SAS® Forecast Studio 1.4

User's Guide

THE POWER TO KNOW®

The correct bibliographic citation for this manual is as follows: SAS Institute Inc. 2006. *SAS ® Forecast Studio 1.4: User's Guide*. Cary, NC: SAS Institute Inc.

SAS® Forecast Studio 1.4: User's Guide

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

1st printing, November 2006

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

Contents

Part 1
About SAS Forecast Studio

Contents

Chapter 1

What's New in SAS Forecast Studio 1.4

Chapter Contents

Chapter 1
What's New in SAS Forecast Studio 1.4

Overview

SAS Forecast Studio 1.4 has the following changes and enhancements:

- improvements to the workspace
- enhancements to overrides
- new reconciliation options
- new forecasting options
- enhancements to events
- changes to the New Project Wizard
- new reports

Improvements to the Workspace

The following enhancements were made to the workspace:

- The Model View now includes the model selection list for the project. The model selection list shows the models that have been fitted to this series. For each model, this list displays the model name and the model selection criteria.
- You can now customize the contents of the Forecasting View. From the **Forecasting View** tab in the Customize View dialog box, you can select the information to show on the forecasting plot and in the data table.
- If SAS Forecast Studio detects that the hierarchy is unreconciled, then a "Reconciliation is out of date" warning appears at the top of the workspace. You can click **Update** for SAS Forecast Studio to reconcile the hierarchy.

Enhancements to Overrides

The following enhancements were made to the override functionality:

- You can now add overrides directly in the data table in the Forecasting View or by using the Override Calculator.
- You can also specify whether the overrides should be locked. SAS Forecast Studio honors locked overrides when reconciling the hierarchy. Unlocked overrides are used as a guideline and might not equal the final forecast after reconciliation.

- The new Override Conflicts dialog box enables you to resolve the override conflicts that SAS Forecast Studio detects during reconciliation.

New Reconciliation Options

The reconciliation functionality has significantly improved. These improvements enable you to get more accurate forecasts than in previous releases of SAS Forecast Studio.

The new advanced reconciliation options that are available when you configure a hierarchy enable you to do the following:

- Specify how you want to resolve conflicts between locked overrides. You can choose to force SAS Forecast Studio to resolve all conflicts before reconciliation, or you can choose to have SAS Forecast Studio ignore the conflicts.
- Specify whether to restrict the direction of the reconciliation. You can choose no restriction on the reconciliation direction. If you choose this option, then SAS Forecast Studio can use an implicit top-down method or bottom-up method of reconciliation. You can also choose to restrict the direction. If you choose this option, then SAS Forecast Studio reconciles the hierarchy based only on the reconciliation method that you selected. This option can result in unreconciled nodes.

Depending on the options that you select, the new Reconciliation Failures report might open. This report shows all of the unreconciled nodes in the hierarchy, and it can also include any override conflicts.

New Forecasting Options for the Project

The dialog box that displays the forecasting options for the project has been redesigned. The following options are new for this release:

- The accumulation options enable you to convert a time series. With the accumulation options, you can convert a time series that has no fixed interval into a time series that has a fixed interval. You can also convert a time series that has a fixed interval into a time series with a lower frequency time interval. Accumulation combines data within the same time interval into a summary value for that time period.
- The time ID options enable you to specify the interval of the values for the time ID variable, the seasonal cycle length, a multiplier, and a shift. You can also specify whether your dates include the time of day, a date format, and any inactive days. You can also edit the format of the time ID. This enables you to specify a user-defined format for the time ID.

Enhancements to Events

The following enhancements were made to events:

- You can now import events from a SAS data set that contains events data. This enables you to import events that are external to SAS Forecast Studio.
- You can also export events. 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 at your site can then import these events into projects that are running on the same SAS Workspace Server as the events repository.

Changes to the New Project Wizard

The following enhancements were made to the New Project Wizard:

- The order of some steps was changed to make it easier to create a project.
- The interface for specifying the forecasting options for the project has been improved. This improvement makes it easier for you to specify these options when you create the project.

New Reports

Several new reports were added to SAS Forecast Studio. These reports are organized into categories so that they are easier to find. The following categories contain new reports:

- In the Time Series Reports category, the Level Descriptive Statistics Table report is new.
- In the Time Series Analysis Reports category, the following reports are new:
  - Node Stepwise Regression Analysis Report
  - Node Stepwise Autoregression Analysis Report
  - Node X12-ARIMA Analysis Report
  - Node Stationary Analysis Report
  - Node GARCH Autoregression Analysis Report
  - Node GARCH Regression Analysis Report
- In the Statistical Model Forecast Reports category, the following reports are new:
  - Node Model Forecast Percent Change Table
  - Node Model Forecast Analysis Report
  - Branch Model Forecast Plot

- In the Reconciled Evaluation Reports category, the Level Reconciled Statistics Excel Table report is new.
- In the Reconciled Forecast Reports category, the Branch Reconciled Forecast Report is new.
- In the Override Reports category, the Level Override Table is new.
- In the Forecast Quality Reports category, the following reports are new:
  - Level Model Forecast Quality Extreme Table
  - Level Reconciled Forecast Quality Extreme Table
  - Level Model Forecast Quality Missing Table
  - Level Reconciled Forecast Quality Missing Table

Chapter 2
Introduction to SAS Forecast Studio

Chapter Contents

Chapter 2
Introduction to SAS Forecast Studio

Overview of SAS Forecast Studio

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

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

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

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

Using this application, you can do the following tasks:

- Generate forecasts automatically 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 1.4 supports with some exceptions the U.S. Section 508 software standards. SAS currently plans for SAS Forecast Studio software to have increased accessibility compliance in a future version. If you have specific questions about the accessibility of SAS Forecast Studio, send them to **accessibility@sas.com** or call SAS Technical Support.

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 that may be contained in a Help file, follow these steps:

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

Printing Help Topics

To print the current topic, 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

Chapter Contents

Chapter 3

Preparing an Input Data Set for SAS Forecast Studio

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 "Specifying the Accumulation Options" on page 111.

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 18.

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

1. The data is aggregated if the input is one of the following:
   - timestamped data that is recorded at no particular frequency (also called transactional data)
   - data recorded at a smaller time interval than needed for forecasting

2. Any gaps in the data are filled in. Gaps appear when there is not an observation for each time period or when the data is not equally spaced. The added observations have the required values of the time ID variable and the value that you specified for missing values. For more information, see "Working with Missing Values" on page 19.
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 or input 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.

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, "SAS Forecast Studio Variables."

For more information about variable roles, see "Assign Data to Roles" on page 60.

Examples

The following 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

For missing values in other variables, including the dependent, independent, and external forecast variables, you can specify how to interpret missing values (regardless of the variable roles) when you create a project. For more information, see "Set Forecasting Options" on page 64.

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 2
Working in SAS Forecast Studio

Contents

Chapter 4
Using the Workspace

Chapter Contents

Chapter 4
Using the Workspace

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.

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

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

Forecasting View
: displays the forecast plot and the data table for the series that you selected in the **Hierarchy** tab or the **Table** tab.

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 on the **Hierarchy** tab or the **Table** tab.

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

Status bar
: displays information about the current procedure, the username, and the SAS Forecast Server Mid-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
	enables you to archive a project
	reconciles the hierarchy
	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 Tab

What Is the Hierarchy Tab?

The **Hierarchy** tab 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. When you select a series on the **Hierarchy** tab, the corresponding series is selected on the **Table** tab.

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 70.

The following 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. For more information, see "Adding Overrides" on page 130.

  Note: The override icon persists even if you have removed the override from the data table. For example, suppose you change the horizon for a series, and this results in SAS Forecast Studio removing override values from the data table. (Remember you can only specify overrides for future values.) Although

the overrides are no longer in the data table, the override icons still appear in the **Hierarchy** tab next to the nodes that had the override values.

- If you add a note for a series, the  icon appears next to the series in the hierarchy. For more information, see "Add a Note to a Series" on page 132.
- If you use a filter to create an exception rule, the series that are identified as exceptions appear in red. For more information, see Chapter 6, "Working with Filters."

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

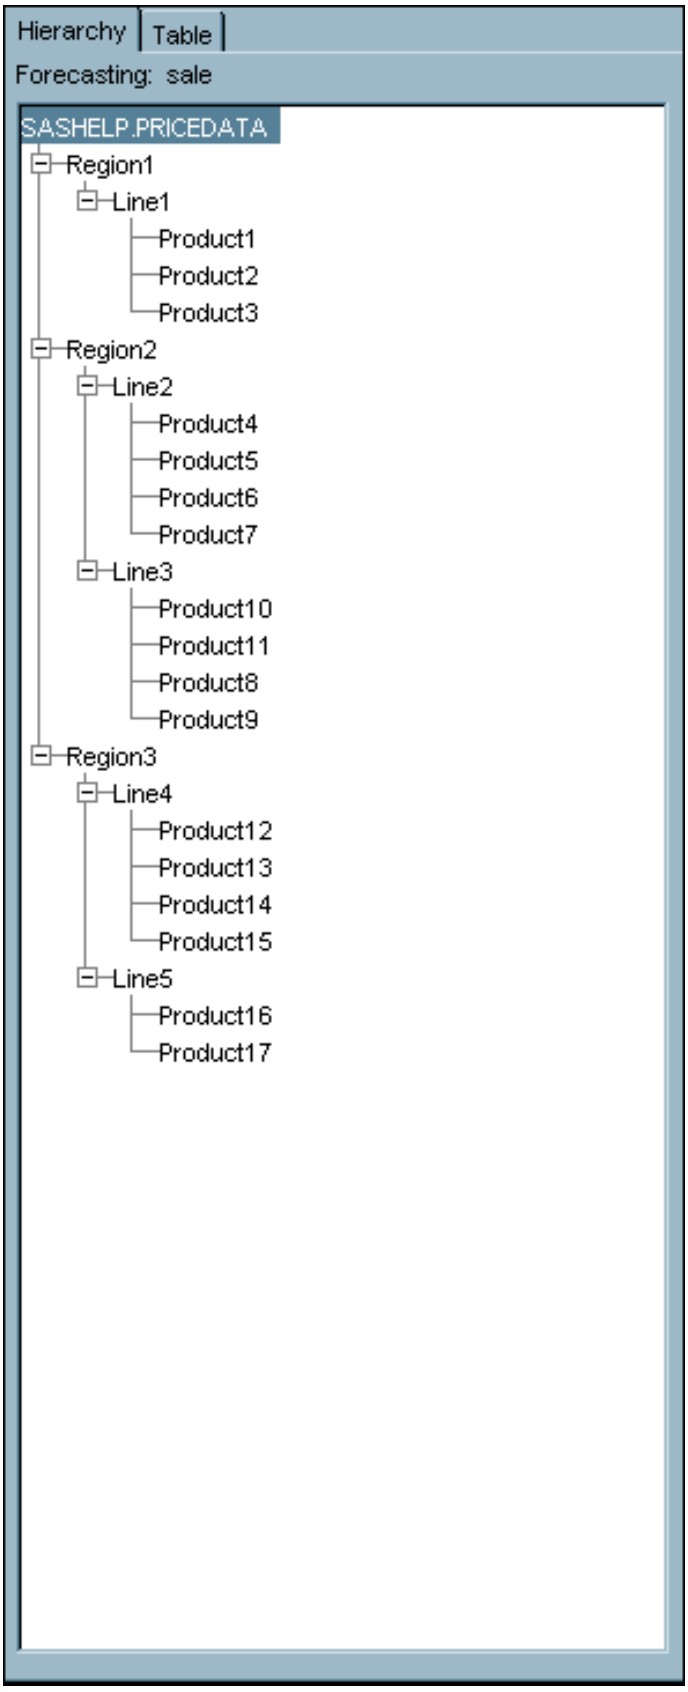

Figure 4.1. Hierarchy Tab

Customizing the Hierarchy Tab

To customize the **Hierarchy** tab, complete the following steps:

1. Select **View** → **Customize View**. The Customize View dialog box opens.
2. Click the **Hierarchy** tab.
3. Select information to display in the project hierarchy and click  .

   By default, the following items are displayed in the hierarchy:

   - **Notes**
   - **Overrides**
   - **Exceptions**

   Note: Click **Reset to Defaults** to return to the default values in the hierarchy.
4. Click **OK** to close the Customize View dialog box and to apply your changes.

About the Table Tab

What Is the Table Tab?

The **Table** tab contains the following information:

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

The **Table** tab enables you to compare the statistics of fit for each series. 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.

If you have used a filter to create an exception rule, the series that are flagged as exceptions appear in red. When you select a series on the **Table** tab, the corresponding series is selected on the **Hierarchy** tab.

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

For more information about how to customize the table, see "Customizing the Table Tab" on page 31.

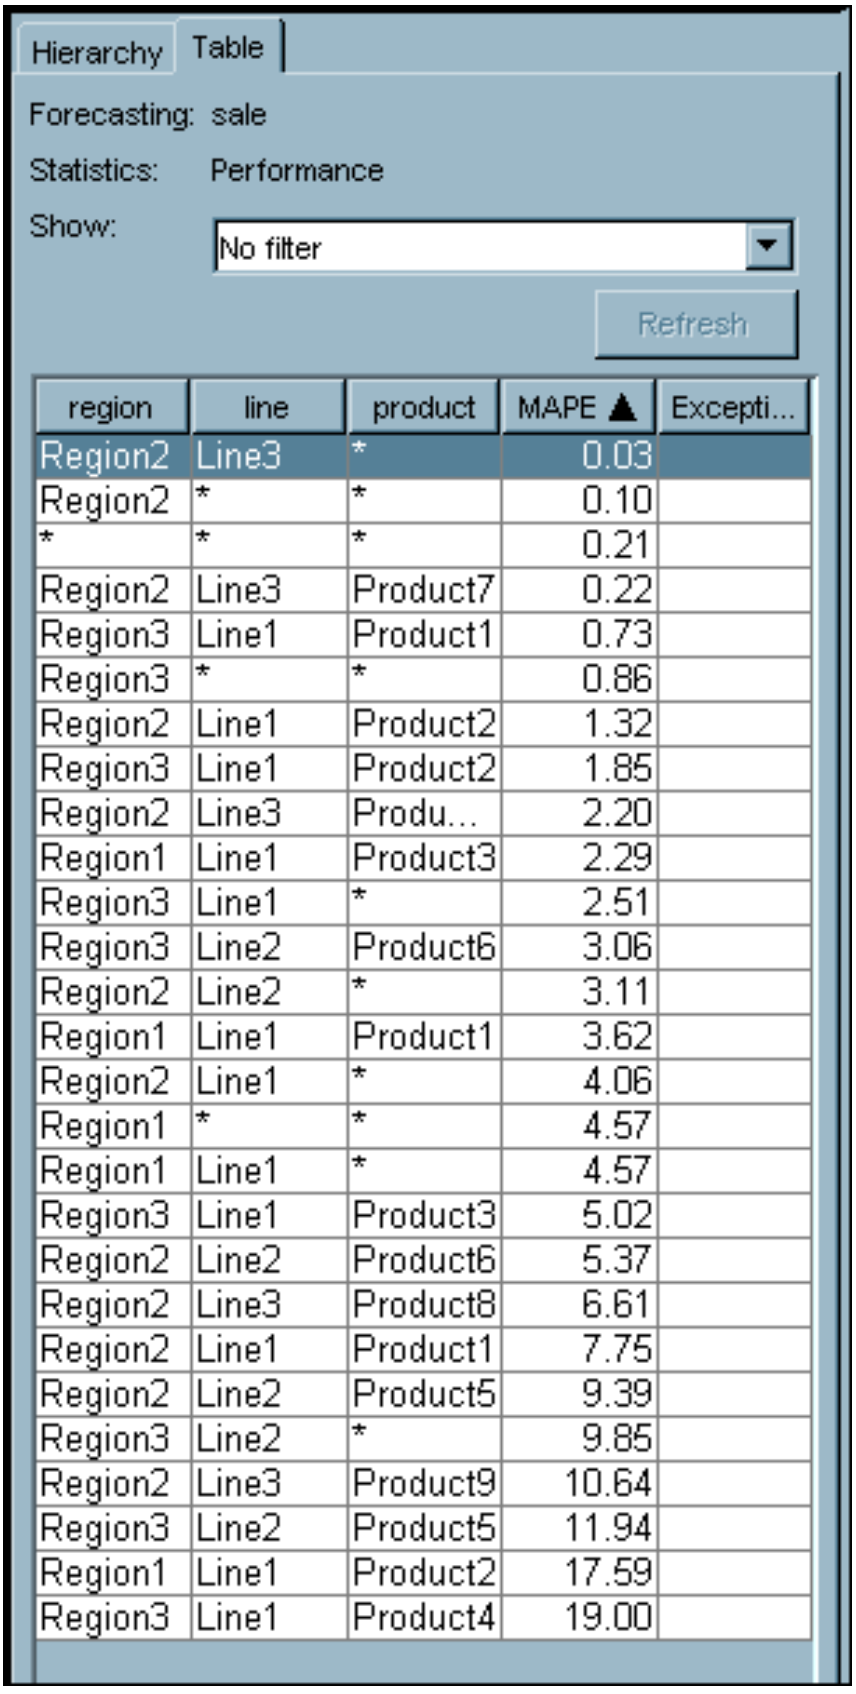

region	line	product	MAPE ▲	Excepti...
Region2	Line3	*	0.03	
Region2	*	*	0.10	
*	*	*	0.21	
Region2	Line3	Product7	0.22	
Region3	Line1	Product1	0.73	
Region3	*	*	0.86	
Region2	Line1	Product2	1.32	
Region3	Line1	Product2	1.85	
Region2	Line3	Produ...	2.20	
Region1	Line1	Product3	2.29	
Region3	Line1	*	2.51	
Region3	Line2	Product6	3.06	
Region2	Line2	*	3.11	
Region1	Line1	Product1	3.62	
Region2	Line1	*	4.06	
Region1	*	*	4.57	
Region1	Line1	*	4.57	
Region3	Line1	Product3	5.02	
Region2	Line2	Product6	5.37	
Region2	Line3	Product8	6.61	
Region2	Line1	Product1	7.75	
Region2	Line2	Product5	9.39	
Region3	Line2	*	9.85	
Region2	Line3	Product9	10.64	
Region3	Line2	Product5	11.94	
Region1	Line1	Product2	17.59	
Region3	Line1	Product4	19.00	

Figure 4.2. Table Tab

Filtering the Contents of the Table Tab

You can filter the contents of the table by using the **Show** drop-down list. The following options are available:

- **No filter** - lists all of the series. No filter is applied. This is the default when you generate a forecast.
- **Notes** - lists only the series that have notes. For more information, see "Add a Note to a Series" on page 132.
- **Overrides** - lists only the series for which you specified an override. For more information, see "Adding Overrides" on page 130.
- **All Exceptions** - lists only the series that are identified as exceptions. You identify exceptions when you create a filter. If you have identified more than one filter as an exception, then any series that are identified as exceptions appear in this list. For more information, see Chapter 6, "Working with Filters."
- a filter for each level of the hierarchy - lists the series for that level of the hierarchy. This option is available only if you forecasted your data hierarchically.
- any user-defined filters.

After selecting a filter from the **Show** drop-down list, you can refresh the contents in the table by clicking **Refresh**.

Customizing the Table Tab

To customize the **Table** tab, complete the following steps:

1. Select **View** → **Customize View**. The Customize View dialog box opens.
2. Click the **Table** tab.
3. Select the statistics of fit to display in the series table and click  . From this list, you can also choose to display the Reconciled *statistic of fit* for the project.
4. Select the series information to display in the series table and click  .

   If you choose the following series information, these items are marked with special icons in the **Table** tab:

   - **Notes** – They are identified in the series table by the  icon.
   - **Overrides** – They are identified in the series table by the  icon.
   - **Exceptions** – They are identified in the series table by ✕. The exceptions are displayed by default.

   Note: You can click **Reset to Defaults** to return to the default values in the series table.
5. Click **OK** to close the Customize View dialog box and to apply your changes.

About the Forecasting View

What Is the Forecasting View?

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

In the Forecasting View, you can view the forecast for the selected series on the **Hierarchy** tab or on the **Table** tab. The Forecasting View has three sections: a forecast plot, the data table, and a Notes section.

