar to the training and testing of neural networks. A portion of the data is withheld from training (fit) and the withheld data (holdout) is used to test performance.

When a particular statistic of fit is used for forecast model selection, it is referred to as the model selection criterion. For example, if the MAPE (an often recommended choice) is used as a model selection criterion, the forecast model with the smallest MAPE in the evaluation region (in-sample or holdout-sample) is chosen as the best model.

When a particular statistic of fit is used to judge how well the forecasting process is predicting the future, it is referred to as the performance statistic.

Descriptions of Statistics of Fit

You can use statistics of fit to measure how well different models fit the data. The statistics of fit for the various forecasting models can be printed or stored in a data set.

The definitions and formulas for the statistics of fit that are available in SAS Forecast Studio are described below. In these formulas, n is the number of nonmissing observations and k is the number of fitted parameters in the model. $APE = |100 * (y_t - \hat{y}_t)/y_t|$ is the absolute percent error. $ASPE = |100 * (y_t - \hat{y}_t)/0.5(y_t + \hat{y}_t)|$ is the absolute symmetric percent error. $APPE = |100 * (y_t - \hat{y}_t)/\hat{y}_t|$ is the absolute predictive percent error. $RAE = |(y_t - \hat{y}_t)/(y_t - y_{t-1})|$ is the relative absolute error. The errors are ignored in the statistical computations when the denominator is zero.

Adjusted R-square (ADJRSQ)
: The adjusted R^2 statistic, $1 - (\frac{n-1}{n-k})(1 - R^2)$.

Akaike Information Criterion (AIC)
: Akaike's information criterion, $n \ln(SSE/n) + 2k$.

Amemiya's Adjusted R-square (AADJRSQ)
: Amemiya's adjusted R^2, $1 - (\frac{n+k}{n-k})(1 - R^2)$.

Amemiya's Prediction Criterion (APC)
: Amemiya's prediction criterion, $\frac{1}{n}SST(\frac{n+k}{n-k})(1 - R^2) = (\frac{n+k}{n-k})\frac{1}{n}SSE$.

Geometric Mean Absolute Error Percent of Standard Deviation (GMAPES)
: The geometric mean of the absolute error as a percentage of the standard deviation.

Geometric Mean Percent Error (GMAPE)
: The geometric mean percent error.

Geometric Mean Predictive Percent Error (GMAPPE)
: The geometric mean absolute predictive percent prediction error.

Geometric Mean Relative Absolute Error (GMRAE)
: The geometric mean of the relative absolute errors.

Geometric Mean Symmetric Percent Error (GMASPE)
: The geometric mean of the absolute symmetric percent errors.

Maximum Absolute Error Percent of Standard Deviation (MAXAPES)
: The maximum of the absolute error as a percentage of the standard deviation.

Maximum Error (MAXERR)
: The largest prediction error.

Maximum Percent Error (MAXPE)
: The largest percent prediction error, $100 \max((y_t - \hat{y}_t)/y_t)$. The summation ignores observations where $y_t = 0$.

Maximum Predictive Percent Error (MAXPPE)
: The maximum of the predictive percent errors.

Maximum Relative Error (MAXRE)
: The maximum of the relative errors.

Maximum Symmetric Percent Error (MAXSPE)
: The maximum of the symmetric percent errors.

Mean Absolute Error (MAE)
: The mean absolute prediction error, $\frac{1}{n}\sum_{t=1}^{n}|y_t - \hat{y}_t|$.

Mean Absolute Error Percent of Standard Deviation (MAPES)
: The mean of the absolute error as a percentage of the standard deviation.

Mean Absolute Percent Error (MAPE)
: The mean of the absolute percent errors.

Mean Absolute Scaled Error (MASE)
: The mean of the absolute scaled errors.

Mean Error (ME)
: The mean prediction error, $\frac{1}{n}\sum_{t=1}^{n}(y_t - \hat{y}_t)$.

Mean Percent Error (MPE)
: The mean percent prediction error, $\frac{100}{n}\sum_{t=1}^{n}\frac{(y_t-\hat{y}_t)}{y_t}$. The summation ignores observations where $y_t = 0$.

Mean Predictive Percent Error (MPPE)
: The mean of the predictive percent error.

Mean Relative Absolute Error (MRAE)
: The mean of the relative absolute errors.

Mean Relative Error (MRE)
: The mean of the relative errors.

Mean Square Error (MSE)
: The mean squared prediction error calculated from the one-step-ahead forecasts. $MSE = \frac{1}{n}SSE$. This formula enables you to evaluate small holdout samples.

Mean Symmetric Percent Error (MSPE)
: The mean of the symmetric percent errors.

Median Absolute Error Percent of Standard Deviation (MDAPES)
: The median of the absolute error as a percentage of the standard deviation.

Median Percent Error (MDAPE)
: The median of the percent errors.

Median Predictive Percent Error (MDAPPE)
: The median of the predictive percent errors.

Median Relative Absolute Error (MDRAE)
: The median of the relative absolute errors.

Median Symmetric Percent Error (MDASPE)
: The median of the symmetric percent errors.

Minimum Absolute Error Percent of Standard Deviation (MINAPES)
: The minimum of the absolute error as a percentage of the standard deviation.

Minimum Error (MINERR)
: The smallest prediction error.

Minimum Percent Error (MINPE)
: The smallest percent prediction error, $100 \min((y_t - \hat{y}_t)/y_t)$. The summation ignores observations where $y_t = 0$.

Minimum Predictive Percent Error (MINPPE)
: The smallest predictive percent error.

Minimum Relative Error (MINRE)
: The smallest relative error.

Minimum Symmetric Percent Error (MINSPE)
: The smallest symmetric percent error.

Number of Missing Actual Values (NMISSA)
: The number of missing actual values.

Number of Missing Predicted Values (NMISSP)
: The number of missing predicted values.

Number of Observations (NOBS)
: The total number of observations used to fit the model, including both missing and nonmissing observations.

Number of Parameters (NPARM)
: The number of parameters fit to the data. For combined forecast, this is the number of forecast components.

Root Mean Square Error (RMSE)
: The root mean square error, $\sqrt{MSE}$.

R-square (RSQUARE)
: The R^2 statistic, $R^2 = 1 - SSE/SST$. If the model fits the series badly, the model sum of squares error, SSE, may be larger than SST and the R^2 statistic will be negative.

Random Walk R-square (RWRSQ)
: The random walk R^2 statistic (Harvey's R^2 statistic using the random walk model for comparison), $1 - (\frac{n-1}{n})SSE/RWSSE$, where $RWSSE = \sum_{t=2}^{n} (y_t - y_{t-1} - \mu)^2$, and $\mu = \frac{1}{n-1} \sum_{t=2}^{n} (y_t - y_{t-1})$.

Sample Size (N)
: The size of the sample.

Schwarz Bayesian Information Criterion (SBC)
: Schwarz Bayesian information criterion,
$n \ln(SSE/n) + k \ln(n)$.

Sum of Squares Error (SSE)
: The sum of the squared prediction errors. $SSE = \sum_{t=1}^{n} (y_t - \hat{y}_t)^2$, where $\hat{y}$ is the one-step predicted value.

Symmetic Mean Absolute Percent Error (SMAPE)
: The symmetric mean of the absolute percent error.

Symmetric Mean Absolute Predictive Percent Error (MAPPE)
: The mean of the absolute symmetric predictive percent error.

Total Corrected Sum of Squares for the dependent variable
: The total sum of squares for the series corrected for the mean: $\sum_{t=1}^{n} (y_t - \overline{y})^2$, where $\overline{y}$ is the series mean.

Total Sum of Squares (SST)
: The total sum of squares for the series, uncorrected for the mean: $\sum_{t=1}^{n} y_t^2$.

Unbiased Mean Square Error (UMSE)
: The unbiased mean square error.

Unbiased Root Mean Square Error (URMSE)
: The unbiased root mean square error.

Additional Information

For more information about statistics of fit, see the "Forecasting Process Details" chapter in the *SAS High-Performance Forecasting User's Guide*.

Appendix D

Default Model Selection Lists

Contents

Overview of the Default Model Selection Lists

Although you can create your own model, SAS Forecast Studio ships with several lists of candidate models that can be added to the project during the project setup. These model selection lists are saved in several SAS catalogs in the SASHELP library.

When you use a model, that model is automatically added to the model repository for the project. After a model is added to the repository, you can add a model to any series in the project. To view the complete list of models that are available for the current project, select **Project** → **Model Repository**.

Best Smoothing Model

The model selection list for the Best Smoothing Model is saved in the SASHELP.HPFDFLT.BEST catalog. The following models are in this selection list:

- Damped Trend Exponential Smoothing
- Double (Brown) Exponential Smoothing
- Linear (Holt) Exponential Smoothing

- Seasonal Exponential Smoothing
- Simple Exponential Smoothing
- Winters (Additive) Method
- Winters (Multiplicative) Method

Best Nonseasonal Smoothing Model

The model selection list for the Best Nonseasonal Model is saved in the SASHELP.HPFDFLT.BESTN catalog. The following models are in this selection list:

- Seasonal Exponential Smoothing
- Winters (Additive) Method
- Winters (Multiplicative) Method

Best Seasonal Smoothing Model

The model selection list for the Best Seasonal Smoothing Model is saved in the SASHELP.HPFDFLT.BESTS catalog. The following models are in this selection list:

- Damped Trend Exponential Smoothing
- Double (Brown) Exponential Smoothing
- Linear (Holt) Exponential Smoothing
- Simple Exponential Smoothing

Time Series Forecasting Models

The model selection list for the time series forecasting models is saved in the SASHELP.HPFDFLT.TSFSELECT catalog. The following models are in this selection list:

- Airline Model
- ARIMA(0,1,1)s NOINT

- ARIMA(0,1,1)(1,0,0)s NOINT
- ARIMA(0,1,2)(0,1,1)s NOINT
- ARIMA(0,2,2)(0,1,1)s NOINT
- ARIMA(2,0,0)(1,0,0)s
- ARIMA(2,1,0)(0,1,1)s NOINT
- ARIMA(2,1,2)(0,1,1)s NOINT
- Damped Trend Exponential Smoothing
- Double (Brown) Exponential Smoothing
- Linear (Holt) Exponential Smoothing
- Linear Trend
- Linear Trend with Autoregressive Errors
- Linear Trend with Seasonal Terms
- Log Airline Model
- Log ARIMA(0,1,1)s NOINT
- Log ARIMA(0,1,1)(1,0,0)s NOINT
- Log ARIMA(0,1,2)(0,1,1)s NOINT
- Log ARIMA(0,2,2)(0,1,1)s NOINT
- Log ARIMA(2,0,0)(1,0,0)s
- Log ARIMA(2,1,0)(0,1,1)s NOINT
- Log ARIMA(2,1,2)(0,1,1)s NOINT
- Log Damped Trend Exponential Smoothing
- Log Double (Brown) Exponential Smoothing
- Log Linear (Holt) Exponential Smoothing
- Log Linear Trend
- Log Linear Trend with Autoregressive Errors
- Log Linear Trend with Seasonal Terms
- Log Mean
- Log Random Walk with Drift
- Log Seasonal Dummy