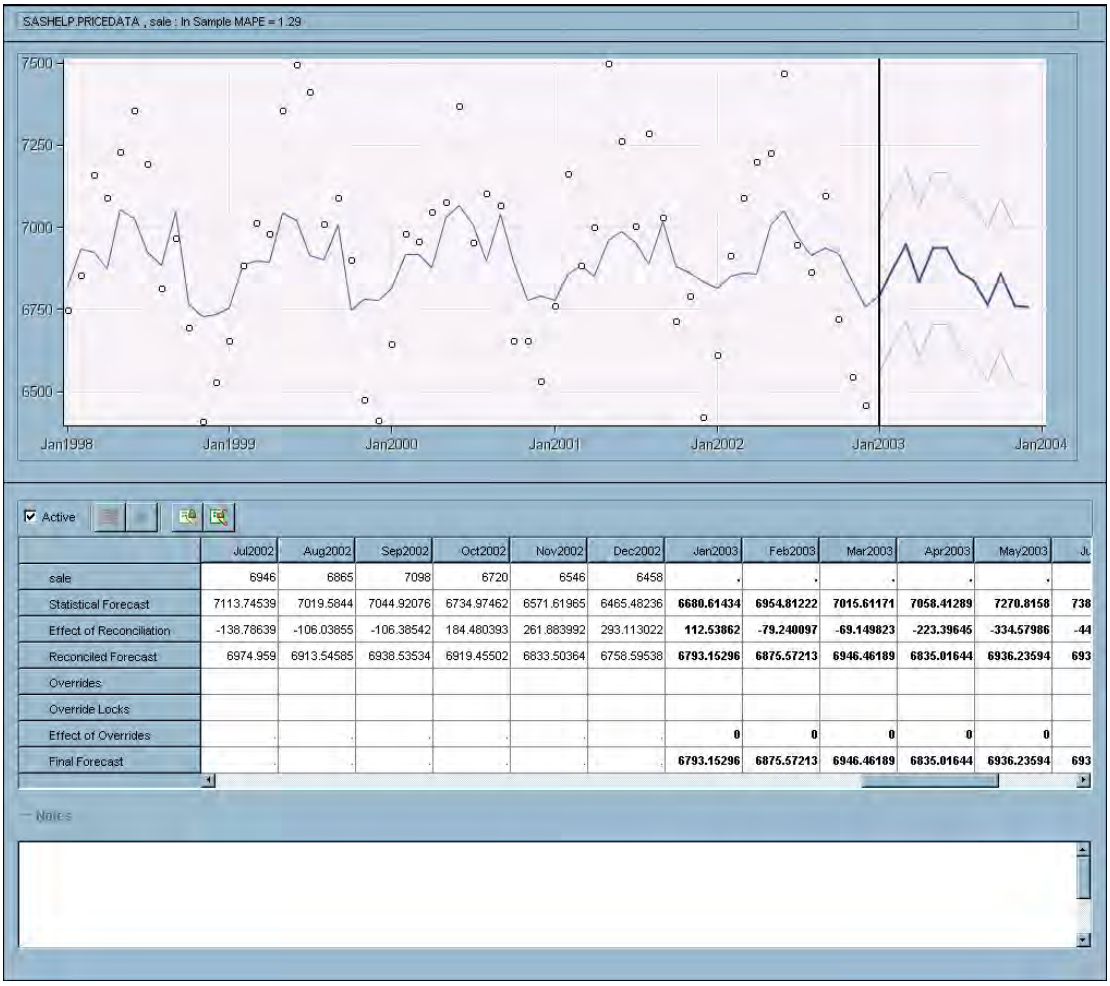

	Jul2002	Aug2002	Sep2002	Oct2002	Nov2002	Dec2002	Jan2003	Feb2003	Mar2003	Apr2003	May2003	Ju
sale	6946	6865	7098	6720	6546	6458	.	.	.	.	.	
Statistical Forecast	7113.74539	7019.5844	7044.92076	6734.97462	6571.61965	6465.48236	**6680.61434**	**6954.81222**	**7015.61171**	**7058.41289**	**7270.8158**	**738**
Effect of Reconciliation	-138.78639	-106.03855	-106.38542	184.480393	261.883992	293.113022	**112.53862**	**-79.240097**	**-69.149823**	**-223.39645**	**-334.57986**	**-44**
Reconciled Forecast	6974.959	6913.54585	6938.53534	6919.45502	6833.50364	6758.59538	**6793.15296**	**6875.57213**	**6946.46189**	**6835.01644**	**6936.23594**	**693**
Overrides												
Override Locks												
Effect of Overrides	.	.	.	.	.	.	**0**	**0**	**0**	**0**	**0**	
Final Forecast	.	.	.	.	.	.	**6793.15296**	**6875.57213**	**6946.46189**	**6835.01644**	**6936.23594**	**693**

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 values:

- historical data
- reconciled statistical forecasts
- reconciled statistical forecast confidence intervals
- any overrides that you have specified
- final forecasts

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 values:

- historical data
- statistical forecasts
- statistical forecast confidence intervals
- any overrides that you have specified
- final forecasts

To customize the forecast plot, see "Customizing the Forecasting View" on page 35.

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.

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 123.

By default, the data table displays the following information:

Row	Description
the name of the dependent variable	the historical data that was in the input data set for the dependent variable.
Statistical Forecast	the values that SAS Forecast Studio forecasted for the dependent variable
Effect of Reconciliation	the value that was added to or subtracted from the statistical forecast to determine the reconciled forecast.
Reconciled Forecast	the total of the statistical forecast plus the effect of the reconciliation
Overrides	the override values that you specified for any time period
Override Locks	displays whether you have specified the override as a locked override. For more information, see "Lock and Unlock Overrides" on page 132.
Effect of Overrides	the increase or decrease to the reconciled forecast to get the final forecast.
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 135. • 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.

Note: The Effect of Reconciliation, the Reconciled Forecast, Override Locks, and the Effect of Overrides rows only appear in the 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 58.

Optionally, you can choose to include the Aggregate event and indicator variable lifts row in the data table. This row shows the aggregate effect that the events that you specified have on the reconciliation of the statistical forecast. You can also choose to show the confidence intervals for the statistical forecast and the reconciled statistical forecast. To customize the content of the data table, see "Customizing the Forecasting View" on page 35.

Historical values are displayed in regular font. Future values are displayed in bold. Cells that you can edit have a white background. For more information, see "Adding Overrides" on page 130.

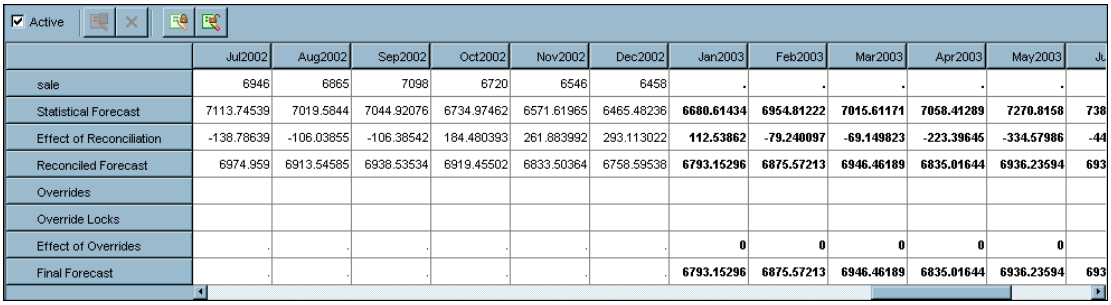

Active

	Jul2002	Aug2002	Sep2002	Oct2002	Nov2002	Dec2002	Jan2003	Feb2003	Mar2003	Apr2003	May2003	Ju
sale	6946	6865	7098	6720	6546	6458	.	.	.	.	.	
Statistical Forecast	7113.74539	7019.5844	7044.92076	6734.97462	6571.61965	6465.48236	**6680.61434**	**6954.81222**	**7015.61171**	**7058.41289**	**7270.8158**	**738**
Effect of Reconciliation	-138.78639	-106.03855	-106.38542	184.480393	261.883992	293.113022	**112.53862**	**-79.240097**	**-69.149823**	**-223.39645**	**-334.57986**	**-44**
Reconciled Forecast	6974.959	6913.54585	6938.53534	6919.45502	6833.50364	6758.59538	**6793.15296**	**6875.57213**	**6946.46189**	**6835.01644**	**6936.23594**	**693**
Overrides												
Override Locks												
Effect of Overrides	.	.	.	.	.	.	**0**	**0**	**0**	**0**	**0**	
Final Forecast	.	.	.	.	.	.	**6793.15296**	**6875.57213**	**6946.46189**	**6835.01644**	**6936.23594**	**693**

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. For more information, see "Add a Note to a Series" on page 132.

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. For more information, see Chapter 9, "Working with Overrides."
- Add notes. For more information, see "Add a Note to a Series" on page 132.

Customizing the Forecasting View

To customize the Forecasting View, complete the following steps:

1. Select **View** → **Customize View**. The Customize View dialog box opens.
2. Click the **Forecasting View** tab.
3. Select the information to show on the forecasting plot and in the data table. You can choose from the following options:

- **Historical Data**
- **Statistical Forecast**
- **Statistical Forecast Confidence Intervals**
- **Reconciled Statistical Forecast**
- **Reconciled Statistical Forecast Confidence Intervals**
- **Manual Overrides**
- **Final Forecast**
- **Aggregate event lifts** (for the plot) or **Aggregate event and indicator variable lifts** (for the table)

For more information about these options, see "About the Data Table" on page 33.

Note: Click **Reset to Defaults** to return to the default values in the forecasting view.

4. Click **OK** to close the Customize View 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 project. You can also use plots and tables to compare how different models fit the data.

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

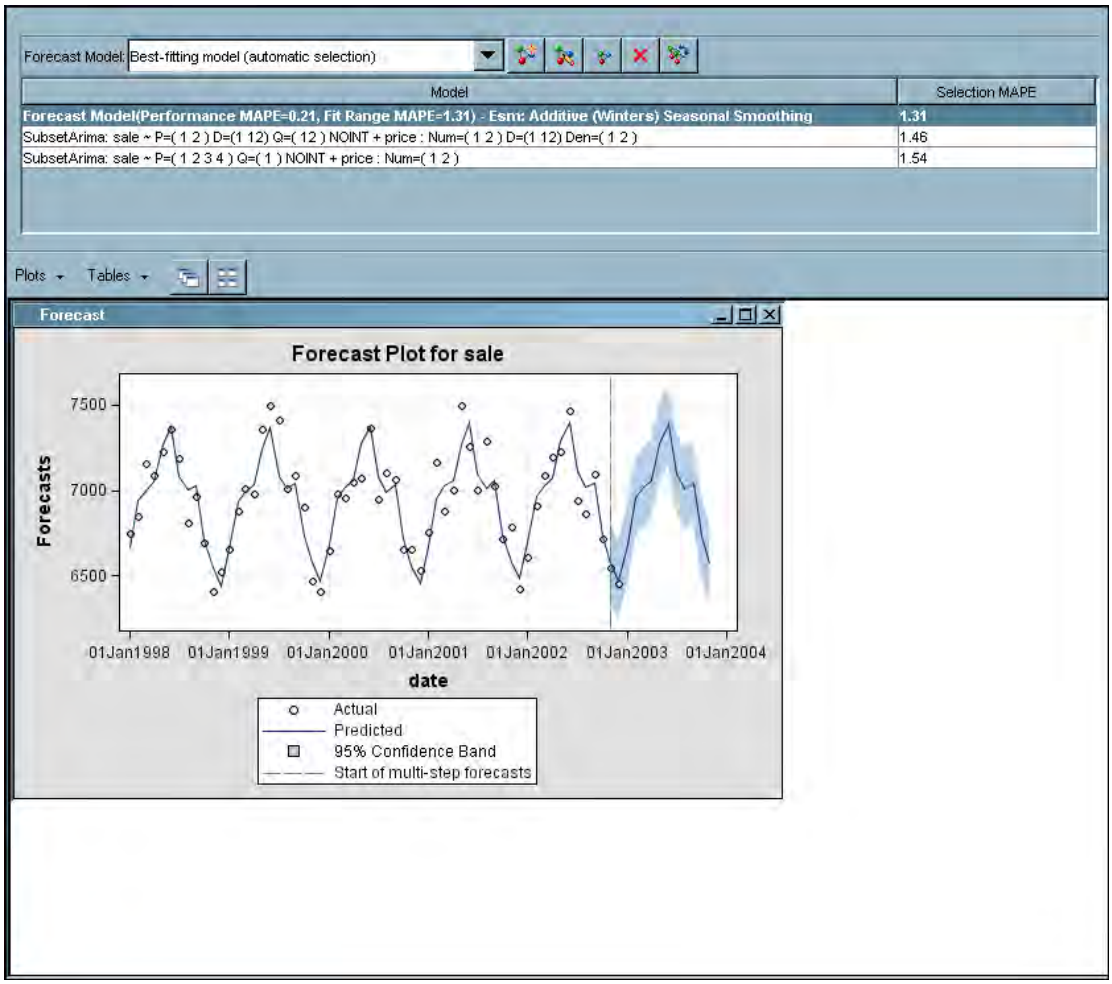

Figure 4.5. Model View

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 and the fit criterion for the model.

The name of the forecast model is in bold and has **Forecast Model** prepended to the name.

For more information, see "Using the Model Selection List" on page 222.

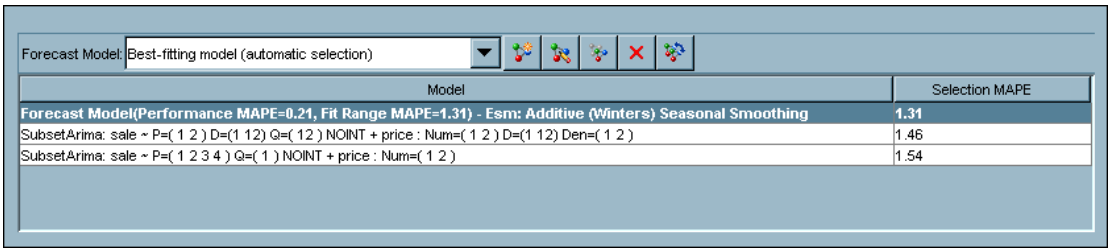

Figure 4.6. Model Selection List in Model View

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 224 and "View the Model Tables" on page 225.

Note: If you select another series, then some of the plots or tables that you created might close. These plots or tables close because they are not available for the series that you selected. If you have selected a series that contains missing values, then you cannot create some plots or tables in the Model View.

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:

- View model plots. For more information, see "View the Model Plots" on page 224.
- View model tables. For more information, see "View the Model Tables" on page 225.
- Select a model to forecast. For more information, see "Select the Forecast Model" on page 226.

- Edit the selected model. For more information, see "Edit a Model" on page 228.
- Create a new model. For more information, see "Create an Entirely New Model" on page 175.
- Copy an existing model. For more information, see "Working with the Model Selection List" on page 223.
- Compare models. For more information, see "Compare Models" on page 230.
- Remove models from the model selection list. For more information, see "Working with the Model Selection List" on page 223.

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 on the **Hierarchy** tab or on the **Table** tab. 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 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: If you select a different series or change the value of one of the time series options, then some of the plots or tables that you created might close. These plots or tables close because they are not available for the series that you selected or the options that you specified.

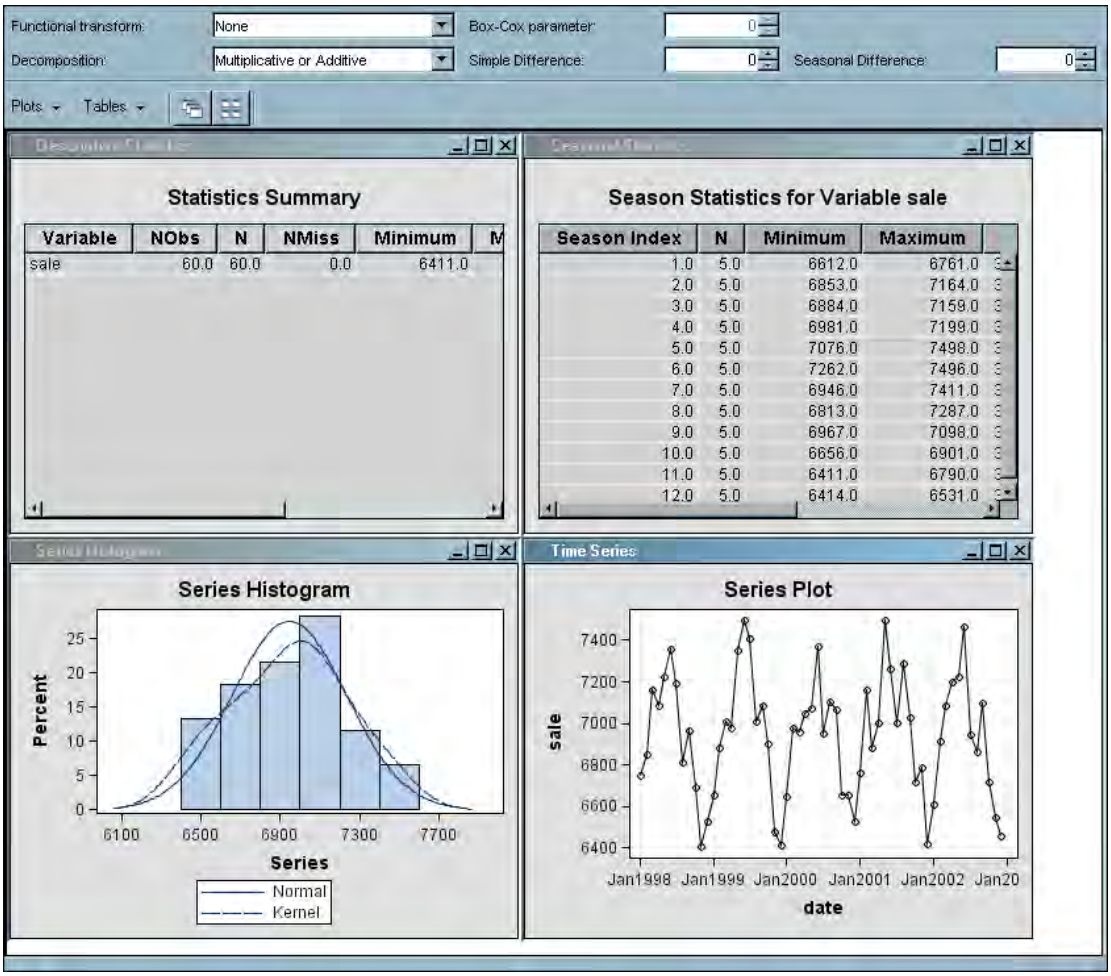

Figure 4.7. Series Analysis View with the Plots Tiled in the Window

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 Select Series window, you choose the second variable from a list of independent and reporting variables.
    - **Cross-series Plot** - displays a two-variable plot of the dependent variable and a second variable. From the Select Series window, you choose the second variable from a list of independent and reporting variables.
    - **Cycle Component** - displays a plot of the cycle component 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 Transform** - specifies the transform for the time series. The available options depend on the open plots.
- **Decomposition** - specifies the decomposition for the time series.
- **Box-Cox parameter** - specifies the Box-Cox 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
• 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 the icon to view the SAS log. • 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.	• Your user ID. • The name of the SAS Forecast Server Mid-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

Chapter Contents

Chapter 5
Working with Projects

Overview of Projects

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
!SASROOT\SASForecastStudio\Projects*project-name*

UNIX operating environments
!SASROOT/SASForecastStudio/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*.

Output Data Sets in Project Directory

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
!SASROOT\SASForecastStudio\Projects*project-name***hierarchy**

UNIX operating environments
!SASROOT/SASForecastStudio/Projects/*project-name*/**hierarchy**

From the project directory, you can view the following output data sets:

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 Forecast Options for the Project" on page 109.

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
- _COMP_ - the name of the component
- _TIMEID_ - the time ID variable
- _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_ - the name of the variable to be forecast
- _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.
- _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
- _TIMEID_ - the time ID values
- 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_ - the name of the dependent variable
- _TIMEID_ - time ID values
- 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.

OUTSEASON= data set

The OUTSEASON= data set contains the variables that you assign as BY variables and the variables listed below. The OUTSEASON= data set records the seasonal statistics for each dependent variable.

The following variables that are related to seasonal statistics are based on the values that you specify for the time interval and seasonality options:

- _NAME_ - variable name
- _TIMEID_ - time ID values
- _SEASON_ - seasonal index
- NOBS - number of observations
- N - number of nonmissing observations

- NMISS - number of missing observations
- MIN - minimum value
- MAX - maximum value
- RANGE - range value
- SUM - summation value
- MEAN - mean value
- STDDEV - standard deviation
- CSS - corrected sum of squares
- USS - uncorrected sum of squares
- MEDIAN - median value

These statistics are computed for each season.

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
- 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
- RSQUARE - R-square
- ADJRSQ - adjusted R-square
- AADJRSQ - Amemiya's adjusted R-square
- RWRSQ - random walk R-square
- AIC - Akaike information criterion
- SBC - Schwarz Bayesian information criterion
- APC - Amemiya's prediction criterion
- MAXERR - maximum error
- MINERR - minimum error
- MINPE - maximum percent error
- MAXPE - minimum percent error
- ME - mean error
- MPE - mean percent error

If the statistics-of-fit step fails for particular variable, no observations are recorded.

OUTSTATSELECT= data set

The OUTSTATSELECT= 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:

- _MODEL_ - the name of the forecasting model
- _SELECT_ - the name of the model selection list
- _NAME_ - variable name
- _REGION_ - statistics-of-fit region
- DFE - degrees of freedom error
- N - number of observations
- NOBS - number of observations used
- NMISSA - number of actual values that are missing
- NMISSP - number of predicted values that are missing
- NPARMS - number of parameters
- TSS - total sum of squares (uncorrected)
- SST - total sum of squares (corrected)
- 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
- RSQUARE - R-square
- ADJRSQ - adjusted R-square
- AADJRSQ - Amemiya's adjusted R-square
- RWRSQ - random walk R-square
- AIC - Akaike information criterion
- SBC - Schwarz Bayesian information criterion
- APC - Amemiya's prediction criterion
- MAXERR - maximum error
- MINERR - minimum error
- MINPE - maximum percent error
- MAXPE - minimum percent error
- ME - mean error
- MPE - mean percent error

If the statistics-of-fit step fails for a particular variable, no observations are recorded. If you specify 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.

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
- _STATUS_ - forecasting status. Nonzero values imply that no forecast was generated for the series.
- NOBS - number of observations
- N - number of nonmissing 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.

OUTTREND= data set

The OUTTREND= data set contains the variables that you assign as BY variables and the variables listed below. The OUTTREND= data set records the trend statistics for each dependent variable.

The following variables that are related to trend statistics are based on the values that you specified for the time interval and seasonality options:

- _NAME_ - variable name
- _TIMEID_ - time ID values
- _SEASON_ - seasonal index
- NOBS - number of observations
- N - number of nonmissing observations
- NMISS - number of missing observations
- MIN - minimum value
- MAX - maximum value
- RANGE - range value
- SUM - summation value
- MEAN - mean value
- STDDEV - standard deviation
- CSS - corrected sum of squares
- USS - uncorrected sum of squares
- MEDIAN - median value

These statistics are computed for each time period.

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 are related to the forecasting step:

- _NAME_ - variable name
- _TIMEID_ - time ID values
- 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 sign 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 can also contain the variables listed below. The following variables contain observations related to the statistics-of-fit step:

- _NAME_ - variable name
- _REGION_ - statistics-of-fit region
- DFE - degrees of freedom error
- N - number of observations
- NOBS - number of observations used
- NMISSA - number of actual values that are missing
- NMISSP - number of predicted values that are missing
- NPARMS - number of parameters
- 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
- RSQUARE - R-square
- ADJRSQ - adjusted R-square
- AADJRSQ - Amemiya's adjusted R-square
- RWRSQ - random walk R-square

- AIC - Akaike information criterion
- SBC - Schwarz Bayesian information criterion
- APC - Amemiya's prediction criterion
- MAXERR - maximum error
- MINERR - minimum error
- MINPE - maximum percent error
- MAXPE - minimum percent error
- ME - mean error
- MPE - mean percent error

If the statistics-of-fit step fails for 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
- _STATUS_ - forecasting status. Nonzero values imply that no forecast was generated for the series.
- NOBS - number of observations
- N - number of nonmissing 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.

Managing Your Projects

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

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

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

- To create a new project, click **New**. The New Project Wizard opens. For more information, see "Overview of the New Project Wizard" on page 58.
- To open an existing project, select the project from the list and click **Open**. You can have only one project open at a time.

  Note: Dimmed projects can only be opened 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 window. For more information, see "View Project Properties" on page 73.

- To 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**.

  You can only delete a project that you created, and when you delete a project, any archived projects that you have created are also deleted. For more information, see "Archiving Projects" on page 71.

  Note: You should not delete the projects from the project directory on the SAS Workspace Server. For more information about the project directory, see "What Is a Project?" on page 47.

Opening a SAS Forecast Studio 1.2 Project

Changes to Projects in SAS Forecast Studio 1.4

You can open a SAS Forecast Studio 1.2 project in SAS Forecast Studio 1.4. When you open a 1.2 project in SAS Forecast Studio 1.4, you will see the following changes:

- The names of the default models that are associated with the project change. For example, a model named DIAG1 will change to a model called ARIMA. For more information about default models, see "Overview of Project Models" on page 219.
- For projects that were created in SAS Forecast Studio 1.2, the sensitivity of the seasonality test is set to 0.01 by default. If you open a 1.2 project in SAS Forecast Studio 1.4, then this sensitivity remains the same because SAS Forecast Studio tries to honor a project's settings when the application is upgraded.

  However, retaining this sensitivity value could cause some of the ARIMA or Subset ARIMA models to fail. This failure is the result of stricter seasonality tests in SAS Forecast Studio 1.4. Series that were previously identified as seasonal might now be identified as nonseasonal, and as a result, any seasonal models that were fit to these series will fail. To prevent these models from failing, you should clear the **Perform seasonality test** check box, so that no seasonality test is performed. For more information, see "Specify the Model Selection Options" on page 120.

  Note: If you create the same project in SAS Forecast Studio 1.4, then no seasonality test is performed by default, so the ARIMA and Subset ARIMA models will not fail.
- If you open a 1.2 project that contains overrides in SAS Forecast Studio 1.4, then the following occurs:
  - All overrides from the 1.2 project are automatically locked.
  - The settings that were used to reconcile the overrides in the 1.2 project are the same as the default settings of the advanced reconciliation options in SAS Forecast Studio 1.4. These advanced reconciliation settings are **No restriction** and **Require that conflicts be resolved before reconciliation**.

– The overrides from the 1.2 project are not immediately reconciled. The "Reconciliation is out of date" message appears at the top of the Forecasting View. The overrides in the 1.2 project need to be reconciled using the new SAS Forecast Studio 1.4 functionality. You must click **Update** to reconcile the overrides. However before you click **Update**, you can change the status of the overrides from the 1.2 project from locked to unlocked or change the settings for the advanced reconciliation options.

Note: After reconciliation, your final forecasts might not be the exact same values as the final forecasts that you had in 1.2. The final forecasts are changed because of improvements to the override reconciliation functionality.

For more information about overrides and reconciliation in SAS Forecast Studio 1.4, see Chapter 9, "Working with Overrides."

Updating the Project

When you open a SAS Forecast Studio 1.2 project in SAS Forecast Studio 1.4, the Update Project dialog box opens.

You can choose how to update the models in the project by using the following options:

- **Recreate automatically-generated models, refit all models, and select a forecast model**.
- **Refit all models, and select a forecast model** - This is the default.

If you select the **Use updated data if available** check box, SAS Forecast Studio checks to see if the data source has been updated, and if so, updates the project data and reforecasts the existing series based on the option that you selected.

After you click **OK**, SAS Forecast Studio automatically converts any 1.2 projects to be compatible with 1.4. Any overrides that you have specified are stored in the new 1.4 format. If you try to open this project again in SAS Forecast Studio 1.2, 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, select **Create a new project** in the Welcome to SAS Forecast Studio dialog box.
- To create a new project when SAS Forecast Studio is already running, select **File** → **New Project** or click  in the toolbar.

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

1. 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. For more information about SAS naming conventions, see *SAS Language Reference: Concepts*.
2. Provide a brief description of the project. A description helps other users better understand the function of your project. This description is optional.
3. Select the **Allow other users to view and edit this project** check box to share this project. This option is not selected by default. When this check box is not selected, only the person who created 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.

   Note: 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 "Required Format" on page 18.

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 66.

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 67.

- **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, make sure that the BY variables are specified in the same order that the data set is sorted.

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

- **Independent Variables** - specifies 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 66.

2. Click **Add Events** to add an event to your project. For more information, see "Managing Events" on page 96.
3. Click **Adjustments** to add an adjustment variable. For more information, see "Managing Adjustment Variables" on page 68.
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 67.

   Note: If you assign a variable to a role and the variable name is reserved, then an error message will appear. For more information about variable names, see "Reserved Variable Names" on page 241.

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. Select whether or not to forecast a hierarchy. 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, the variable that you assign to the Dependent Variable role is sorted into individual series in a hierarchy using the BY variables that you assign. If you choose to create a forecast hierarchy, then only one dependent variable can be forecast.

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. The Hierarchy Preview window opens. For more information, see "Previewing the Hierarchy" on page 71.

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 most variable types, you can choose from the following aggregation statistics:

   - **Maximum of values**
   - **Sum of values**, which is the default
   - **Number of non-missing values**

- **Median of values**
- **Corrected sum of squares of values**
- **Last value**
- **Total number of values**
- **First value**
- **Uncorrected sum of squares of values**
- **Standard deviation of values**
- **Number of missing values**
- **Average of values**
- **Minimum of values**

For more information about these statistics, see "Understanding the Accumulation and Aggregation Methods" on page 112.

Note: If you are aggregating data for the dependent variable, then only the **Sum of values** and **Average of values** options are available.

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

4. Select the **Reconcile hierarchy** check box to reconcile the hierarchy.

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

5. 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 generate reconciled 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 generate the reconciled forecasts for the higher levels in the hierarchy.

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

6. 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 70.

7. Click **Advanced** to specify the advanced reconciliation settings. For more information, see "Specify the Advanced Reconciliation Options" on page 136.
8. 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 Options dialog box. For more information about these options, see "Working with the Forecast Options for the Project" on page 109.
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 Summary window after SAS Forecast Studio is finished. If any errors occur during the run, these errors appear in the Summary window. For more information, see "Reviewing a Summary of the Forecast Run" on page 65.

Reviewing a Summary of the Forecast Run

Once SAS Forecast Studio has finished generating the forecasts for the project, the Forecast Summary dialog box appears. This dialog box contains the following information:

- the name of the last task that was run
- the number of series processed
- the number of series that completed successfully
- the number of series that completed unsuccessfully
- the number of errors

Note: The number of series that complete unsuccessfully and the number of errors appear only if this number is greater than zero. Click **View Log** to view the SAS log. For more information, see "View the SAS Log" on page 75.

Click **Close** to close the window and view the forecast results.

Open a Project

You can open a project in the following ways:

- When you first invoke SAS Forecast Studio, you can open a project from the Welcome to SAS Forecast Studio window.
- After SAS Forecast Studio is running, you can manage your projects and open existing projects from the Projects dialog box. For more information, see "Managing Your Projects" on page 56.

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

The Welcome to SAS Forecast Studio window opens by default when you start SAS Forecast Studio. From this window, you can create a new project or open an existing project. The existing projects are sorted alphabetically. However, you can choose these projects by date.

You can choose from the following options:

- **Open an existing project** - enables you to select a project from the list. Dimmed projects can be opened only by the person who created the project. To share this project, select the **Allow other users to view and edit this project** check box in the Properties window. For more information, see "View Project Properties" on page 73.

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

- **Create a new project** - opens the New Project Wizard. For more information, see "Overview of the New Project Wizard" on page 58.

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

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:

	numeric variable
	character variable
	date/time variable

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

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

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 60.

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 66.
3. Click **OK**.

Note: When you add an independent variable, SAS Forecast Studio performs a non-destructive 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 "Understanding Time Series Data" on page 17.

Common errors include the following:

- SAS Forecast Studio is unable to determine the time interval for the data. If this happens, the Specify Time Interval dialog box opens. From this dialog box, you must select from one of the following options:
  - **Select another variable as a time ID** - enables you to return to the Assign Data to Roles step in the New Project Wizard and select a different variable to use as the time ID.
  - **Manually set the time interval and other time ID properties** - enables you to create the time ID. The options that are available from this dialog box are the same as the time ID options in the Options dialog box. For more information about the options that you need to set, see "Specifying Options for the Time ID" on page 115.
- You must specify a time interval when the input data is transactional. This is when the input data is timestamped at no particular frequency.
- If the time interval is too small, then a large number of the observations will have missing values. For example, you select WEEK as the interval, but there is not enough data for most of the weeks. A larger interval, such as MONTH, should be selected.
- Assigning variables to incorrect roles is a common source of errors. **Note:** Incorrectly assigning variables to the BY variables role does not always generate errors. For example, if you forget to assign a BY variable, then that variable is not included in the hierarchy and the data is not aggregated over that variable.

Working with Adjustment Variables

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

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 68.

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 **Add**. The New Adjustment Variable dialog box opens.
3. Select the adjustment variable and move it to the **Selected** box. The name of this variable will appear in the **Adjustment Variables** list in the Adjustments dialog box. The variable assigned as the adjustment variable cannot have any other role in the project.
4. Specify the operations to perform before and after forecasting. You can choose from add, subtract, multiply, divide, minimum, and maximum. By default, no operation is selected.
5. Select the dependent variable to adjust and move it to the **Selected** box. You must select one variable to adjust.
6. Click **OK** to return to the Adjustments dialog box. For more information about this dialog box, see "Managing Adjustment Variables" on page 68.

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 68.

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 68.

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 68.

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 **Move Up** or **Move Down**. 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 68.

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 62.

   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 recreated.
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

Previewing the Hierarchy

After you have specified the order of the levels in the hierarchy in the New Project Wizard or Hierarchy Settings dialog box, you can see how the hierarchy will be arranged. (The levels are the variables that you assign to the BY variables role.) The Hierarchy Preview window enables you to view the levels in the hierarchy and the order in which they will appear.

Note: The Hierarchy Preview window displays only the nodes in the first four levels of the hierarchy. The remaining levels are represented by dotted lines.

This preview is of the organizational hierarchy; it is not the actual hierarchy of the data. To change the order of the hierarchy, see "Configure the Hierarchy" on page 62.

Archiving Projects

Managing Archived Projects

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

To view more information about an archived project, select the project name from the list. The name, description, and date that the archive was created appears in the Archive Information box.

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

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

Create an Archived Project

To create an archived project, complete the following steps:

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

SAS Forecast Studio saves the archived project on the SAS Workspace Server. An Archives subdirectory is created for your project. For example, if you archived a project named "MyProject," then your archived project is saved in the following locations:

Windows operating environments

!SASROOT\SASForecastStudio\Archives\MyProject\Archive.far

UNIX operating environments

!SASROOT/SASForecastStudio/Archives/MyProject/Archive.far

Restore an Archived Project

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

1. Select **File** → **Archives** to open the Archives dialog box.
2. Select the project name from the list and click **Restore**. You are prompted to confirm this restoration. Click **OK** to restore the selected project.

Note: Any changes that you make to the current project are lost when you restore an archive. Consider creating an archive of the current project before restoring an archive.

Delete an Archived Project

To delete an archived project, complete the following steps:

1. Select **File** → **Archives** to open the Archives dialog box.
2. Select the project from the list in the Archives dialog box and click **Delete**. You are prompted to confirm this deletion. Click **OK** to delete the selected project.

Update the Project Data

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

To update the project data, complete the following steps:

1. Select **File** → **Update Project Data**. The Update Project Data dialog box opens.
2. Select the action that you want to perform. You can choose from the following options:

   - **Recreate automatically-generated models, refit all models, and select a forecast model**.
   - **Refit all models, and select a forecast model** - This is the default.

- **Refresh the current forecast model, updating the parameter values**.
- **Refit the current forecast model, using the same parameter values**.
- **Do not update the project data**.

3. 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 **Recreate automatically-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 "Best Smoothing Model" on page 253.

Viewing Properties

View Project Properties

To view and edit the project properties, complete the following steps:

1. Select **File** → **Project Properties**. The Project Properties window opens.

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

2. To add a project description, type the description in the **Description** box.
3. Select the **Allow other users to view and edit this project** check box to share this project. This option is not selected by default. When this check box is not selected, only the person who created the project can open it.
4. Click **OK** to save your changes and to close the Project Properties window.

View Series Properties

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

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

You can see the following information for the selected series:

- the name of the series
- start date
- end date

- number of values
- number of missing values
- mean
- standard deviation
- minimum value
- maximum value

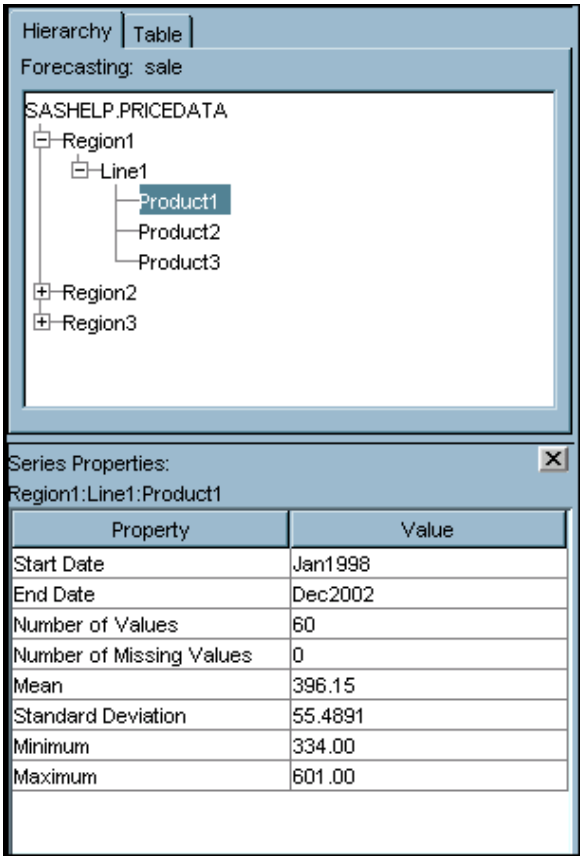

Figure 5.1. Hierarchy Tab with Series Properties Pane

Working with the SAS Log

What Is 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**.

Note: You can also view the SAS log from the Summary window that opens after SAS Forecast Studio has produced the forecasts. For more information about this window, see "Reviewing a Summary of the Forecast Run" on page 65.

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 75.

Save the SAS Log

To save the SAS log, complete the following steps:

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

Working with the SAS Code

View 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 75.

From the SAS Code window, 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 76.

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 window opens.
2. Select the version of the code that you want to view from the **Task** drop-down list. The following options are available:
   - **Recreate Project** - displays the code that SAS Forecast Studio uses to recreate 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 recreate 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 **Recreate Project** from the **Task** drop-down list.

The version of the code that you select appears in the dialog box. You can expand and collapse sections of code by using the plus and minus signs next to the line numbers.

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 window opens.
2. Select the version of the project code that you want to view. For more information, see "View the SAS Code" on page 75.
3. In the SAS Code dialog box, click **Save As**. The Save As window opens.
4. Select the location where you want to save this file and click **Save**.

Running Projects in Batch Mode

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.

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 recreate the project. This code includes all of the data preparation and forecasting steps.

  These files include the following project code:

  - Library declarations that create libnames, 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 recreates the model repositories.
  - 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 libnames, 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 recreates the model repositories.
- 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 libnames, 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.
- 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 libnames, 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 libnames, 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 libnames, 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 recreates 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.

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

Chapter Contents

Chapter 6
Working with Filters

What Is a Filter?

Filters enable you to specify criteria to subset your data. In SAS Forecast Studio, you create filters to help manage your forecasts. You can specify whether to use a filter to identify exceptions. These exceptions help you identify forecasts that do not meet a specific criterion. After the filter is applied, series that are identified as exceptions appear in red in the **Hierarchy** tab.

For an example of how to use a filter to identify exceptions, see "Examples of Filters" on page 85.

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**.

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

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

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

Create a Filter

To create a filter, complete the following steps:

1. Select **Tools** → **Filters**. The Filters dialog box opens.
2. Click **New**. The New Filter 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 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.
5. To use this filter to identify exceptions, select the **This filter will be used to identify exceptions** check box. This option is selected by default.
6. From the **Series Property** drop-down list, select the series property to evaluate. You can choose from the following:

   - **None** (which is the default)
   - **Level**
   - all of the available levels in the hierarchy
   - all of the fit criteria
   - the series properties

   Note: The level options are not available if you did not forecast your data hierarchically.
7. From the **Condition** drop-down list, select the conditional operator to use to evaluate the series property. The conditional operators that are available depend on the value that you have selected from the **Series Property** drop-down list:

   - If you have selected **Level** or a level in the hierarchy from the **Series Property** drop-down list, then only the **Equal to** option is available.
   - If you have selected any of the fit criteria or any of the series properties from the **Condition** drop-down list, then the available options include **Equal to**, **Greater than** (which is the default), and **Less than**.

   Note: The **Condition** drop-down list is not available if you select **None** as the series property.
8. From the **Value** drop-down list or text box, select or type the value to evaluate against the series property. How this option is displayed depends on what you have selected in the **Series Property** drop-down list.

   - If you have selected a series property (such as a level) that is a fixed value, then this option appears as a drop-down list.
   - If you have selected a series property (such as a fit criterion) that is a continuous value, then this option appears as a text box.

   Note: The **Value** option is not available if you select **None** as the series property.
9. Click **Add Row** if you need to add additional filter expressions.

   To remove a row, click  .
10. 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 85.

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 Edit *Filter-name* dialog box opens. The options in this dialog box are the same as those in the New Filter dialog box. For more information about the available options, see "Create a Filter" on page 83.
3. 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**.

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 New Filter 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. Select the **This filter will be used to identify exceptions** check box.
6. From the **Series Property** drop-down list, select **Mean Absolute Percent Error**.
7. From the **Condition** drop-down list, select **Greater than**.
8. In the **Value** text box, type **`7`**.