- Log Seasonal Exponential Smoothing
- Log Simple Exponential Smoothing
- Log Winters Method – Additive
- Log Winters Method – Multiplicative
- Mean
- Random Walk with Drift
- Seasonal Dummy
- Seasonal Exponential Smoothing
- Simple Exponential Smoothing
- Winters (Additive) Method
- Winters (Multiplicative) Method

Appendix E

Sample Reports in SAS Forecast Studio

Contents

Overview of Reports

The sample reports that ship with SAS Forecast Studio are available from the Reports and Stored Processes dialog box. These reports are organized into the following categories:

- Getting Started
- Transactional Statistical Reports
- Time Series Reports
- Time Series Analysis Reports

- Statistical Model Parameter Analysis Reports
- Statistical Model Evaluation Analysis Reports
- Statistical Model Forecasts Reports
- Statistical Model Forecast Summary Reports
- Reconciled Evaluation Reports
- Joint Statistical Forecasts Reports
- Reporting Variables Reports
- Statistical Model Component Reports
- OLAP
- Reconciled Forecasts Reports
- Final Forecast Reports
- Override Reports
- Forecast Quality Reports

For more information about each report, see "Description of Sample Reports" on page 244.

Description of Sample Reports

Getting Started Reports

The following table gives a brief explanation of the Getting Started reports and lists the parameters that you can customize at run time.

Name of Report	Description
Simple Example of a Stored Process	Demonstrates an example of a stored process by creating a simple HTML report. When you run this report, you can customize the title of the report.
Import Excel Table	Enables you to import a Microsoft Excel spreadsheet to a SAS data set. When you run this report, you must specify the location of this spreadsheet. **NOTE:** To run this report, you must have the FS_LIBNAME defined. If you do not have this LIBNAME statement defined, then you cannot run this report.

Name of Report	Description
Node Macro Definition Export File	Enables you to export to an external file the macro definitions for the selected node of the hierarchy. When you run this report, you must specify the path and name of this external file.
Getting Started with Forecast Studio Reports	Creates an HTML report that contains the results of the SAS Forecast Studio project. The report contains the following data sets: Time Series, Time Series Summary, Model Selection, Model Parameter, Model Components, Model Extended, Model Forecast, Model Statistics, Reconciled Forecasts, and Reconciled Statistics.

NOTE: The macros that are listed in the **Macros** node are sample macros and are not intended to be run.

Transactional Statistics Reports

The following table gives a brief explanation of the Transactional Statistics reports and lists the parameters that you can customize at run time.

Name of Report	Description
Node Trend Statistics Table	Creates an HTML report that contains a table of the trend statistics (for example, the mean, the minimum value, the maximum value) for the selected node in the hierarchy. The trend statistics are displayed for each time period for the series that generated the data at the selected node. The report also includes the global statistics for the dependent variable. When you run this report, you can specify the name of this report, and you can select the trend statistics to include in the report.
Node Seasonal Statistics Table	Creates an HTML report that contains a table of the seasonal statistics (for example, the mean, the minimum value, the maximum value) for the selected node of the hierarchy. The seasonal statistics are displayed for each time period for the series that generated the data at the selected node. The report also includes the global statistics for the dependent variable. When you run this report, you can specify the name of this report, and you can select the seasonal statistics to include in the output.

Time Series Reports

These reports contain plots of the time series or dependent variable to be forecast. The following table gives a brief explanation of the Time Series reports and lists the parameters that you can customize at run time.

Name of Report	Description
Node Series Plot	Creates an HTML report that contains a time series plot for the selected node. When you run this report, you can specify the title of the report.
Node Cross Series Plot	Creates an HTML report that contains a plot of the time series and independent variables for the selected node. When you run this report, you must select one independent variable to include in the report. Optionally, you can change the title of the report. **NOTE:** To run this report, you must have assigned at least one independent variable. For more information, see "Modify Variable Assignments" on page 49.
Branch Series Plot	Creates an HTML report that contains a time series plot for the selected node and for the lower-level nodes in the same branch of the hierarchy. When you run this report, you can change the title of the report. You can also select whether SAS Forecast Studio should use the same scale for the Y-axis of each graph. Using the same scale enables you to compare values across graphs. By default, SAS Forecast Studio scales the series.
Path Series Plot	Creates an HTML report that contains a time series plot for the selected node and for the upper-level nodes in the same path of the hierarchy. When you run this report, you can change the title of the report.
Path Cross Series Plot	Creates an HTML report that contains a cross-series plot for the selected node and for the upper-level nodes in the same path of the hierarchy. When you run this report, you must select one independent variable to include in the report. Optionally, you can change the title of the report. **NOTE:** To run this report, you must have assigned at least one independent variable. For more information, see "Modify Variable Assignments" on page 49.