The final window should look like the following:

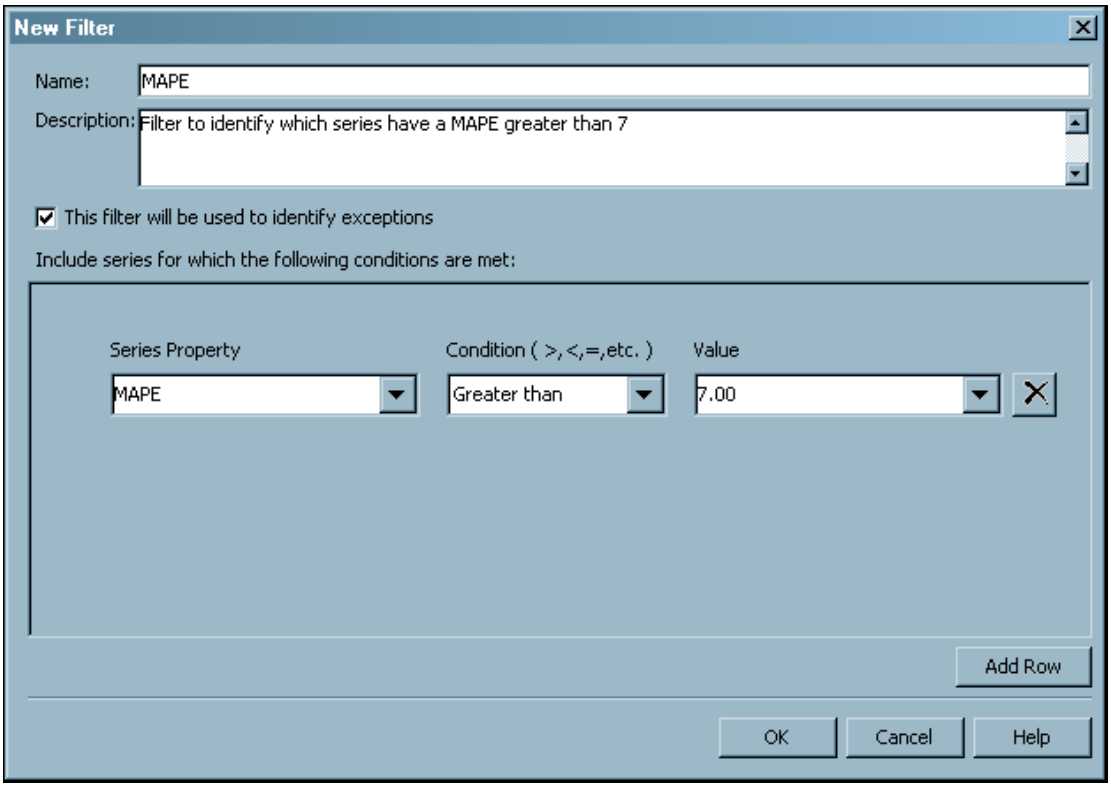

Figure 6.1. Completed New Filter Dialog Box for Filter Example

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

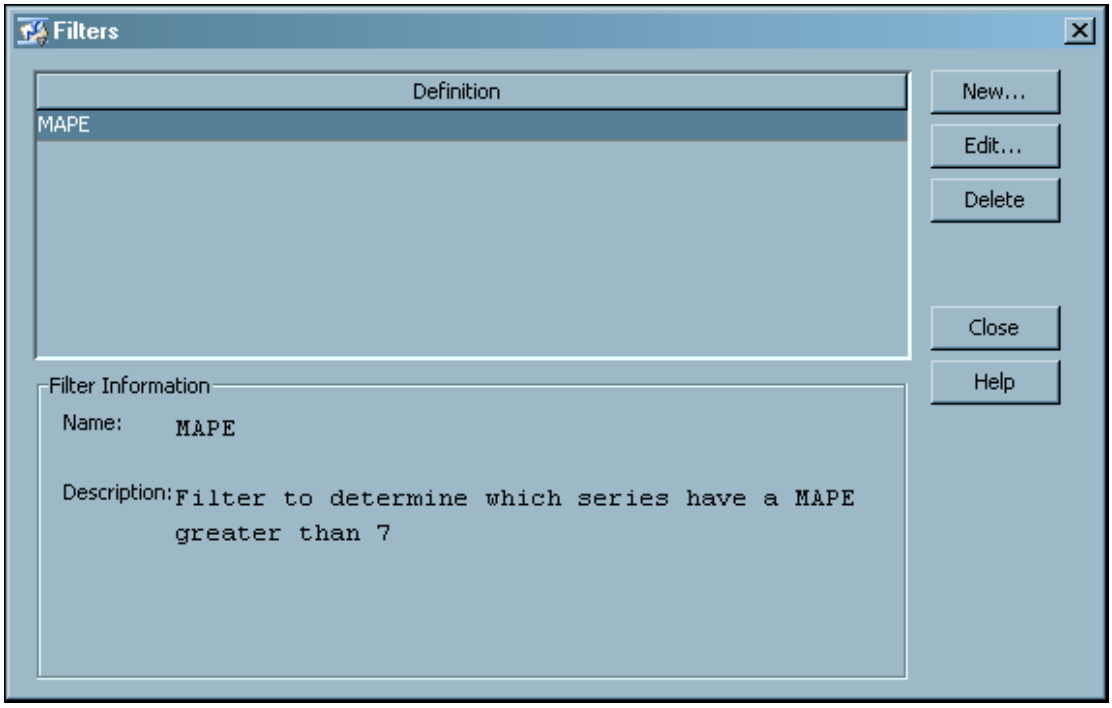

Figure 6.2. New Filter in Filters Dialog Box

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

Because you selected the **This filter will be used to identify exceptions** check box, any series that has a MAPE greater than 7 appears in bold red in the **Hierarchy** tab and has a red ✕ in the Exceptions column on the **Table** tab.

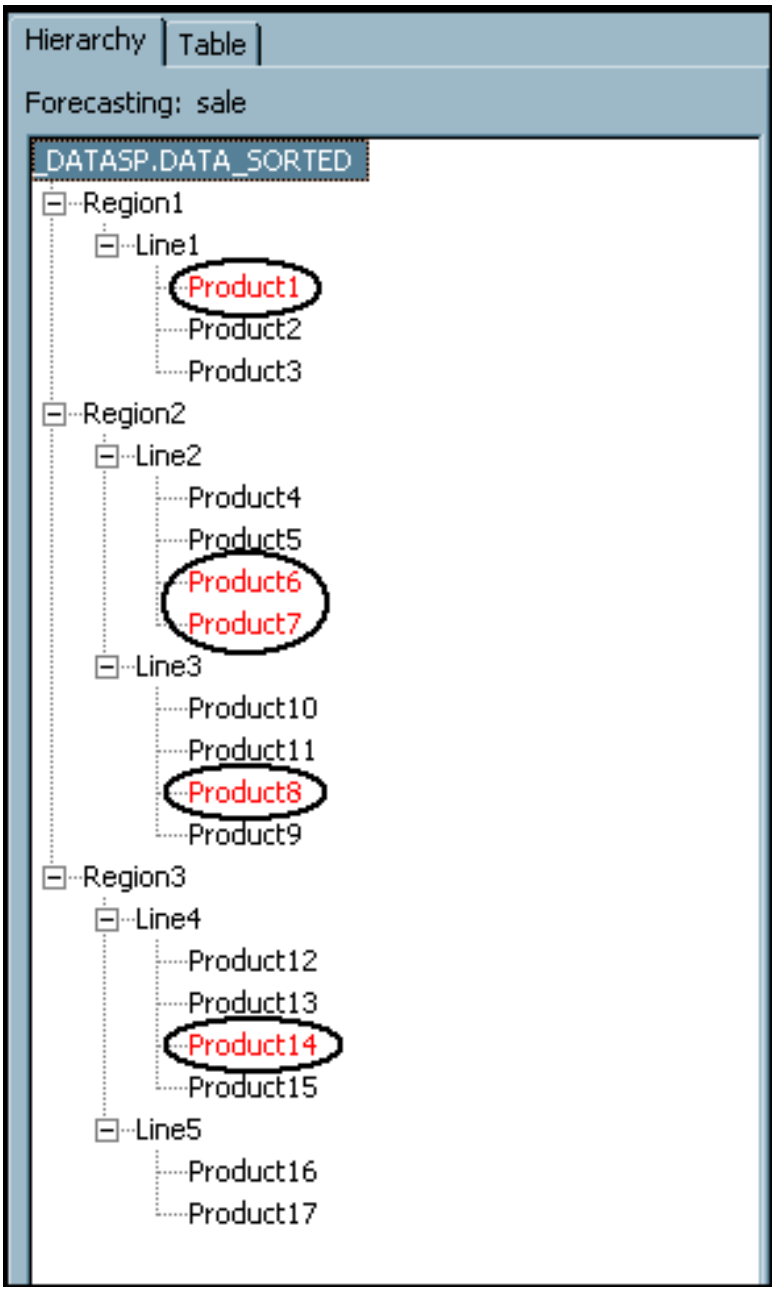

Figure 6.3. Series in Hierarchy Tab That Are Identified As Exceptions

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 New Filter 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. Select the **This filter will be used to identify exceptions** check box.
6. From the **Series Property** drop-down list, select **Mean Absolute Percent Error**.
7. From the **Condition** drop-down list, select **Greater than**.
8. In the **Value** text box, type **`7`**.
9. Click **Add Row**. A new row appears in the New Filter dialog box. SAS Forecast Studio automatically joins the filter conditions by AND.
10. From the **Series Property** drop-down list, select **Min**.
11. From the **Condition** drop-down list, select **Less than**.
12. In the **Value** text box, type **`300`**.

    The final window should look like the following:

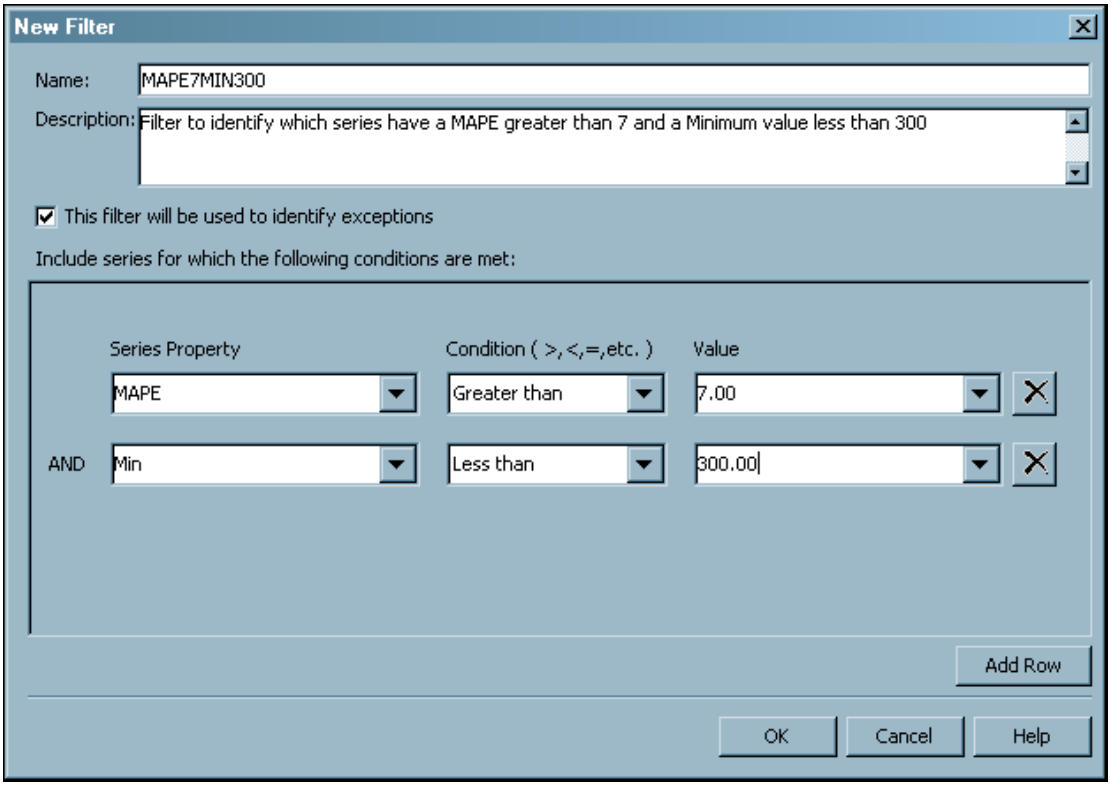

Figure 6.4. Completed New Filter Dialog Box for Filter Example

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

The filter that you created is now available from the **Show** drop-down list in the **Table** tab. Select **MAPE7MIN300** to view the series that meet these filter criteria.

Because you selected the **This filter will be used to identify exceptions** check box, any series that meets these criteria appears in bold red in the **Hierarchy** tab and has a red ✕ in the Exceptions column on the **Table** tab.

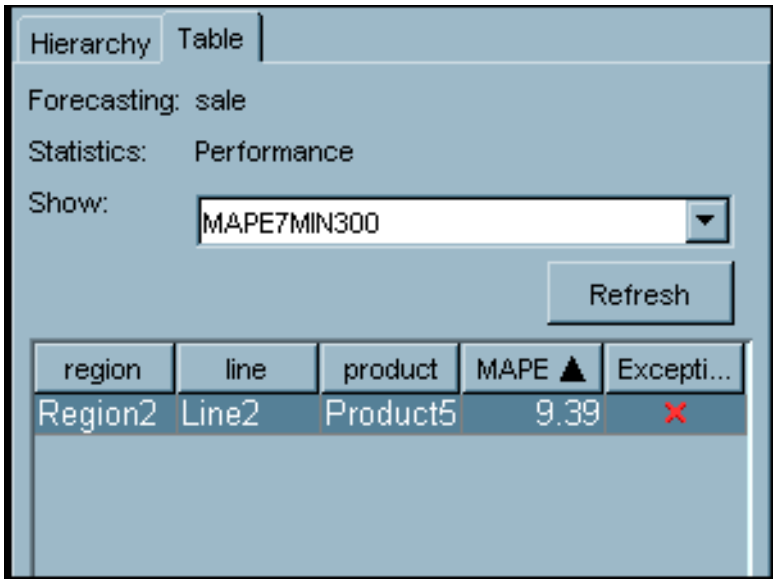

Figure 6.5. Series Identified As Exceptions in the Table Tab

Chapter 7
Working with Events

Chapter Contents

Chapter 7
Working with Events

What Is an Event?

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

For example, daily retail sales data 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.

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 Time Series	Description
pulse		Temporary change in the magnitude of a time series process. The magnitude returns to the former level immediately after the change.
level shift		Persistent change in the level of a time series process.
ramp	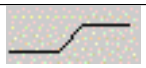	Persistent change in the trend or slope of a time series process.
temporary change		Temporary change in the magnitude. The magnitude decays to the former level after the change.

After you have created at least two events, you can combine events. For more information, see "Combine Events" on page 103.

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 60.
- Once you have created a project, you can manage your events by selecting **Project** → **Events**.

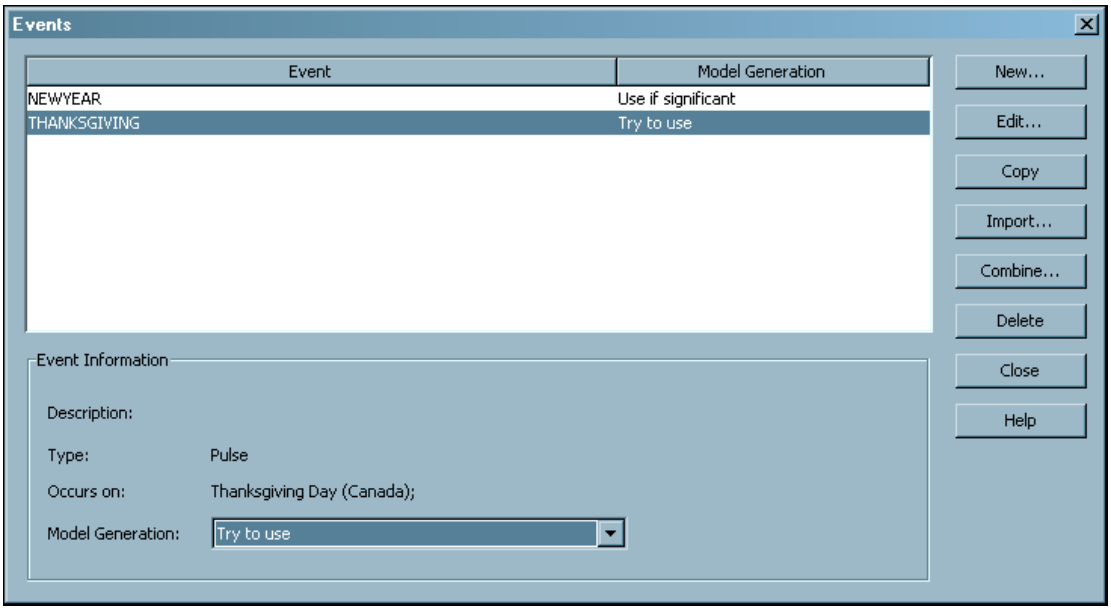

Figure 7.1. Events Dialog Box

The **Event** list displays the events that you have created and whether the event is included in the models. To view more information about an event, select the event from the list. The Event Information pane displays the description, the event type, and the date when the event occurs. You can also specify whether the event is included in the models. For more information, see "Include an Event in Model Generation" on page 104.

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

- create a new event
- edit an existing event
- copy an event
- import an event
- export an event

- combine events
- specify whether to include the event in model generation
- delete an event

Creating Events

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 New Event dialog box opens.

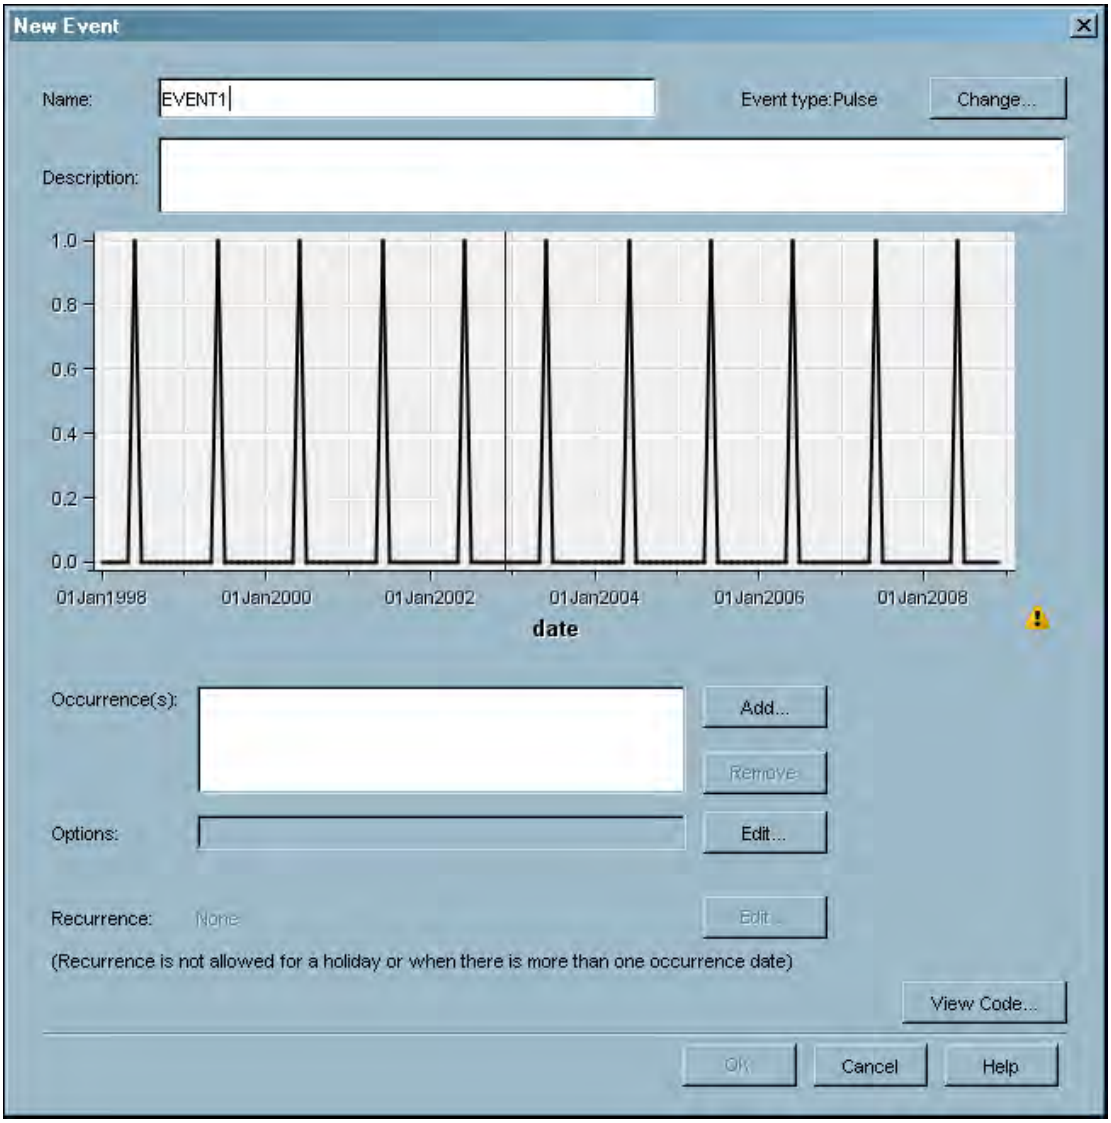

Figure 7.2. New Event Dialog Box

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.
5. Click **Change** to specify an event type. For more information, see "Specify an Event Type" on page 99.
6. Click **Add** to add an occurrence. For more information, see "Add an Occurrence" on page 99.

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

7. Click **Edit** to specify the options for the occurrence. For more information, see "Specify Options for the Occurrence" on page 100.
8. Click **Edit** to specify the recurrence. For more information, see "Specify an Event Recurrence" on page 100.

   Note: You cannot specify the recurrence of an event if the occurrence dates are noncontiguous.
9. (Optional) Click **View Code** to view the SAS code for the event.
10. Click **OK** to save the event. The new event appears in the Events dialog box.

After you have created two or more events, you can combine events. For more information, see "Combine Events" on page 103.

Specify an Event Type

When you create a new event, the default event type is Pulse. In 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 95.

Add an Occurrence

When you create or edit an event, you can specify the time period that the event occurs. The **Occurrence** box in the New Event or Edit *Event-name* dialog box displays a list of the occurrences for the event.

To add an occurrence, complete the following steps:

1. Click **Add** in the New Event or Edit *Event-name* dialog box.
2. Select the time period for the occurrence.

   Note: The dialog box that opens depends on the timestamp of your data.
3. Click **OK**.

The occurrence now appears in the **Occurrence** box. To specify the duration and time of the occurrence, see "Specify Options for the Occurrence" on page 100. To remove an occurrence, select the occurrence from the **Occurrence** box and click **Remove**.

Specify Options for the Occurrence

After you have added an occurrence to an event, you can specify the duration and time of the occurrence. These options apply to all the occurrences that you specified for the event.

To specify these options for each occurrence, complete the following steps:

1. Next to the Options box, click **Edit**. The Event Options dialog box opens.
2. Specify the duration of the occurrences.
3. Specify the time of the occurrences. You can choose from the following options:

   - **On the specified period** - specifies that the event occurs on the selected period.

     Note: If you select this option, then other options on the dialog box are not available.
   - **Relative to the specified period** - specifies that the event occurs before or after the selected period.

     If you select this option, then you can specify the direction and magnitude of the offset.

Specify an Event Recurrence

If you have specified a single occurrence for the event, then you can apply a recurrence.

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

1. Select event recurrence from the **Recurrence** drop-down list. The options available from the drop-down list depend on the interval of your data. For a list of the available options, see the PROC HPFEVENTS documentation in the *SAS High-Performance Forecasting User's Guide*.
2. If available, specify the end period for the recurrence. You can choose from the following options:

   - After a specified number of occurrences. The default number of occurrences is 12.
   - At a specified date and time. The default date is 12 occurrences from the first event in the data.
3. Click **OK**.

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 Edit *Event-name* dialog box opens. The options in this dialog box are the same as the New Event dialog box. For more information about the available options, see "Create an Event" on page 97.

   Note: After you have created an event, you cannot change the name of the event from the Edit *Event-name* 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 101.
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 Edit *Event-name* 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 101.

Import Events

You might want to include the events from an external event repository in SAS Forecast Studio. You can do this by importing the event repository.

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. To view the libraries on that server, double-click the name of the server.

4. To view the data sets within that library, double-click the name of the library.

   Note: To return to a previous level, click  .

5. 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.

   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. When you import the event, you should select the data set called *data-set-name*.

6. Click **OK**.

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

Export 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, SASUSER, or WORK libraries.

   Note: To return to a previous level, click  .

5. Specify a name for the data set in the **Data set name** 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. When you import the event, you should select the data set called *data-set-name*.

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 New Event dialog box opens.

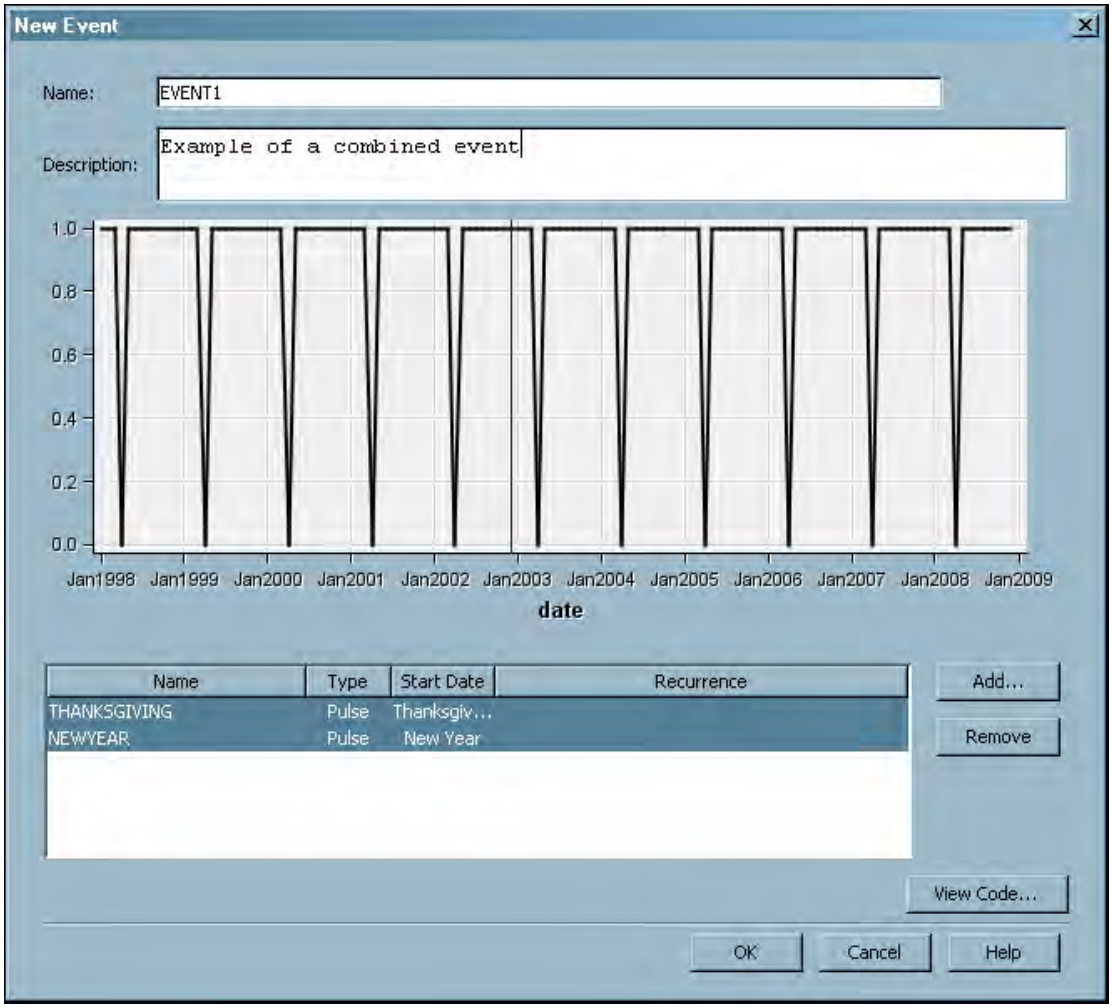

Figure 7.3. New Event Dialog Box

3. Specify a name for the new combined event.
4. (Optional) Provide a brief description for the new event.
5. From the list, select the events that you are combining. This table automatically includes the event(s) that you selected in the Events dialog box.

   To add an event, click **Add**. From the Select Event dialog box that opens, 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**. The Edit *Event-Name* dialog box opens. The options that are available from this dialog box are the same as when you created the combined event.

Include an Event in Model Generation

The **Model Generation** drop-down list in the Events dialog box specifies whether an event is used when SAS Forecast Studio generates a model. These models can be for the entire project or for a series. You specify the models that SAS Forecast Studio generates from the Model Generation panel in the Options dialog box. For more information, see "Specify the Model Generation Options" on page 118.

To specify the status of an event when SAS Forecast Studio generates a model, complete the following steps:

1. Select **Project** → **Events**. The Events dialog box opens.
2. Select the event.

   Note: To select multiple events, press CTRL and then select the events.
3. From the **Model Generation** drop-down list, you can select from the following options:

   - **Do not use in model generation** - 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 causal.
   - **Use if significant** - specifies to include the event in the model if feasible, if the resulting model is stable and causal, and if the event is significant.
   - **Force use** - specifies to include the event in the model if feasible.
4. Click **OK** in the Events dialog box.

The Model Generation column for all the events that you selected is updated.

You can also specify whether to include events in the models that you create. For more information, see "Include Events in User-Defined Models" on page 211.

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

Chapter Contents

Chapter 8

Generating Forecasts

Working with the Forecast Options for the Project

Overview of Forecasting Options

All forecasting options are initially set to a default value. To view the forecasting options for your project, select **Project** → **Options**.

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

- Data Preparation
- Accumulation
- Time ID
- Model Generation
- Model Selection
- Forecast
- Files

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

Specify the Data Preparation Options

To specify the data preparation options for your project, complete the following steps:

1. Select **Project** → **Options**. The Options dialog box opens.
2. In the selection pane, select **Data Preparation**.
3. In the Data Preparation panel, you can specify the following options:

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

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

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

Note: When the value of the **Missing value interpretation** option is changed, then SAS Forecast Studio performs a nondestructive diagnosis. SAS Forecast Studio preserves any user-created models. Any system-generated models are recreated.

Leading/trailing zero interpretation

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 the value of the **Leading/trailing zero interpretation** option is changed, then SAS Forecast Studio performs a nondestructive diagnosis. SAS Forecast Studio preserves any user-created models. Any system-generated models are recreated.

Leading/trailing missing value interpretation

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 the value of the **Leading/trailing missing value interpretation** option is changed, then SAS Forecast Studio performs a nondestructive diagnosis. SAS Forecast Studio preserves any user-created models. Any system-generated models are recreated.

Out of Sample range (periods from end)

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.

Start date

specifies the first date to use in the data preparation. You can choose from the following options:

- **Use earliest date in data set** - uses datetime value at the beginning of the data. This is the default.
- **Specify start date** - enables you to specify the first date to use.

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 change the 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.

End date

specifies the last date to use in the data preparation. You can choose from the following options:

- **Use latest date in data set** - uses datetime value at the end of the data. This is the default.
- **Specify end date** - enables you to specify the last date to use.

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

Note: You cannot change the end date if the time interval for your data is by hour, by minute, or by second. By default, SAS Forecast Studio uses the latest date in the data set.

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

4. Click **OK** to save your changes and to reforecast the data.

Specifying the Accumulation Options

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.

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.

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

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

Let $\boldsymbol{a} = Accumulate(\boldsymbol{R})$ 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.

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

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

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

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

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

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

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

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

Maximum

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

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

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

Median

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

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

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

Number of Non-Missing Observations

accumulates the vector values based on the number of non-missing values.

$\boldsymbol{a} = \boldsymbol{Q_N}$

Number of Observations

accumulates the vector values based on the number of values.

$\boldsymbol{a} = \boldsymbol{Q}$

Number of Missing Observations

accumulates the vector values based on the number of missing values.

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

First Occurrence

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

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

Last Occurrence

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

$\boldsymbol{a} = \boldsymbol{r_Q}$

Standard Deviation

accumulates the vector values based on their standard deviation.

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

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

Uncorrected Sum of Squares

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

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

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

Corrected Sum of Squares

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

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

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

Specify the Accumulation Options

To specify the accumulation options, complete the following steps:

1. Select **Project** → **Options**. The Options dialog box opens.
2. In the selection pane, select **Accumulation**.

   In the Accumulation panel, you can specify the method for accumulating data in the case of more than one observation per time period.
3. In the table, for each dependent variable, independent variable, and reporting variable select the accumulation method to use. (The reporting variables are listed as auxiliary in the Type column.) You can choose between the following options:

   - **Number of missing values**
   - **Median of values**
   - **Number of non-missing values**
   - **None** - No accumulation occurs
   - **Maximum of values**
   - **First value**
   - **Minimum of values**
   - **Last value**
   - **Corrected sum of squares of values**
   - **Uncorrected sum of squares of values**
   - **Total number of values**
   - **Sum of values**
   - **Standard deviation of values**
   - **Average of values**

   Note: The default value for each variable depends on the value of the aggregation option that you selected when you created the project. In the New Project Wizard, you specify the aggregation options if you selected to forecast hierarchically. If you did not forecast your data hierarchically, then the default is **Sum of values**.

   For more information about each of the accumulation methods, see "Understanding the Accumulation and Aggregation Methods" on page 112.
4. Click **OK** to save your changes and to reforecast the data.

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 recreated.

Specifying Options for the Time ID

Understanding SAS Time Intervals

You assign the time ID variable when you set up the data for the project by 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 60.

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 lenth 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 if they are spaced according to the interval that you specified.

Time intervals are specified in SAS by using character strings. Each of these strings are 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.

Specify the Time Interval

To specify the time interval, complete the following steps:

1. Select **Project** → **Options**. The Options dialog box opens.
2. In the selection pane, select **Time ID**.
3. In the Time ID panel, you can specify the following options for the time interval:

   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**. An interval is selected automatically if SAS Forecast Studio can determine the interval from the time ID variable.

Note: By default, the values in the **Interval** drop-down list and the **Seasonal cycle length** box are determined by the variable that you assign to the time ID role. For example, if the time ID variable contains data values that are spaced one month apart, then **MONTH** is the default selection in the **Interval** drop-down list, and the **Seasonal cycle length** box contains the corresponding default seasonal cycle length.

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.

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.

Interval Name	Units for the Shift	Maximum Shift Value	Example
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	HOUR8.7 specifies eight-hour intervals starting at 6:00am, 2:00pm, and 10:00pm.
MINUTE	MINUTE	multiplier	
SECOND	SECOND	multiplier	

Dates include time of day
specifies that the dates in the time ID variable include the time of day.

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**. For more information, see "Edit the Format of the Time ID Variable" on page 117.

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 recreated.

Select which days to define as "weekend" (inactive) days
enables you to select which days are the weekend or inactive days in the week. You can only select these days if you selected **WEEKDAY** as the interval.

Note: To reset the options for the time ID variable to their default values, click **Reset to Defaults**. Any changes that you have made are lost.

4. Click **OK** to save your changes and to reforecast the data.

Edit the Format of the Time ID Variable

You can edit the format of the time ID variable by using the Options dialog box. Editing the time ID format enables you to specify a user-defined format for the time ID variable.

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 recreated.

To edit the format of the time ID variable, complete the following steps:

1. Select **Project** → **Options**. The Options dialog box opens.
2. In the selection pane, select **Time ID**.
3. In the Time ID panel, click **Edit** next to the **Format** field. The Edit Time ID Format dialog box opens.
4. In this dialog box, select from the following options:

   - **System recommended** - specifies the format the SAS Forecast Studio determined from the date format of the time ID variable. This is the default.
   - **Select from list** - enables you to select a predefined format from the drop-down list.
   - **Custom** - enables you to specify a custom format in the text box.

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

5. Click **OK** to apply the format to the time ID variable.

Specify the Model Generation Options

From the Model Generation 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.

You must select at least one system-generated model (ARIMA, exponential smoothing, or unobserved components) that SAS Forecast Studio will generate. You can optionally select a model selection list to use when generating the models.

To specify the model generation options, complete the following steps:

1. Select **Project** → **Options**. The Options dialog box opens.
2. In the selection pane, select **Model Generation**.
3. In the Model Generation panel, you can specify the system-generated models for SAS Forecast Studio to create for each series. You can choose from the following options:

   ARIMA
   : 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.

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

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

4. If you choose to create a system-generated model, then you can also specify the following options:

   Dependent variable transformation
   specifies the transformation to apply to the dependent variable. You can select from the following transformations:

   - no transformation
   - 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. Forecasts can be based on the mean or median. By default, the mean value is used.

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

5. If you want to generate models from an existing model list, select the **Try models from an existing model list (INSELECT)** check box. This option 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 "Time Series Forecasting Models" on page 254.

   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 **Try models from an existing model list (INSELECT)** check box. SAS Forecast Studio will continue to use the external model selection list that you selected.

6. Click **OK** to save your changes and to reforecast the data.

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

Specify the Model Selection Options

From the Model Selection panel, you can specify the options to use when fitting the models.

To specify the model selection options, complete the following steps:

1. Select **Project** → **Options**. The Options dialog box opens.
2. In the selection pane, select **Model Selection**.
3. In the Model Selection panel, you can specify the following options:

Detect outliers
specifies whether SAS Forecast Studio should detect outliers in the data when fitting an ARIMA model.

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

Use holdout sample
specifies the holdout sample to use when fitting the models. You can specify values for the following options:

- **Number of periods to use** 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 1.
- **Maximum percentage of series to use** specifies the size of the holdout sample as a percentage of the length of the time series. The default is 100 (100%), which means no restriction on the holdout sample size based on the series length. 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*.

At least this number of observations are needed for a trend model
specifies minimum number of observations for a trend model.

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

At least this number of observations are needed 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 HPFDIAGNOSE procedure in the *SAS High-Performance Forecasting User's Guide*.

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 B, "Statistics of Fit."

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

4. Click **OK** to save your changes and to reforecast the data.

Specify the Forecast Options

To specify the forecast options, complete the following steps:

1. Select **Project** → **Options**. The Options dialog box opens.
2. In the selection pane, select **Forecast**.
3. In the Forecast panel, you can specify the following options:

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.

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.

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

4. Click **OK** to save your changes and to reforecast the data.

Specify the Files Options

To specify the files options, complete the following steps:

1. Select **Project** → **Options**. The Options dialog box opens.
2. In the selection pane, select **Files**.

3. In the Files panel, you can select the **Generate index files** option. When selected, SAS Forecast Studio indexes the series in the project so that it is faster to navigate between series in the application. If you only want to run the project code in batch mode, do not select this option.

   Note: To reset the files option to its default values, click **Reset to Defaults**. Any changes that you have made are lost.

4. Click **OK** to save your changes and to reforecast the data.

Reforecast the Data

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

To reforecast your data, complete the following steps:

1. Select whether to reforecast the current series or all the series in the project.

   - To reforecast the current series, select the series in the **Hierarchy** tab or **Table** tab and select **Series** → **Reforecast**. The Forecast Series dialog box opens.
   - To reforecast all the series in the project, select **Project** → **Reforecast** or click  in the toolbar. The Forecast Project dialog box opens.

2. Select how to forecast the series or project. The following options are available:

   - **Recreate automatically-generated models, refit all models, and select a forecast model**.
   - **Refit all models, and select a forecast model** - This is the default.
   - **Refresh the current forecast model, updating the parameter values**.
   - **Refit the current forecast model, using the same parameter values**.

   Note: If you select the **Refresh the current forecast model, updating the parameter values** option, then the OUTSTATSELECT= data set is not created. As a result, the statistics of fit that appear in the model selection list and are the basis for selecting the best fitting model are not updated. For more information about this data set, see "Output Data Sets in Project Directory" on page 47.

3. If you are reforecasting the entire project, you can select the **Use updated data if available** check box. SAS Forecast Studio checks to see if the data source has been updated, and if so, updates the project data and reforecasts the existing series based on the option that you selected.

   Note: If you choose any option other than **Recreate automatically-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 "Best Smoothing Model" on page 253.

4. Click **OK** to reforecast the series or project. When finished, the Forecast Summary dialog box opens. For more information, see "Reviewing a Summary of the Forecast Run" on page 65.

Note: From the Model View, you can view the model selection list for the selected series in the **Hierarchy** or **Table** tab. The model selection list shows the forecast model for the selected series. By default, SAS Forecast Studio selects the best model for each series when a project is created.

However, using the **Forecast Model** drop-down list in the Model View, you can change the forecast model for a series. SAS Forecast Studio selects the forecast model based on the following criteria:

- 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 model selection list, see "What Is the Model Selection List?" on page 222.

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 deselect 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 recreates 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 **Table** tab, select the series that you want to make inactive.
2. Select **View** → **Forecasting View** to open the Forecasting View.
3. In the data table, clear the **Active** check box.

   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

Chapter Contents

Chapter 9
Working with Overrides

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** tab 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 132.

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 **End date** option in the Options 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.

  If you specify an end date (that is not the default value of the **End date** option), then the start of the horizon is affected as follows:

  - If the end date that you specify is after the end date for the series, then the horizon is not affected by the end date that you specified. You can enter overrides for those time ID values that occur one interval after the end date for the series.
  - If the end date that you specify is earlier than the end date of the series, then the horizon starts one interval after the value that you specified for the end date. However, if you specify an end date that is not in the input data set, then the horizon is the first value in the data set that is one interval after the specified end date. For example, in your input data set, the interval of the time ID values is month, and the dates for the time ID values are the first of the month. If you specified November 5, 2005, as the end date, then the horizon starts at January 1, 2006. However, if you specified November 1, 2005, as the end date, then the horizon starts at December 1, 2005.

  For more information about the **End date** option, see the documentation for the Data Preparation options in "Working with the Forecast Options for the Project" on page 109.

- You cannot add overrides when the value for the **Horizon** option is less than the number of the **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** tab. 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.

  For more information about these options, see the documentation for the Forecast and Data Preparation options in "Working with the Forecast Options for the Project" on page 109.

- 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 Forecast panel of the Options dialog box. For more information, see "Specify the Forecast Options" on page 121.

How Are Overrides Processed?

How SAS Forecast Studio processes overrides depends on the advanced reconciliation options that you selected and whether conflicts are detected. For more information about the advanced reconciliation options, see "Understanding the Reconciliation Options" on page 137. For more information about conflicts, see "What Is an Override Conflict?" on page 133.

If you selected the **Require that conflicts be resolved before reconciliation** option as the advanced reconciliation option, then SAS Forecast Studio uses the following process:

1. You add an override. A "Reconciliation is out of date" message appears at the top of the Forecasting View.
2. You click **Update** to reconcile the hierarchy. SAS Forecast Studio checks for conflicts in the locked overrides.
3. 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. A message about unreconcilied nodes appears in the Forecasting View. You can click **Report** to view the Reconciliation Failure Report. For more information, see "About the Reconciliation Failure Report" on page 139.