Name of Report	Description
Hierarchical Series Plot	Creates an HTML report that contains the time series plots for the selected node and for lower-level and upper-level nodes in the hierarchy. This report contains a series plot for the current node, a series plot for the child branch, and a table that displays the value of the BY variables for that branch. When you run this report, you can specify the titles of the graphs. You can also select whether to use the same scale for the child series. Using the same scale enables you to compare values across graphs. By default, SAS Forecast Studio scales the series.
Level Descriptive Statistics Table	Creates an HTML report that contains a descriptive statistics table for the selected level of the hierarchy. When you run this report, you can specify the following information: • the title of the table • the descriptive statistic that you want to use to sort the table • the sort order (either ascending or descending) • the minimum and maximum number of observations, nonmissing observations, and missing observations • the minimum and maximum of the minimum series values, the maximum series values, and the mean series values

Time Series Analysis Reports

These reports contain general diagnostic plots and statistics that are associated with time series modeling. The following table gives a brief explanation of the Time Series Analysis reports and lists the parameters that you can customize at run time.

Name of Report	Description
Node Time Series Analysis Report	Creates an analysis report of the time series for the selected node in the hierarchy. The HTML report contains the following information: • season statistics for the dependent variable • descriptive statistics for the time series • seasonal decomposition for the dependent variable • series plot for the dependent variable • series histogram • seasonal cycle plot for the dependent variable • an ACF plot and a standardized ACF plot for the dependent variable • a PACF plot and a standardized PACF plot for the dependent variable • an IACF plot and a standardized IACF plot for the dependent variable • a white noise probability plot and a white noise probability plot (log scale) for the dependent variable • a seasonally adjusted series plot for the dependent variable • a percent change adjusted series plot for the dependent variable • a trend-cycle component plot for the dependent variable • a seasonal-irregular component plot for the dependent variable • a seasonal component plot for the dependent variable

Name of Report	Description
Node Time Series Analysis Report (cont.)	The HTML report also contains the following information: • a trend-cycle-seasonal plot for the dependent variable • an irregular component plot for the dependent variable • a trend component plot for the dependent variable • a cycle component plot for the dependent variable • a statistics summary for the dependent variable When you run this report, you can customize the title of the report.
Node Decomposition Analysis Report	Creates a decomposition analysis report for the selected node in the hierarchy. The HTML report contains the following information: • the seasonal decomposition plots that you selected • a table that shows for each time period the value of the original series and the value of the seasonally adjusted series When you run this report, you can customize or select the following parameters: • the title of the report • the plots to include in the report • the decomposition mode • the Hodrick-Prescott Filter Parameter • the decomposition components • the series transformation • the simple differencing order • the seasonal differencing order • how to interpret missing values • how to interpret zero values

Name of Report	Description
Node Autocorrelation Analysis Report	Creates an autocorrelation analysis report for the selected node in the hierarchy. The HTML report contains the following information: • the plots that you selected for the dependent variable • a table that shows the autocorrelations for each time lag When you run this report, you can customize or select the following parameters: • the title of the report • the plots to include in the report • the autocorrelation statistics • the number of lags • the number of parameters • the series transformation • the simple differencing order • the seasonal differencing order • how to interpret missing values • how to interpret zero values

Name of Report	Description
Node Cross Correlation Analysis Report	Creates a cross-correlation analysis report for the selected node of the hierarchy. When you run this report, you can customize the following parameters: • the title of the report • the independent variable to include in the report • the cross-correlation statistics • the number of lags • the number of parameters • the series and the cross series transformation • the simple differencing order • the seasonal differencing order • how to interpret missing values • how to interpret zero values
Node Seasonal Cycles Plot	Creates an HTML report that contains the seasonal cycles plot for the selected node in the hierarchy. When you run this report, you can customize the following parameters: • the title of the report • the decomposition mode • the decomposition component • whether to scale the series
Level Seasonal Cycles Plot	Creates an HTML report that contains the seasonal cycles plot for the selected level in the hierarchy. When you run this report, you can customize the following parameters: • the title of the report • the decomposition mode • the decomposition component • whether to scale the series

Name of Report	Description
Node Stepwise Regression Analysis Report	Creates an HTML report that contains the stepwise regression analysis report for the selected level in the hierarchy. When you run this report, you can customize the following parameters: • the title of the report • whether to include the intercept in the analysis • the number of dependent and independent variable lags • the series transformation and cross-series transformation options • the selection parameters, such as the selection method and the significance levels for entry and for staying
Node Stepwise Autoregression Analysis Report	Creates an HTML report that contains the stepwise autoregression analysis report for the selected level in the hierarchy. When you run this report, you can customize the following parameters: • the title of the report • whether to include the intercept in the analysis • the number of dependent variable lags • the series transformation options • the selection parameters, such as the selection method and the significance levels for entry and for staying

Name of Report	Description
Node X12-ARIMA Analysis Report	Creates an HTML report that contains the X12-ARIMA analysis report for the selected level in the hierarchy. When you run this report, you can customize the following parameters: • the title of the report • the transformation options • the X11 options • the ARIMA options • the regression options • the outlier options • the estimation options
Node Stationarity Analysis Report	Creates an HTML report that contains the stationarity analysis report for the selected level in the hierarchy. When you run this report, you can customize the following parameters: • the title of the report • the number of dependent variable lags • the series transformation options
Node GARCH Autoregression Analysis Report	Creates an HTML report that contains the GARCH autoregression analysis report for the selected level in the hierarchy. When you run this report, you can customize the following parameters: • the title of the report • whether to include the intercept in the analysis • the number of dependent variable lags • the series transformation options • the selection parameters • the GARCH options • the estimation options