     Note: Unreconciled nodes occur only if you also selected **Restrict to direction implied by reconciliation method** option as an advanced reconciliation option. If you selected the **No restriction** option instead, then this message does not appear.
4. If conflicts are detected, the Override Conflicts dialog box opens. You view and resolve any conflicts. For more information, see "Resolve Override Conflicts" on page 135.

   SAS Forecast Studio does not reconcile the hierarchy till all the conflicts are resolved. 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. A message about unreconcilied nodes appears in the Forecasting View. You can click **Report** to view the Reconciliation Failure Report. For more information, see "About the Reconciliation Failure Report" on page 139.

     Note: Unreconciled nodes occur only if you also selected **Restrict to direction implied by reconciliation method** option as an advanced reconciliation option. If you selected the **No restriction** option instead, then this message does not appear.

If you selected the **Ignore conflicts** option as the advanced reconciliation option, then SAS Forecast Studio uses the following process:

1. You add an override. A "Reconciliation is out of date" message appears at the top of the Forecasting View.
2. 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. A message about unreconcilied nodes appears in the Forecasting View. You can click **Report** to view the Reconciliation Failure Report. For more information, see "About the Reconciliation Failure Report" on page 139.

  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 130.
- 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 131.

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 "Specify the Forecast Options" on page 121.

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 132.

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 132.

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 127.

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 127.

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

   - Select the **Adjust the reconciled statistical forecast** option to base the override values on the specified increase or decrease of the reconciled statistical forecast.

     Note: If you did not forecast your data hierarchically or there is no reconciliation forecast available, then the **Adjust the statistical forecast** option is available. You can select this option to base the override values on the specified increase or decrease of the statistical 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 **Increase** or **Decrease**. In the text box, specify the value and in the second drop-down list, specify whether this value is in units or percent.

     For example, you want to create an override that is 10% greater than the reconciled statistical forecast. In the first drop-down list, select **Increase**. 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 fixed 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 this value equally**
     – **Split this value proportional to reconciled statistical forecast**
     Note: If you did not forecast your data hierarchically or there is no reconciliation forecast available, then the **Split this value proportional to statistical forecast** option is available.

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

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  .

You can also unlock overrides from the Override Conflicts dialog box. For more information, see "Resolve Override Conflicts" on page 135.

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 **Table** tab.

Series that have an associated note are idenitified in the **Hierarchy** and **Table** tabs 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 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 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.

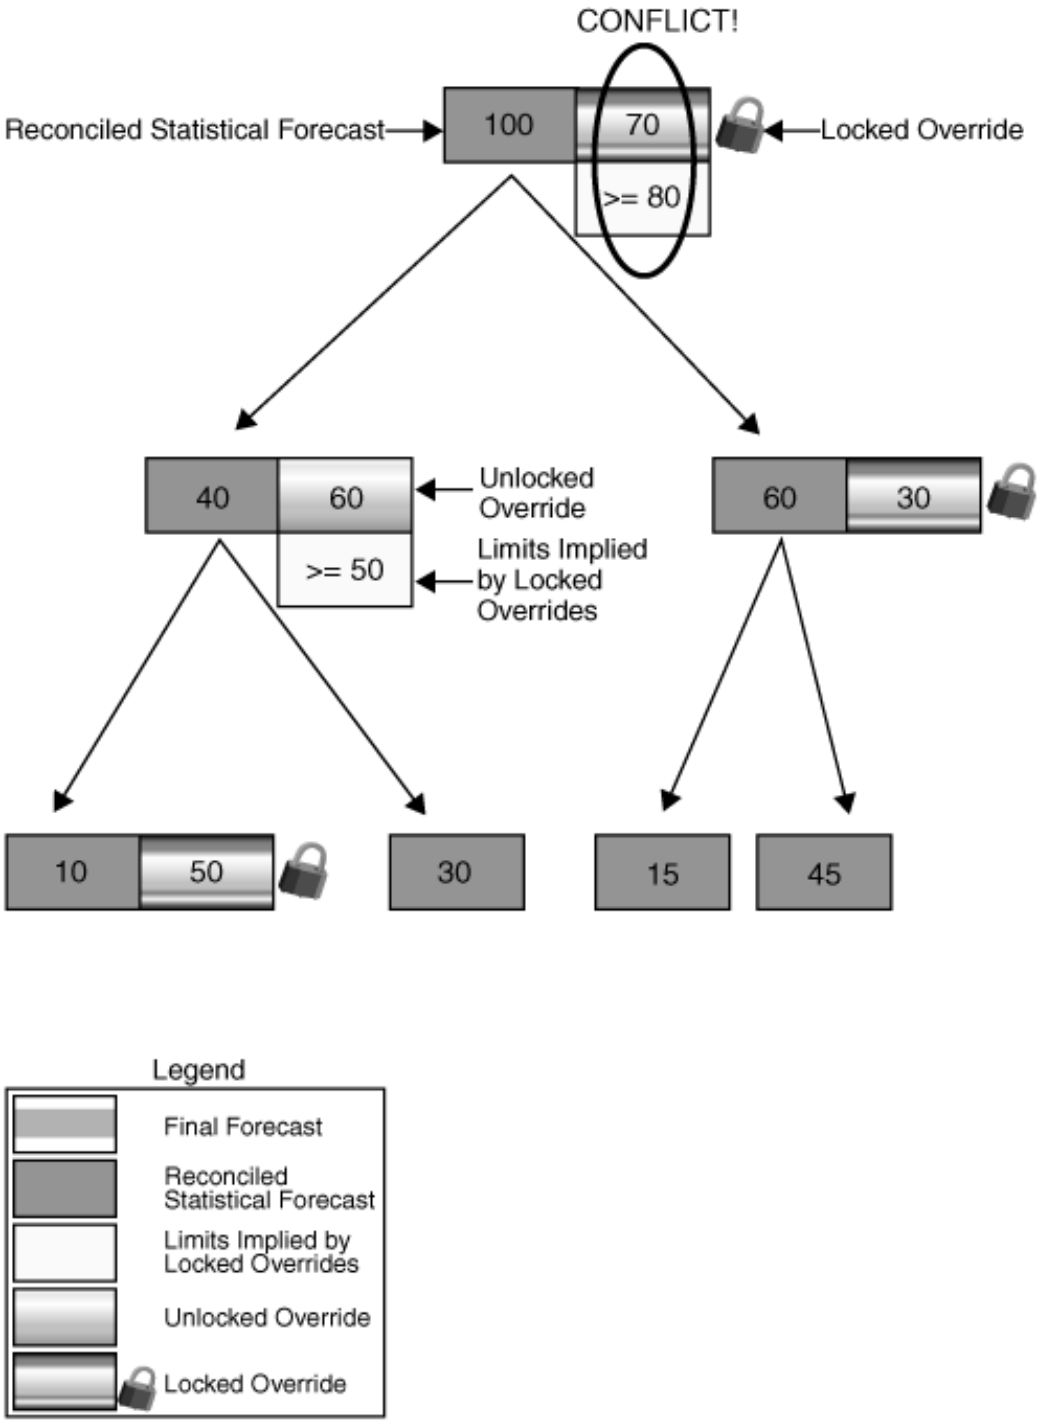

Figure 9.1. Example of an Override Conflict

Resolve Override Conflicts

SAS Forecast Studio detects conflicts between locked overrides. If a conflict occurs, then this conflict appears in the Override Conflicts dialog box.

To resolve conflicts, complete the following steps:

1. From the **Select date** pane, select a date. The **Select Date** pane lists by date all the locked overrides that have conflicts.

   The right pane of the dialog box shows the conflict for the date 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 127.
3. Click **Apply**.

   If there are still conflicts, the window does not close. You might need to select another date from the **Select date** pane.

   If there are no more conflicts, then 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 choose to forecast hierarchically, then you need to 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 136.

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 62.
   - 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 method to use when resolving 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.
4. 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 overriden. 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 137.

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 136.

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 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, in which case the hierarchy remains unreconciled. • You can resolve any conflicts, in which case SAS Forecast Studio reconciles the hierarchy when you close this dialog box.
Ignore conflicts	**No restriction**	When you click **Reconcile** in the Forecasting View, SAS Forecast Studio determines if 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, in which case the hierarchy remains unreconciled. • You can resolve any conflicts, in which case 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 if 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 Forecasting View. Click **Report** to open the Reconciliation Failure Report.

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

This report can contain the following tabs:

Override Conflicts

shows the conflicts between overrides that were detected before reconciliation. By resolving these conflicts and reconciling the hierarchy again, you reduce the number of reconciliation failures that SAS Forecast Studio detects. The relationship between the number of override conflicts and the number of unreconciled nodes depends on how you set up your hierarchy. For more information about the advanced reconciliation options, see "Specify the Advanced Reconciliation Options" on page 136.

The functionality that is available from the Override Conflicts tab is the same as the functionality in the Override Conflicts dialog box. For more information about how to resolve override conflicts from within this tab, see "Resolve Override Conflicts" on page 135.

Note: If you selected the **Require that conflicts be resolved before reconciliation** option when you configured the hierarchy, then this tab is not available.

Unreconciled Nodes

shows the nodes in the hierarchy that could not be reconciled. The table that appears on this tab displays the following information:

- The time period for the unreconciled node.
- The name of the node.
- The final forecast of the unreconciled node. This forecast could have been set by a locked override.
- Any limits that were implied by a locked override. The unreconciled node could violate this limit.

  Note: This information is available only if you selected the **Ignore conflicts** and the **No restriction** options when you configured the hierarchy. For more information about the advanced reconciliation options, see "Specify the Advanced Reconciliation Options" on page 136.
- The final forecast of the unreconciled node based on the final forecasts of the nodes in the lower levels of the hierarchy.

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 62.
- **Sum of Values** is selected as the aggregation method. For more information about aggregation methods, see "Configure the Hierarchy" on page 62.
- The **Allow negative forecasts** option is not checked in the Forecast panel of the Options dialog box. For more information, see "Specify the Forecast Options" on page 121.

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 70.

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 140.

Explanation of the Reconciliation Process

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

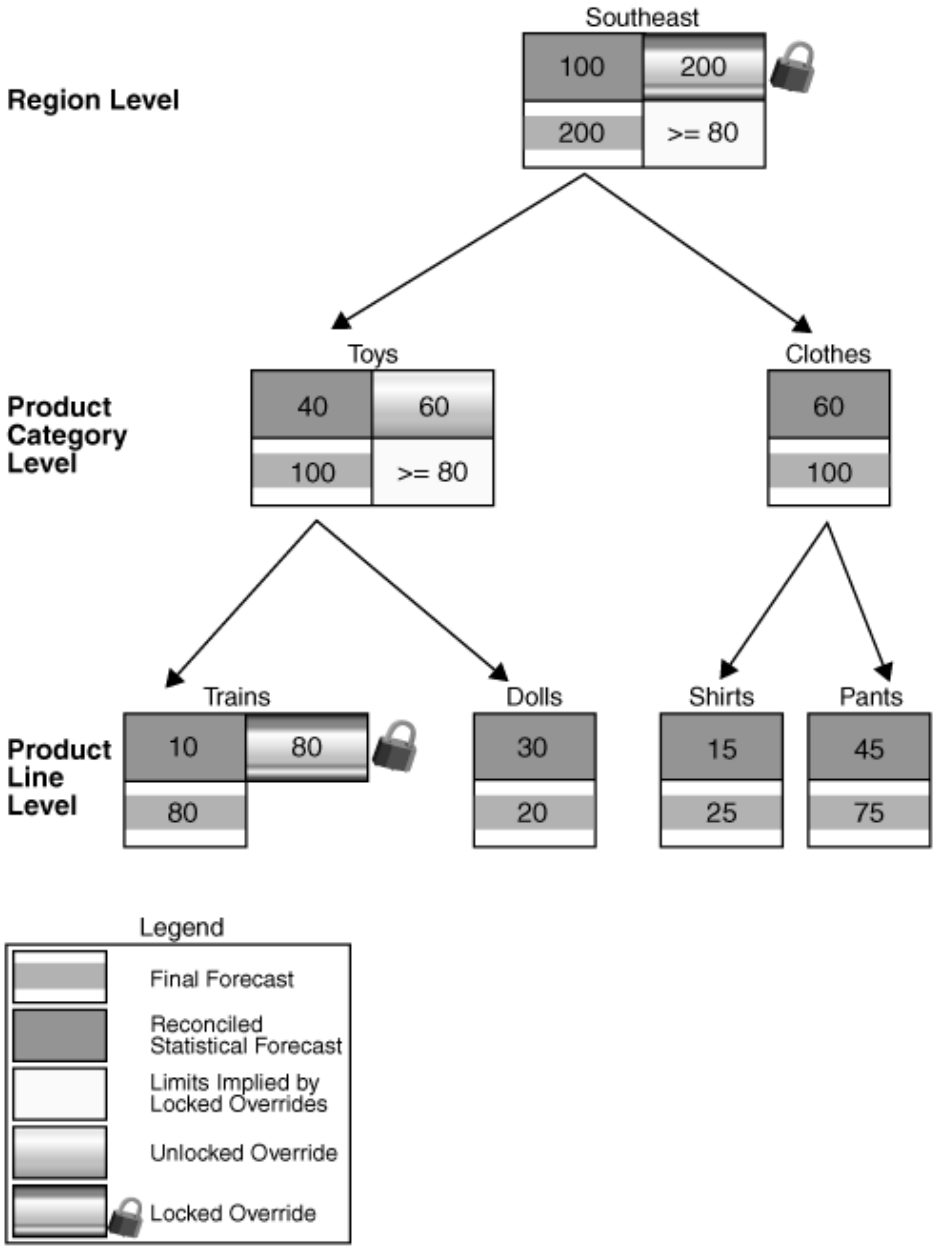

Figure 9.2. Example: SAS Forecast Studio Honors All Locked Overrides

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.

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. 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 70.

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 140.

Explanation of the Reconciliation Process

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

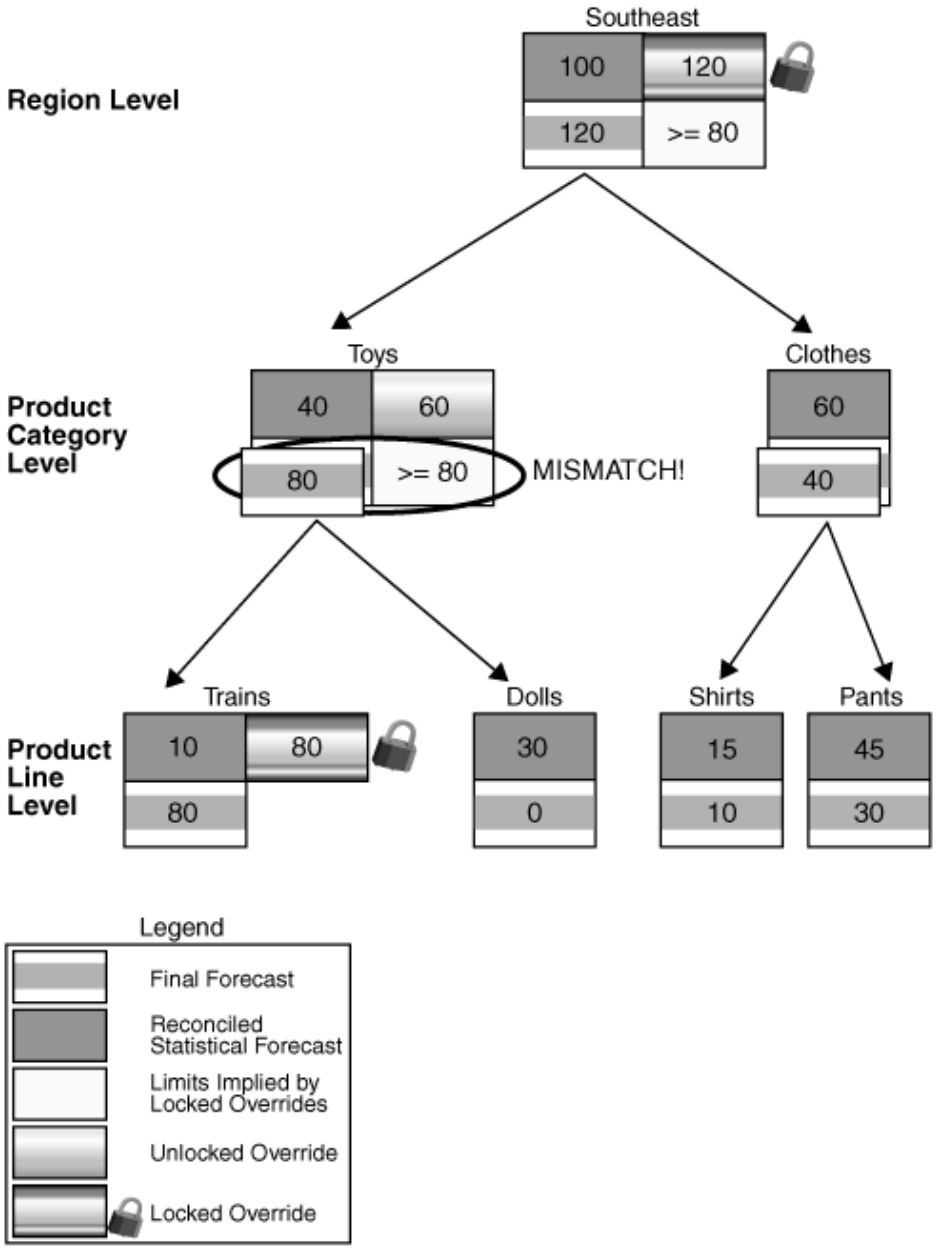

Figure 9.3. Example: Implicit Bottom-Up Method Used to Reconcile Hierarchy

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 70.

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 140.

Explanation of the Reconciliation Process

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

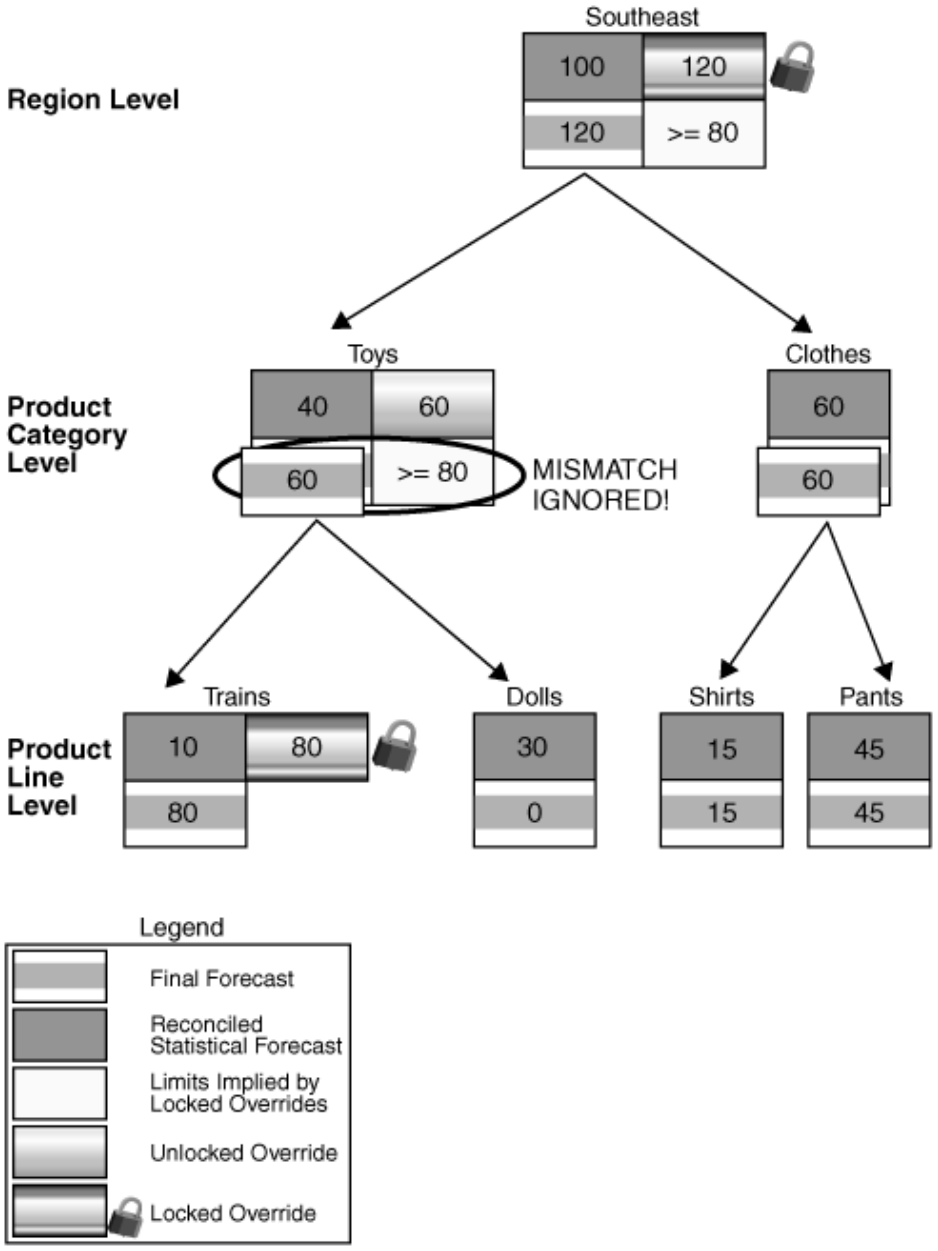

Figure 9.4. Example: Implicit Bottom-Up Method Used to Reconcile Hierarchy

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 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 70.