Name of Report	Description
Node GARCH Regression Analysis Report	Creates an HTML report that contains the GARCH regression analysis report for the selected level in the hierarchy. When you run this report, you can customize the following parameters: • the title of the report • whether to include the intercept in the analysis • the number of dependent variable and independent variable lags • the series transformation options • the cross-series transformation options • the selection parameters • the GARCH options • the estimation options

Statistical Model Parameter Analysis Reports

These reports contain the parameters estimates that are associated with the forecast models. The following table gives a brief explanation of the Statistical Model Parameter Analysis reports and lists the parameters that you can customize at run time.

Name of Report	Description
Node Parameter Estimates Table	Creates an HTML report that contains a table of the parameter estimates for the selected node in the hierarchy. This report lists the parameter estimates for each component (LEVEL, TREND, and SEASON). The report also includes the standard error, *t*-values, and *p*-values of the parameter estimates. When you run this report, you can customize the title of the report.
Level Parameter Estimates Table	Creates an HTML report that contains a table of the parameter estimates for the selected level in the hierarchy. This report lists the parameter estimates for each component (LEVEL, TREND, and SEASON). The report also includes the standard error, *t*-values, and *p*-values of the parameter estimates. When you run this report, you can customize the title of the report.

Name of Report	Description
Level Parameter Estimates Univariate Plot	Creates an HTML report that contains a univariate plot of the parameter estimates that are associated with a common independent variable for the forecast models at the selected level in the hierarchy. When you run this report, you can customize or select the following parameters: • The title of the report. • The independent variables to include in the report. You must select at least one independent variable, but you can have no more than two independent variables in the report. • The significance level.
Level Parameter Estimates Excel Table	Exports to Microsoft Office a table of the parameter estimates for the selected level in the hierarchy. The content of this table is the same as the output when you run the report for the Level Parameter Estimates Table. However, instead of being displayed in HTML, the output is displayed in a Microsoft Excel spreadsheet. By default, this content is saved as a CSV file.

Statistical Model Evaluation Analysis Reports

These reports contain the evaluation criteria that are associated with the forecast models. The following table gives a brief explanation of the Statistical Model Evaluation Analysis reports and lists the parameters that you can customize at run time.

Name of Report	Description
Node Model Statistics Table	Creates an HTML report that contains a table of the model statistics (for example, the fit statistics) for the time series and the forecast model at the selected node. When you run this report, you can customize the title of the report.
Level Model Statistics Table	Creates an HTML report that contains a table of the model statistics for each node in the selected level in the hierarchy. When you run this report, you can customize the title of the report.

Name of Report	Description
Level Model Statistics Univariate Plot	Creates an HTML report that contains a plot of the distribution of fit statistics that are associated with the forecast models at the selected level. When you run this report, you can customize the title of the report, and you select the model statistics that you want to plot.
Level Model Statistics Excel Table	Exports to Microsoft Excel a table of the model estimates for the selected level in the hierarchy. The content of this table is the same as the output when you run the report for the Level Model Estimates Table. However, instead of being displayed in HTML, the output is displayed in a Microsoft Excel spreadsheet. By default, this content is saved as a CSV file.
Hierarchical Model Statistics Table	Creates an HTML report that contains a table of the model statistics for each node in the hierarchy. This report contains the statistics for the parent node, the current node, and the branches in the child node. When you run this report, you can customize the title for each table in the report.

Statistical Model Forecasts Reports

These reports contain the forecast values and associated statistics that SAS Forecast Studio generates for the project. The following table gives a brief explanation of the Statistical Model Forecasts reports and lists the parameters that you can customize at run time.

Name of Report	Description
Node Model Forecast Plot	Creates an HTML report that contains a plot of the one-step ahead and lead forecasts for the selected node. When you run this report, you can customize the title of the report.
Node Model Forecast Table	Creates an HTML report that contains the forecast values of the lead for the selected node in the hierarchy. This report contains the predicted values for each time period in the selected node. The output also includes the lower confidence limits, the upper confidence limits, and the prediction standard errors. When you run this report, you can customize the title of the report, and you can select whether to include the historical data for the node in the report. If you include the historical data, then the report includes the actual values and the prediction errors.