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 140.

Explanation of the Reconciliation Process

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

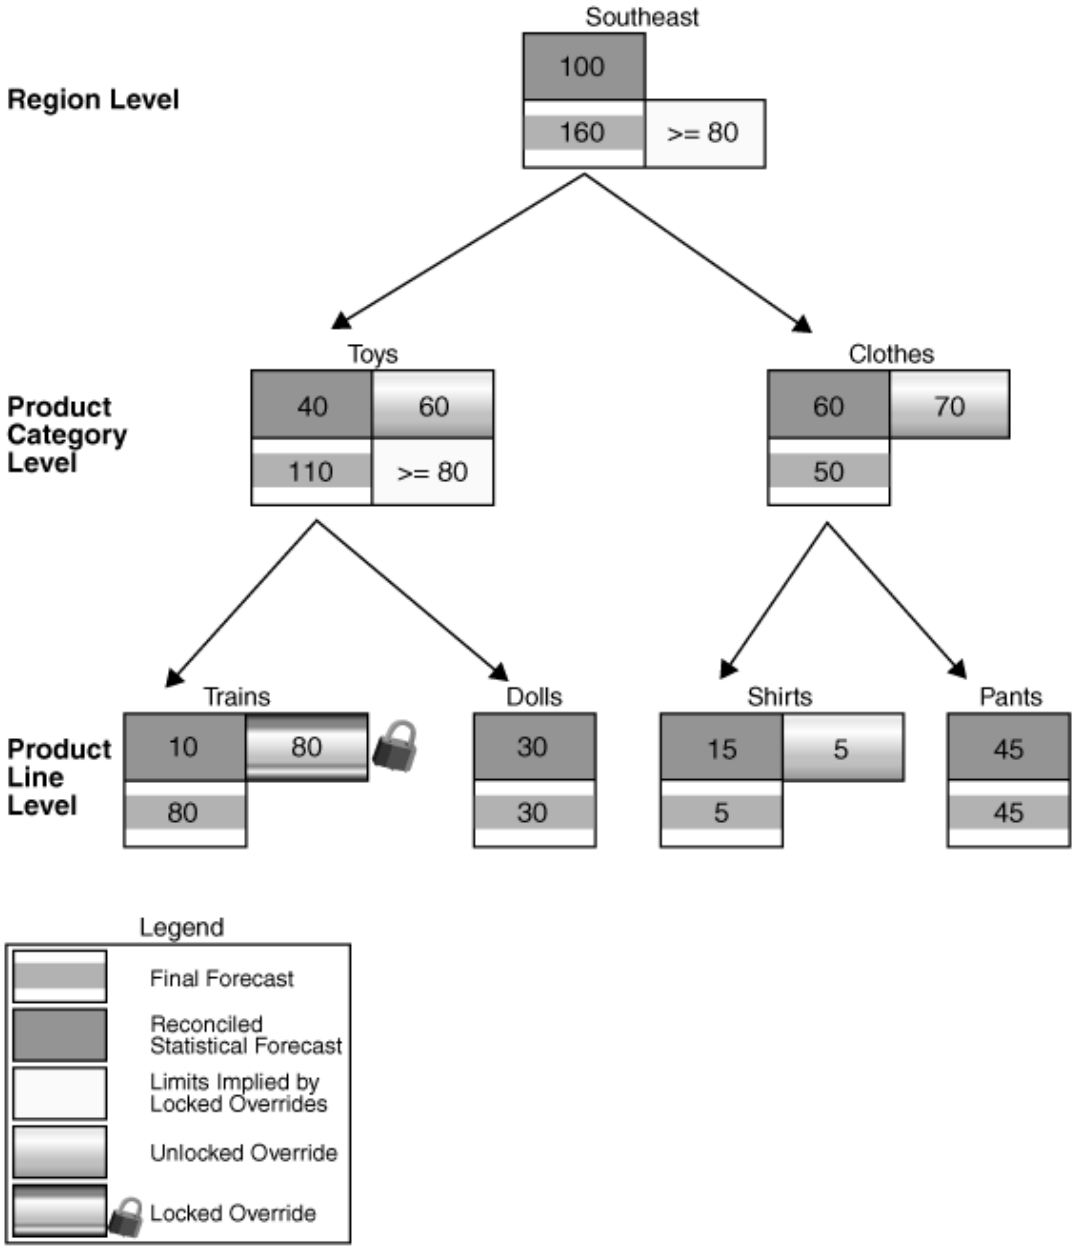

Figure 9.5. Example: Bottom-Up Method of Reconciliation with No Conflicts

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.

Level in the hierarchy	Reconciliation process for the nodes in this level
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 70.

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 140.

Explanation of the Reconciliation Process

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

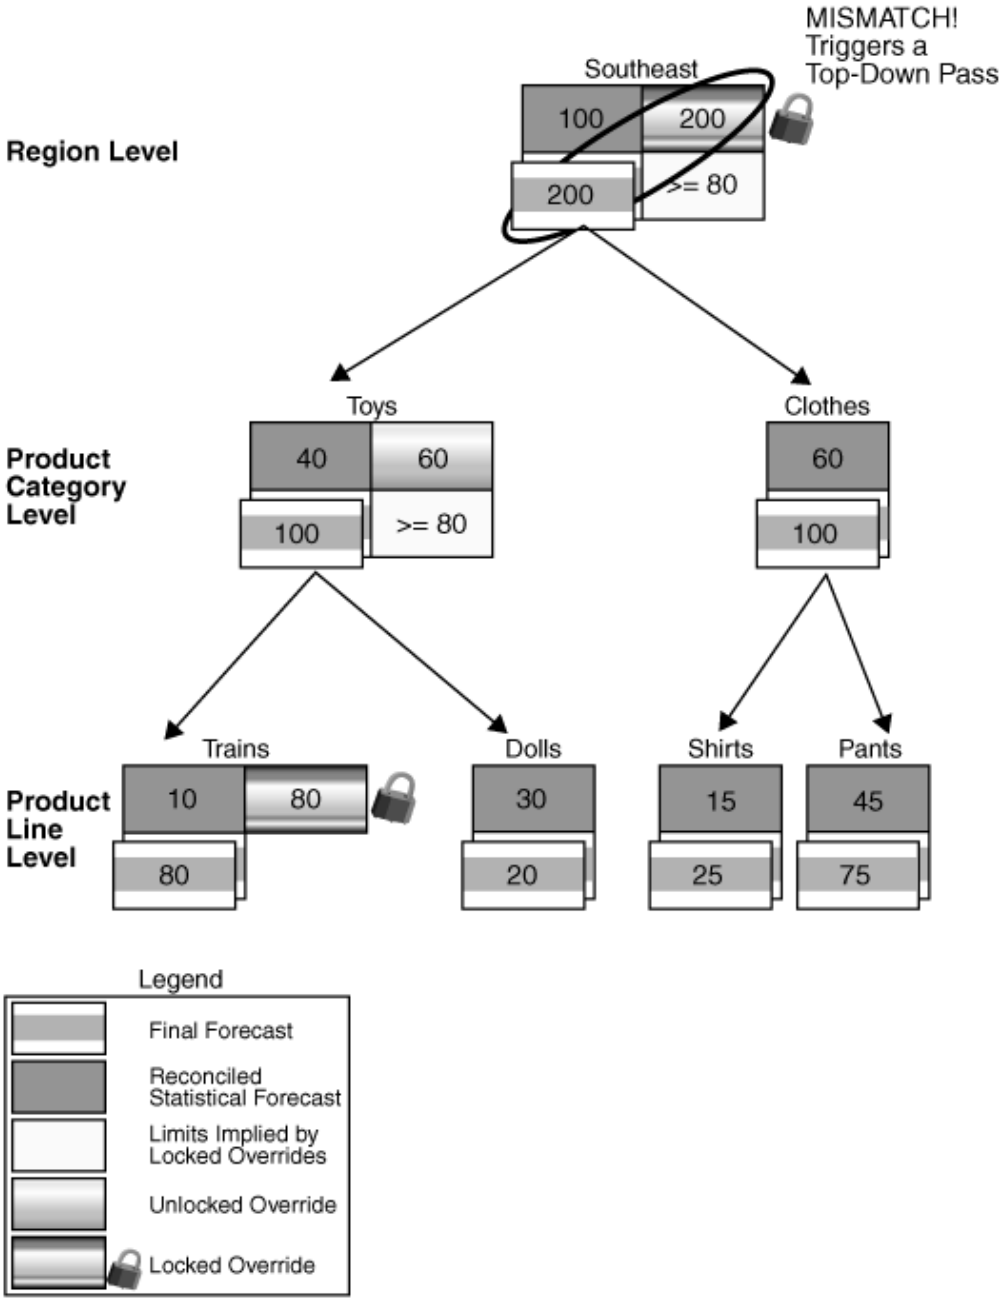

Figure 9.6. Example: Implicit Top-Down Pass in the Bottom-Up Method of Reconciliation

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, 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 restrictions 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 70.

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 140.

Explanation of the Reconciliation Process

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

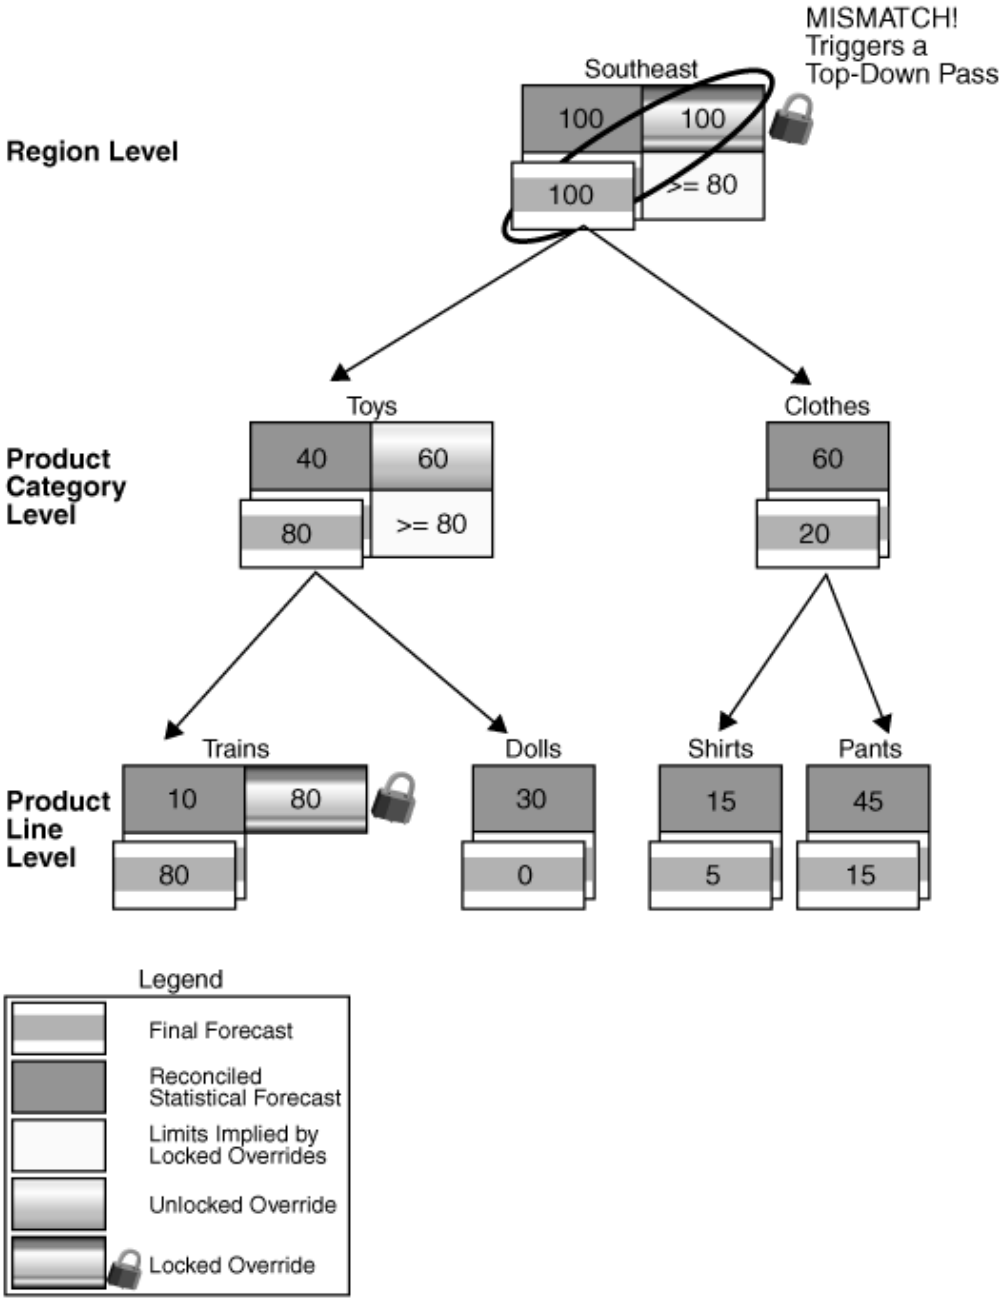

Figure 9.7. Example: Implied Limits Have Precedence over an Unlocked Override

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 70.

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 140.

Explanation of the Reconciliation Process

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

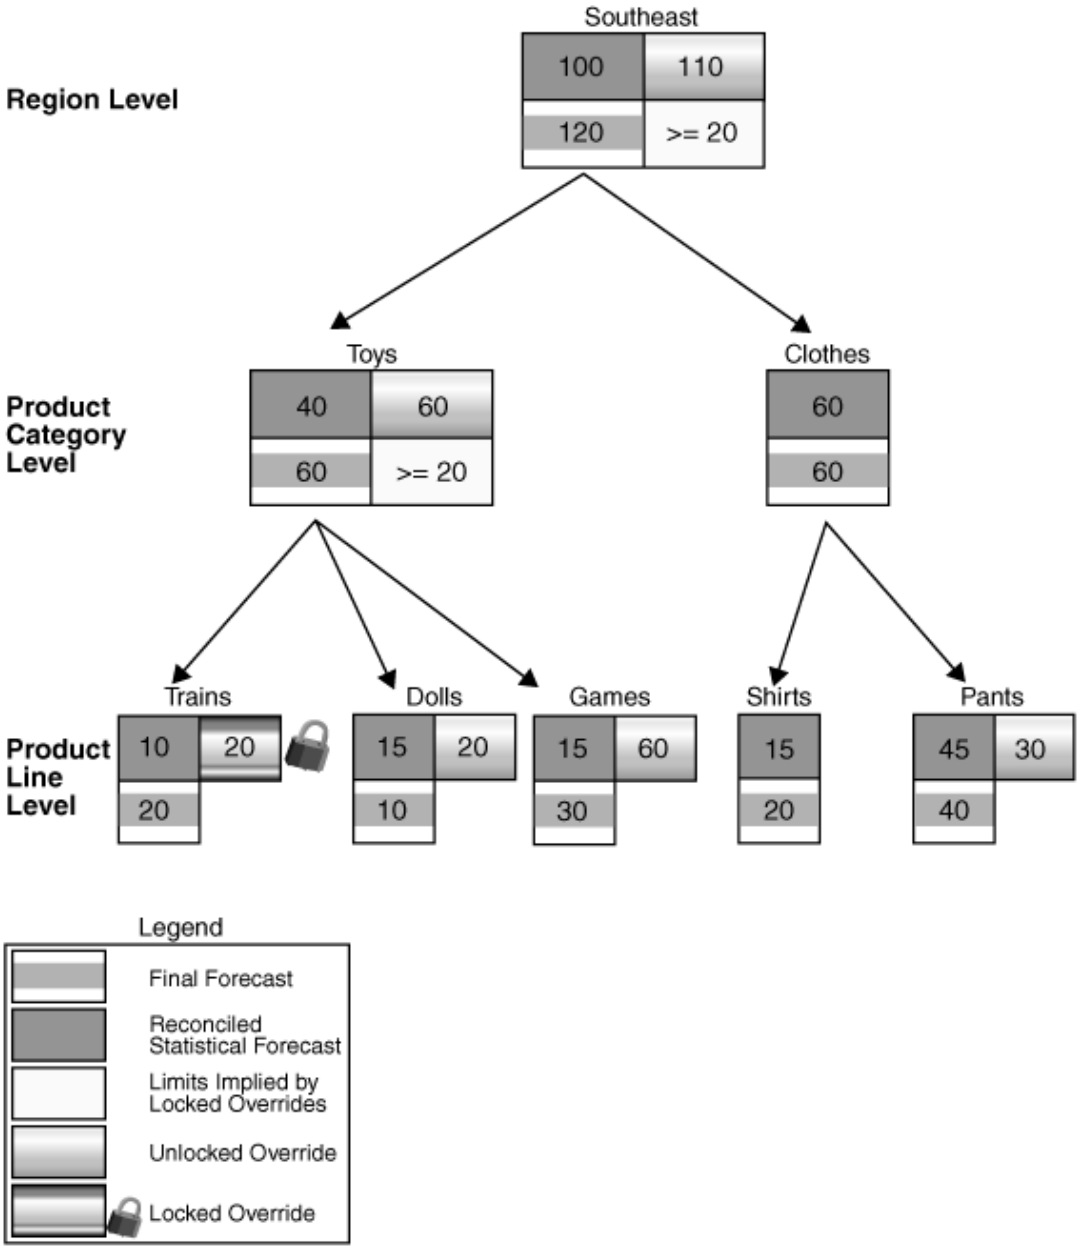

Figure 9.8. Example: Middle-Out Method of Reconciliation

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 the middle-out method of reconciliation, 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 70.

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 140.

Explanation of the Reconciliation Process

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

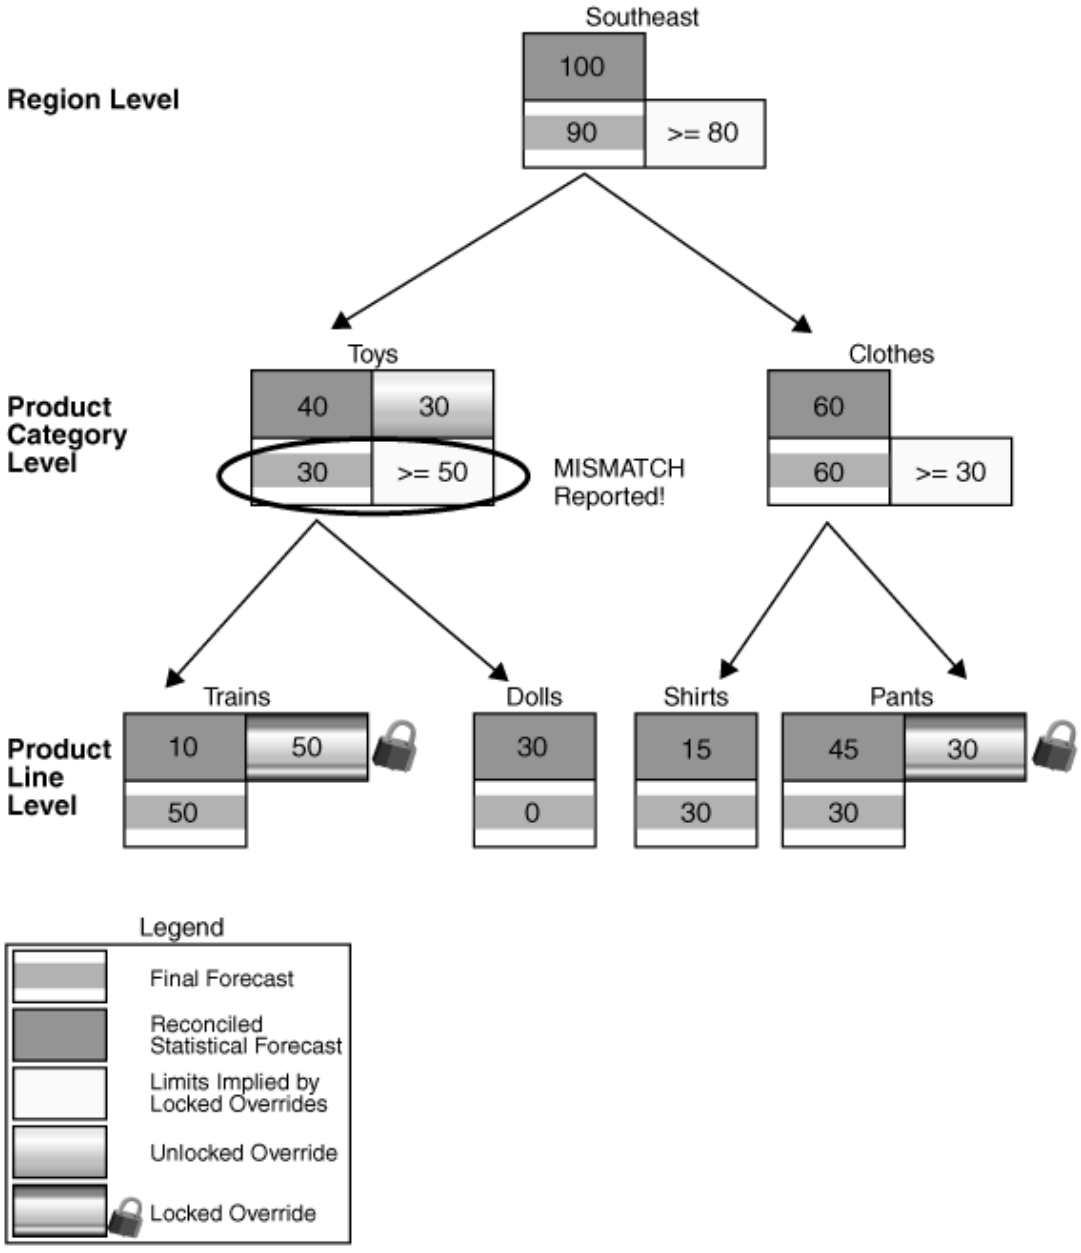

Figure 9.9. Example: Middle-Out Method of Reconciliation

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 the middle-out method of reconciliation, 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 forecast 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

Chapter Contents

Chapter 10
Creating User-Defined Models

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 Project Models dialog box. For more information, see "Create a New Custom Model Based on a Default Model" on page 221.