Name of Report	Description
Node Model Forecast Excel Table	Exports to Microsoft Excel a forecast table for the selected node in the hierarchy. The content of this table is the same as the output when you run the report for the Node Model Forecast Table. However, instead of being displayed in HTML, the output is displayed in a Microsoft Excel spreadsheet. By default, this content is saved as a CSV file. When you run this report, you can select whether to include the historical data for the node in the report. If you include the historical data, then the report includes the actual values and the prediction errors.
Branch Model Forecast Plot	Creates an HTML report that contains a plot of the one-step ahead and lead forecasts for the selected node and for the lower-level nodes in the same branch of the hierarchy. When you run this report, you can customize the title of the report. You can also select whether to scale the forecasts. Scaling the forecasts enables you to compare values across graphs. By default, SAS Forecast Studio scales the forecast values.
Path Model Forecast Plot	Creates an HTML report that contains a plot of the one-step ahead and lead forecasts for the selected node and for the upper-level nodes in same path of the hierarchy. When you run this report, you can customize the title of the report.
Node Model Forecast Percent Change Plot	Creates an HTML report that contains a plot that displays the percentage change of the difference between the mean and the forecasts of the time series. This report displays a plot of the percent change for the future and historical data. When you run this report, you can customize the title of each plot.
Node Model Forecast Percent Change Table	Creates an HTML report that contains a table that displays the percentage change of the difference between the mean and the forecasts of the time series. When you run this report, you can customize the title of the report, and you can select whether to include the historical data in the report.
Node Model Forecast Analysis Report	Creates an HTML report that contains an analysis report of the forecasts for the selected node in the hierarchy. When you run this report, you can customize the title of the report, and you can select what plots and tables to include in the report. You can also choose whether to include detailed information in the tables.

Name of Report	Description
Branch Model Forecast Plot	Creates an HTML report that contains a forecast plot for the selected node and for the lower-level nodes in the same branch of the hierarchy. When you run this report, you can customize the title of the report, and you can specify the number of future time periods to include in the plot.

Statistical Model Forecast Summary Reports

These reports contain the forecast values for the project. The following table gives a brief explanation of the Statistical Model Forecast Summary reports and lists the parameters that you can customize at run time.

Name of Report	Description
Node Model Forecast Summary Table	Creates an HTML report that contains a summary table of the forecast values for the selected node in the hierarchy. This report lists the forecast values for the dependent variable for each time period. When you run the report, you can customize the title of the report.
Level Model Forecast Summary Table	Creates an HTML report that contains a summary table of the forecast values for the selected level in the hierarchy. This report lists the forecast values for each time period. When you run the report, you can customize the title of the report.
Level Model Forecast Summary Excel Table	Exports to Microsoft Excel a summary table of the forecast values for the selected level in the hierarchy. The content of this table is the same as the output when you run the report for the Level Model Forecast Summary Table. However, instead of being displayed in HTML, the output is displayed in a Microsoft Excel spreadsheet. By default, this content is saved as a CSV file.

Reconciled Evaluation Reports

These reports contain information that is associated with the reconciliation of the forecasts. The following table gives a brief explanation of the Reconciled Evaluation reports and lists the parameters that you can customize at run time.

Name of Report	Description
Node Reconciled Statistics Table	Creates an HTML report that contains a table of the statistics (for example, fit statistics) for the time series and reconciled forecasts at the selected node. When you run this report, you can customize the title of the report.
Level Reconciled Statistics Table	Creates an HTML report that contains a table of the statistics (for example, fit statistics) for the time series and reconciled forecasts for each node in the selected level. When you run this report, you can customize the title of the report.
Level Reconciled Statistics Univariate Plot	Creates an HTML report that contains a plot of the distribution of fit statistics for the reconciled forecasts for the selected level. This report also contains a univariate plot for each reconciled statistic that you select. When you run this report, you can customize the title, and you can select the reconciled statistics that you want to plot.
Level Reconciled Statistics Excel Table	Exports to Microsoft Excel a table that contains the statistics for the time series and reconciled forecasts for the selected level. The content of this table is the same as the output when you run the report for the Level Reconciled Statistics Table. However, instead of being displayed in HTML, the output is displayed in a Microsoft Excel spreadsheet. By default, this content is saved as a CSV file.
Hierarchy Reconciled Statistics Table	Creates an HTML report that contains tables of the statistics for the time series and reconciled forecasts for the hierarchy. This report includes tables for the parent node, the current node, and the branches in the child node. When you run this report, you can customize the title for each table in the report.

Joint Statistical Forecasts Reports

These reports contain information about the difference between the reconciled and unreconciled forecasts. The following table gives a brief explanation of the Joint Statistical Forecasts reports and lists the parameters that you can customize at run time.

Name of Report	Description
Node Joint Forecast Plot	Creates an HTML report that contains the following graphs for the selected node in the hierarchy: • a forecast plot that shows the predicted and reconciled values • a plot that shows the reconciliation adjustment for each data point • a plot that shows the percent change of reconciliation adjustment for each data point When you run this report, you can customize the title of the plots.
Node Joint Statistics Table	Creates an HTML report that contains the following tables for the selected node in the hierarchy: • a model statistics table • a table that contains the statistics for the reconciled and unreconciled forecasts When you run this report, you can customize the title of the tables.

Reporting Variables Reports

These reports contain the statistics for the reporting variable. To run this report, you must have assigned at least one reporting variable. For more information, see "Modify Variable Assignments" on page 49.

The following table gives a brief explanation of the Reporting Variables reports and lists the parameters that you can customize at run time.

Name of Report	Description
Node Report Series Plot	Creates an HTML report that contains a series plot for the reporting variable in the selected node in the hierarchy. When you run this report, you must select one reporting variable to include in the output. Optionally, you can also customize the title of the report.

Name of Report	Description
Node Report Series Table	Creates an HTML report that contains a series table for the reporting variables in the selected node in the hierarchy. When you run this report, you must select at least one reporting variable to include in the output. Optionally, you can also customize the title of the report and specify whether to include historical data in the output.
Node Report Series Excel Table	Exports to Microsoft Excel a table that contains the values of the reporting variables for the selected node in the hierarchy. The content of this table is the same as the output when you run the report for the Node Report Series Table. However, instead of being displayed in HTML, the output is displayed in a Microsoft Excel spreadsheet. By default, this content is saved as a CSV file.

Statistical Model Component Reports

These reports contain the predicted values for the input and output variables that are used in the project. The following table gives a brief explanation of the Statistical Model Component reports and lists the parameters that you can customize at run time.

Name of Report	Description
Node Component Table	Creates an HTML report that contains a component table for the selected node in the hierarchy. This report displays the predicted value for the time series and independent variables for the selected node. When you run this report, you can customize the title of the report, and you can select whether to include historical estimates.
Node Component Plot	Creates an HTML report that contains a component plot for the selected node in the hierarchy. This report displays a plot of the predicted values for the model inputs and outputs at the selected node. When you run this report, you can customize the title of the report.

OLAP

The following table gives a brief explanation of the OLAP reports and lists the parameters that you can customize at run time.

Name of Report	Description
Export as OLAP Cube	Exports the project as an OLAP cube. When you run this report, you must specify a name for the cube, a description of the cube, the OLAP schema, and a hierarchy label.

Reconciled Forecasts Reports

These reports contain information about the reconciled forecasts in the project. The following table gives a brief explanation of the Reconciled Forecasts reports and lists the parameters that you can customize at run time.

Name of Report	Description
Node Reconciled Forecast Plot	Creates an HTML report that contains a plot of the reconciled forecast for the selected node. This plot shows the predicted values and shows the upper and lower confidence limits for the future values. When you run this report, you can customize the title of the plot.
Node Reconciled Forecast Table	Creates an HTML report that contains a table of the reconciled forecast values and the associated statistics (for example, standard errors and confidence limits) for the selected node. This table includes the predicted values, the lower and upper confidence limits for these values, the prediction standard errors, and the reconciliation status. When you run this report, you can customize the title of the plot. You can also select whether to include the historical data. If you include the historical data, then the table includes the actual values for each time period and the predicted errors.
Node Reconciled Forecast Excel Table	Exports to Microsoft Excel a table that contains the predicted values for the selected node in the hierarchy. The content of this table is the same as the output when you run the report for the Node Reconciled Forecast Table. However, instead of being displayed in HTML, the output is displayed in a Microsoft Excel spreadsheet. By default, this content is saved as a CSV file. When you run this report, you can select whether to include the historical data. If you include the historical data, then the table includes the actual values for each time period and the predicted errors.

Name of Report	Description
Path Reconciled Forecast Plot	Creates an HTML report that contains a plot of the reconciled forecast for the selected node and each upper-level node in the same path. When you run this report, you can customize the title of the report.
Node Reconciled Forecast Analysis Report	Creates an analysis report for the reconciled forecast values for the selected node in the hierarchy. This HTML report also contains the diagnostic plots and fit statistics for the reconciled forecasts. The number of plots and tables depends on the options that you select when you create the report. When you run this report, you can customize the following parameters: • the title of the report • the plots to include in the report • the tables to include in the report • whether to include the details in the tables
Branch Reconciled Forecast Plot	Creates an HTML report that contains a plot of the reconciled forecasts for the selected node and the lower-level nodes in the same branch of the hierarchy. When you run this report, you can customize the title of the report, and you can specify the number of future time periods to include in the plot.

Final Forecast Reports

These reports contain information on the final forecasts that SAS Forecast Studio generated for the project. The following table gives a brief explanation of the Final Forecast reports and lists the parameters that you can customize at run time.

Name of Report	Description
Node Final Forecast Plot	Creates an HTML report that contains plot of the final forecasts for the selected node in the hierarchy. When you run this report, you can customize the title of the plot.

Name of Report	Description
Node Final Forecast Table	Creates an HTML report that contains a table of the final forecasts for the selected node in the hierarchy. This table lists the predicted values, the lower and upper confidence limits for these values, the prediction standard errors, the reconciliation status, and the final predicted values. When you run this report, you can customize the title of the plot. You can also select whether to include the historical data. If you include the historical data, then the table includes the actual values for each time period and the predicted errors.
Node Final Forecast Excel Table	Exports to Microsoft Excel a table that contains the final forecasts for the selected node in the hierarchy. The content of this table is the same as the output when you run the report for the Node Final Forecast Table. However, instead of being displayed in HTML, the output is displayed in a Microsoft Excel spreadsheet. By default, this content is saved as a CSV file. When you run this report, you can select whether to include the historical data. If you include the historical data, then the table includes the actual values for each time period and the predicted errors.

Override Reports

These reports contain information about any overrides that you added to the project. The following table gives a brief explanation of the Override reports and lists the parameters that you can customize at run time.

Name of Report	Description
Level Override Table	Creates an HTML report that contains a table of all the overrides for the selected level in the hierarchy. When you run this report, you can customize the title of the table, and you can choose whether to include historical overrides.

Forecast Quality Reports

The following table gives a brief explanation of the Forecast Quality reports and lists the parameters that you can customize at run time.

Name of Report	Description
Level Model Forecast Quality Extreme Table	Creates an HTML report that contains a table of the number of extreme forecast values in the model for the selected level in the hierarchy. When you run this report, you can customize the title of the report. You can also specify the number of standard deviations to use in the analysis and the horizon slope.
Level Reconciled Forecast Quality Extreme Table	Creates an HTML report that contains a table of the number of extreme reconciled forecast values for the selected level in the hierarchy. When you run this report, you can customize the title of the report. You can also specify the number of standard deviations to use in the analysis and the horizon slope.
Level Model Forecast Quality Missing Table	Creates an HTML report that contains a table of the number of missing forecast values in the model for the selected level in the hierarchy. When you run this report, you can customize the title of the report.
Level Reconciled Forecast Quality Missing Table	Creates an HTML report that contains a table of the number of missing reconciled forecast values for the selected level in the hierarchy. When you run this report, you can customize the title of the report.

Glossary

accumulation
: either of two processes that are used to convert a time series. (1) Accumulation converts a time series that has no fixed interval into a time series that does have a fixed interval (such as hourly or monthly). (2) Accumulation converts a time series that has a fixed interval into a time series with a lower frequency time interval (such as hourly into daily). Accumulation combines data within the same time interval into a summary value for that time period.

aggregation
: the process of combining more than one time series to form a single series. Aggregation combines data within the same time interval. For example, you can aggregate data into a total or average.

aggregation statistic
: the mathematical operation used to combine forecasts across levels in the hierarchy. The reconciliation method that you choose determines the levels where the aggregation statistic is used. See also reconciliation method.

autocorrelation
: the correlation between observations at different lags in a time series. Autocorrelation coefficient values range from -1 to +1. When the autocorrelation coefficient value at a given lag is positive, the observations that are separated by that lag tend to move together.

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.

backtrace
: See stack trace.

bottom-up method of reconciliation
: a reconciliation method that uses the forecasts at the lowest level of the hierarchy to adjust forecasts for the higher levels in the hierarchy. See also middle-out method of reconciliation, reconciliation method, top-down method of reconciliation.

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.

disaggregation method
: a method that specifies how the forecasts in the lower level of the hierarchy are reconciled when the reconciliation method is top-down or middle-out. The disaggregation method can reconcile the forecasts in either of the following ways: (1) by using the proportion that each lower-level forecast contributes to the higher-level forecast; or (2) by splitting equally the difference between the higher-level forecast and the lower-level forecasts. See also middle-out method of reconciliation, top-down method of reconciliation.

dummy variable
: a numeric variable with a value of either 1 or 0. Dummy variables are used to indicate

whether or not unusual events occur. The variable takes the value of 1 during the event and 0 otherwise.

event
: an incident that disrupts the normal flow of any 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.

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. The holdout sample can be used to evaluate the forecasting performance of a candidate model.

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.

locked override
: a user-supplied value for a forecast that is honored when the hierarchy is reconciled. The final forecast value for that level is the value that you specified for the locked override. Locked overrides can generate override conflicts. See also unlocked override.

middle-out method of reconciliation
: a reconciliation method that combines the bottom-up method above the reconciliation level and the top-down method below the reconciliation level. When a hierarchy has more than one middle level, you need to specify which of those levels to use as the reconciliation level. The forecasts at the reconciliation level are used to generate forecasts for both the higher and lower levels. See also bottom-up method of reconciliation, reconciliation method, top-down method of reconciliation.

model selection criterion
: the statistic of fit 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.

override conflict
: a condition that occurs when the value of one locked override is incompatible with the value of another locked override in the same branch of the hierarchy. Override conflicts that are not resolved prior to reconciliation can result in unreconciled nodes. See also unreconciled node.

partial autocorrelation
: the internal correlation between observations in a time series that causes the effect of all intervening lags to be 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.

performance statistic
: a statistic of fit that is used to determine how well a forecasting process is predicting the future. See also statistic of fit.

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

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 method that specifies the level in the hierarchy where the process of reconciliation starts. 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 removing seasonality from time series data.

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

selection criterion
: a statistical value that is used to evaluate how well a forecasting model $x + y = z$ performs by comparing the actual data to the predictions. See also statistic of fit.

stack traces
: a log of the active stack frames that is created when a program ends abnormally. The last few stack frames often indicate where the error occurred that caused the program to abnormally terminate. Depending on how complex your application is, a stack trace can contain thousands of lines of diagnostics.

statistic of fit
: a statistical value that is used to evaluate how well a forecasting model fits the historical series by comparing the actual data to the predicted values.

time series data
: timestamped 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 of reconciliation
: a reconciliation method that uses the forecasts at the highest level of the hierarchy to adjust

the forecasts for the lower levels. See also bottom-up method of reconciliation, middle-out method of reconciliation, reconciliation method.

transactional data
: timestamped 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.

unlocked override
: a user-supplied value for a forecast that acts as a guideline for the final forecast value. The final forecast for the level reflects the value of the unlocked override, but the final forecast and the unlocked override are often not identical. Because these overrides can be overridden when the hierarchy is reconciled, unlocked overrides do not generate override conflicts. See also locked override.

unreconciled node
: a node in the hierarchy that cannot be reconciled. A node can be unreconciled if (1) the final forecast of an upper level is not equal to the aggregate of the final forecasts of the lower levels; or (2) a final forecast violates the limits that are implied by a locked override in one of the lower levels in the same branch of the hierarchy.

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

Your Turn

We want your feedback.

- If you have comments about this book, please send them to **yourturn@sas.com**. Include the full title and page numbers (if applicable).
- If you have comments about the software, please send them to **suggest@sas.com**.