Create an Entirely New Model

To create a new model, complete the following steps:

1. Select **Project** → **Project Models**. The Project Models dialog box opens. For more information, see "Overview of Project Models" on page 219.

   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 227.
2. Click the **Custom Models** tab in the Project Models dialog box.
3. On the **Custom Models** tab, click **New**. The New Models dialog box opens.
4. In the New Models 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 176.

   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 180.

   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 183.

   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 193.

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 196.

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 199.

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 203.

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 202.

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 205.

5. Click **OK** to select an existing model or to create a new model. The dialog box that opens depends on the options that you selected. After you select or 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 optionally specify one or more input variables.

To create an ARIMA model, complete the following steps:

1. Select **Project** → **Project Models**. The Project Models dialog box opens. For more information, see "Overview of Project Models" on page 219.

   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 227.

2. Click the **Custom Models** tab in the Project Models dialog box.

3. On the **Custom Models** tab, click **New**. The New Models dialog box opens.
4. In the New Models dialog box, select **Create an entirely new model**.
5. From the **Model type** drop-down list, select **ARIMA** as the model type. The New ARIMA Model window opens.

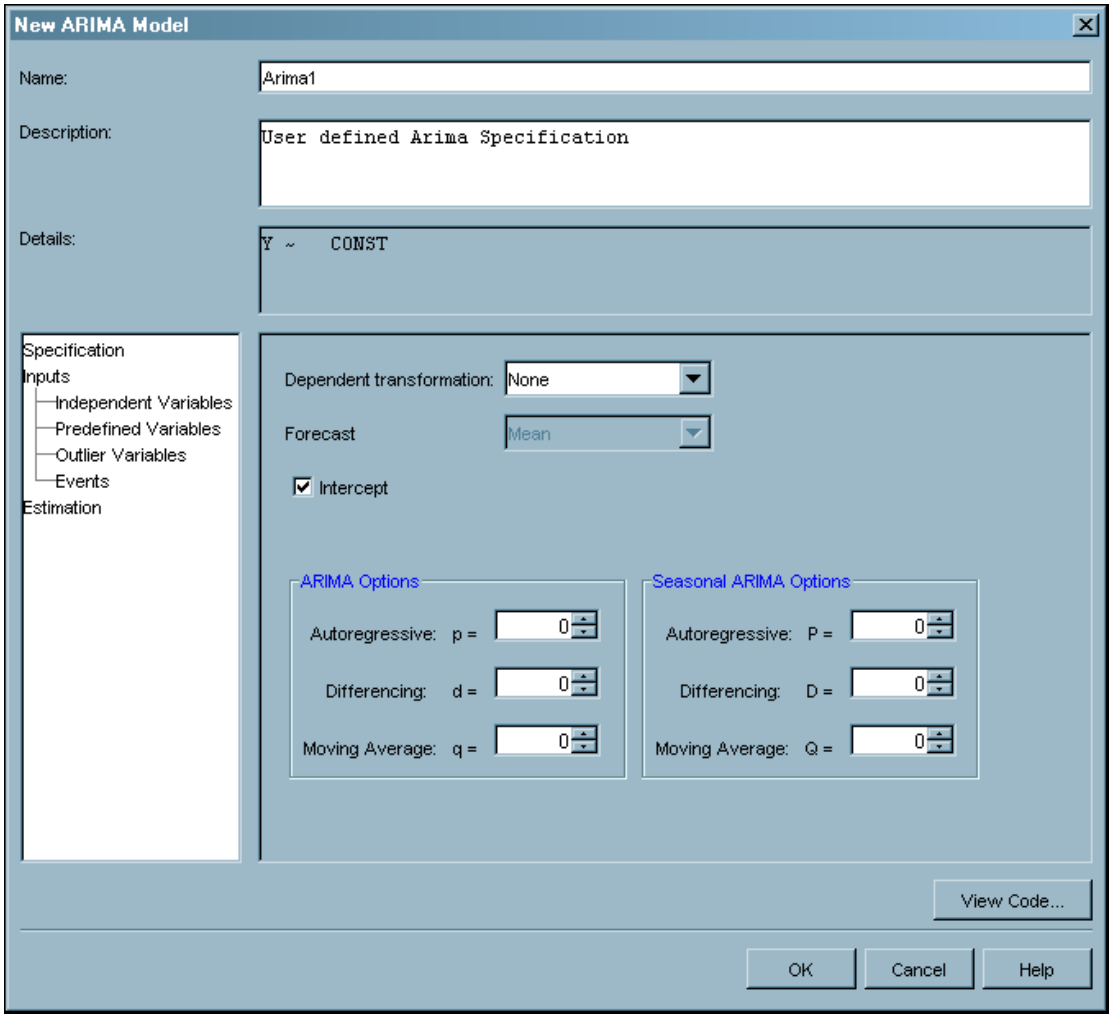

Figure 10.1. New ARIMA Model Dialog Box

Note: The Edit ARIMA Model window 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 207.

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 window. You cannot edit this text.

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

   - specification
   - independent variables
   - predefined variables
   - outlier variables
   - events
   - 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 window, select **Specification** in the selection pane to specify a basic nonseasonal (ARIMA(p, d, q)) or seasonal (ARIMA(p, d, q)(P, D, Q)) model.

You can specify the following options:

- The transformation to apply to the dependent series from the **Dependent Transformation** drop-down list. You can select from the following transformations:
  - no transformation
  - logarithmic transformation
  - square-root transformation
  - logistic transformation
  - Box-Cox transformation where the parameter number is between -5 and 5. Specify this parameter value in the Box-Cox Transformation window. For more information, see "Specifying a Box-Cox Transformation" on page 210.

  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 (p d q)(P D Q). 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 the following: 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 the following: 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 the following: 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, etc.) 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 207.

Independent Variables in ARIMA Models

By default, none of the variables that you assign to the Independent variables role are inputs in the model. However, using the ARIMA Model window, you can specify to include these independent variables in the model. You can also specify the input series options for each independent variable. For more information, see "Include Independent Variables in User-Defined Models" on page 211.

Predefined Variables in ARIMA Models

By default, none of the predefined variables from the input data set are inputs in the model. However, using the ARIMA Model window, you can specify to include these predefined variables in the model. For more information, see "Include Predefined Variables in User-Defined Models" on page 211.

Outlier Variables in ARIMA Models

By default, any outliers that are automatically detected during the creation of the model are inputs in the model. However, using the ARIMA Model window, you can specify to delete these outliers from the model. For more information, see "Include Outlier Variables in User-Defined Models" on page 212.

Events in ARIMA Models

By default, none of the events in the project are inputs in the model. However, using the ARIMA Model window, 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 211.

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 214.

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 window is similar to the ARIMA Model window, 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** → **Project Models**. The Project Models dialog box opens. For more information, see "Overview of Project Models" on page 219.

   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 227.
2. Click the **Custom Models** tab in the Project Models dialog box.
3. On the **Custom Models** tab, click **New**. The New Models dialog box opens.
4. In the New Models 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 New Subset ARIMA Model window opens.

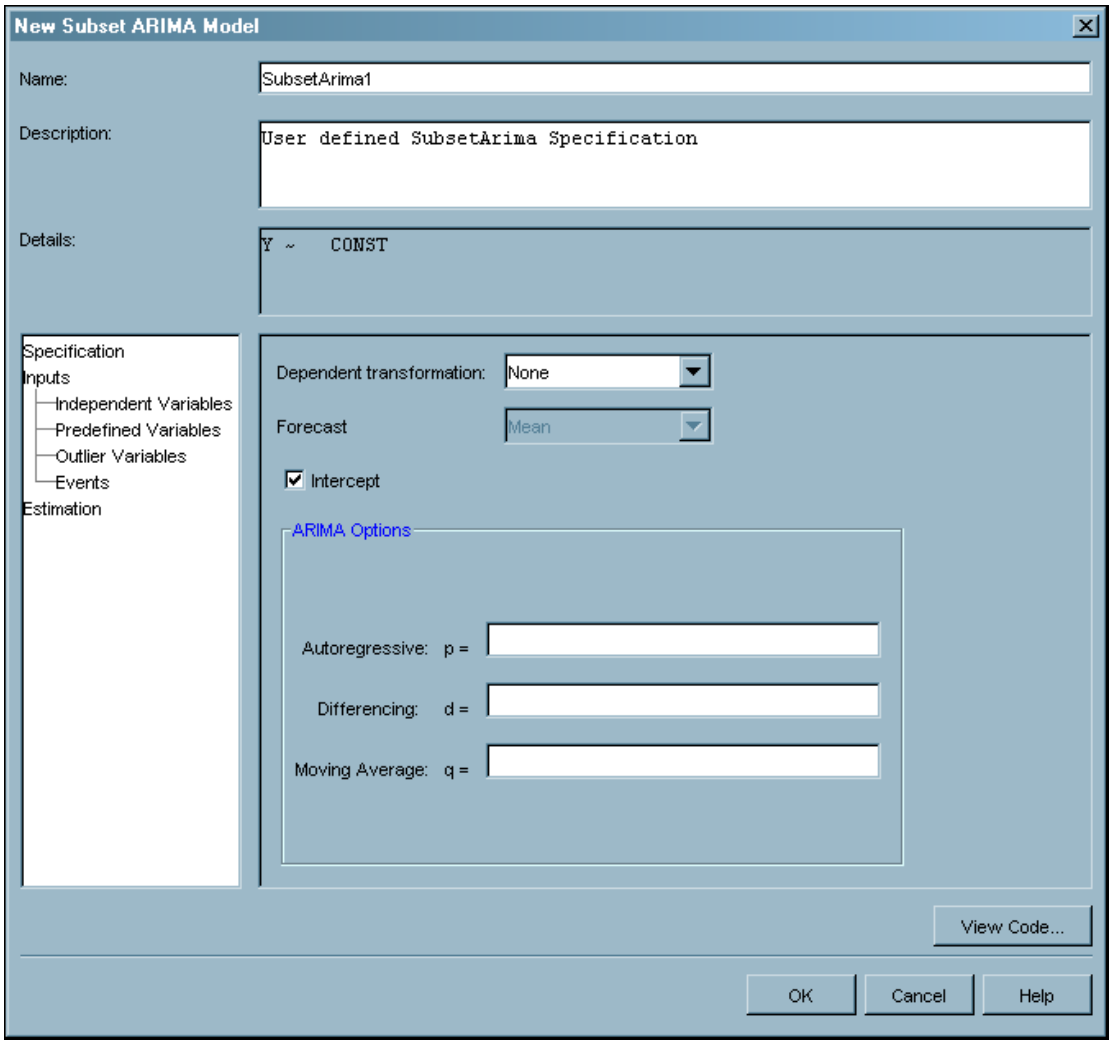

Figure 10.2. New Subset ARIMA Model Dialog Box

Note: The Edit Subset ARIMA Model window 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 207.

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 window. You cannot edit this text.

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

   - specification
   - independent variables
   - predefined variables
   - outlier variables
   - events
   - 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 window, select **Specification** in the selection pane to access the following options:

- The transformation to apply to the dependent series from the **Dependent transformation** drop-down list. You can select from the following transformations:
  - no transformation
  - logarithmic transformation
  - square-root transformation
  - logistic transformation
  - Box-Cox transformation where the parameter number is between -5 and 5. Specify this parameter value in the Box-Cox Transformation window. For more information, see "Specifying a Box-Cox Transformation" on page 210.
- Specify 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 207.

Independent Variables for Subset ARIMA Models

By default, none of the variables that you assign to the Independent variables role are inputs in the model. However, using the Subset ARIMA Model window, you can specify to include these independent variables in the model. You can also specify the input series options for each independent variable. For more information, see "Include Independent Variables in User-Defined Models" on page 211.

Predefined Variables for Subset ARIMA Models

By default, none of the predefined variables from the input data set are inputs in the model. However, using the Subset ARIMA Model window, you can specify to include these predefined variables in the model. For more information, see "Include Predefined Variables in User-Defined Models" on page 211.

Outlier Variables for Subset ARIMA Models

By default, any outliers that are automatically detected during the creation of the model are inputs in the model. However, using the Subset ARIMA Model window, you can specify to delete these outliers from the model. For more information, see "Include Outlier Variables in User-Defined Models" on page 212.

Events for Subset ARIMA Models

By default, none of the events in the project are inputs in the model. However, using the Subset ARIMA Model window, 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 211.

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 214.

Unobserved Components Models

Create an Unobserved Components Model

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

1. Select **Project** → **Project Models**. The Project Models dialog box opens. For more information, see "Overview of Project Models" on page 219.

   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 227.

2. Click the **Custom Models** tab in the Project Models dialog box.

3. On the **Custom Models** tab, click **New**. The New Models dialog box opens.
4. In the New Models dialog box, select **Create an entirely new model**.
5. From the **Model type** drop-down list, select **Unobserved Components** as the model type. The New Unobserved Components Model window opens.

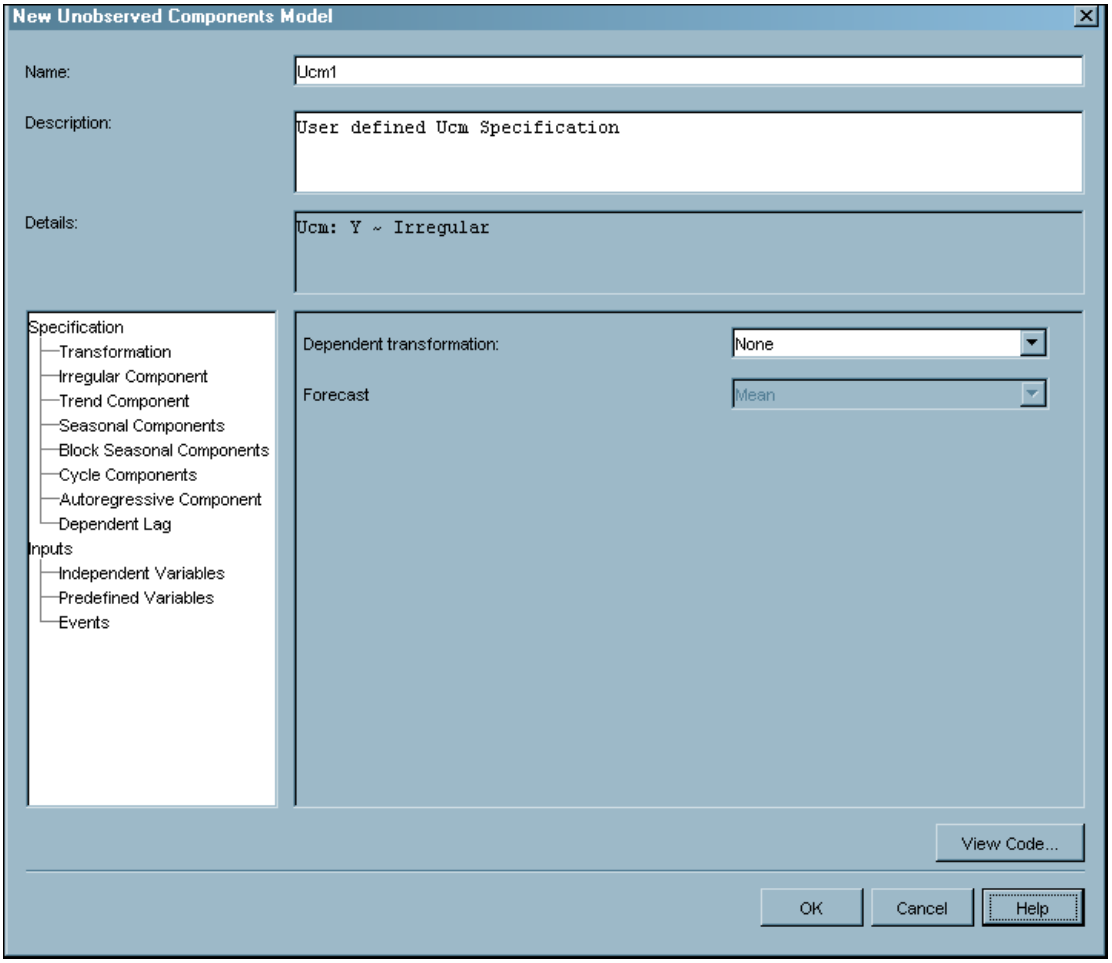

Figure 10.3. New Unobserved Components Model Dialog Box

Note: The Edit Unobserved Components Model window 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 207.

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 window. You cannot edit this text.

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

   - transformation
   - irregular component
   - trend component
   - seasonal components
   - block seasonal components
   - cycle components
   - autoregressive component
   - dependent lag
   - independent variables
   - predefined variables
   - events

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 window, select **Transformation** in the selection pane.
2. From the **Dependent transformation** drop-down list, select the transformation for the model.

   You can select from the following transformations:

   - No transformation. This is the default.
   - Logarithmic transformation.
   - Square-root transformation.
   - Logistic transformation.
   - Box-Cox transformation where the parameter number is between -5 and 5. Specify this parameter value in the Box-Cox Transformation window. For more information, see "Specifying a Box-Cox Transformation" on page 210.

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 window, 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** 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 window, select **Trend Component** in the selection pane.
2. Select the check box for the **Level** or **Slope** component.
3. (Optional) Select the **Define initial variance** check box and specify the initial value for the variance.
4. If you want the variance to be set at the initial value, select the **Do not perform estimation** check box. This option is not available unless you specify a value for the initial variance.

Seasonal Components for Unobserved Components Models

Working with Seasonal Components

The table in the Seasonal Components pane lists each seasonal component and its specified values. For a model, you can specify a maximum of three seasonal components.

From this pane, you can complete the following tasks:

- add a new seasonal component
- edit an existing seasonal component
- delete an existing seasonal component

Create a Seasonal Component

To create a seasonal component, complete the following steps:

1. In the Unobserved Components Model window, 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 60.
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 window, 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 187.
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 window, 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 each component and its specified values. For a model, you can specify a maximum of three block seasonal components.

From this pane, you can complete the following tasks:

- add a new block seasonal component
- edit an existing block seasonal component
- delete a component

Create a Block Seasonal Component

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

1. In the Unobserved Components Model window, 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 window, 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 188.
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 window, 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 each component and its specified values. You can specify up to 50 cycles in a model. By default, the cycle components are estimated from the data. You can optionally create additional cycle components.

From this pane, you can complete the following tasks:

- add a new cycle component
- edit an existing 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 window, select **Cycle Components** in the selection pane.
2. In the Cycle Components pane of the Unobserved Components Model window, 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 window, select **Cycle Components** in the selection pane.
2. In the Cycle Components pane of the Unobserved Components Model window, 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 189.
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 window, select **Cycle Components** in the selection pane.
2. In the Cycle Components pane of the Unobserved Components Model window, 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 window, select **Autoregressive Component** in the selection pane.
2. Select the **Autoregressive component** check box.
3. Optionally, to specify an initial damping value to use during the parameter estimation process, select the **Define initial damping** check box, and then specify a value in the box. This value must be between -1 and 1, including -1 but excluding 1.
4. Optionally, to specify an initial variance to use during the parameter estimation process, select the **Define initial variance** check box, and then specify a value in the box. This value must be greater than or equal to 0.

For the damping and variance values, you can select the **Do not perform estimation (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 window, 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	(5)	(5)
Lag with multiplier	(1 2 3)(2 7)12	(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.

Independent Variables for Unobserved Components Models

By default, none of the variables that you assign to the Independent variables role are inputs in the model. However, using the Unobserved Components Model window, you can specify to include these independent variables in the model. You can also specify the input series options for each independent variable. For more information, see "Include Independent Variables in User-Defined Models" on page 211.

Predefined Variables for Unobserved Components Models

By default, none of the predefined variables from the input data set are inputs in the model. However, using the Unobserved Components Model window, you can specify to include these predefined variables in the model. For more information, see "Include Predefined Variables in User-Defined Models" on page 211.

Events for Unobserved Components Models

By default, none of the events in the project are inputs in the model. However, using the Unobserved Components Model window, 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 211.

Exponential Smoothing Models

Overview of Exponential Smoothing Models

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

Using SAS Forecast Studio, you can create an exponential smoothing model. You can choose from 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** → **Project Models**. The Project Models dialog box opens. For more information, see "Overview of Project Models" on page 219.

   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 227.
2. Click the **Custom Models** tab in the Project Models dialog box.
3. On the **Custom Models** tab, click **New**. The New Models dialog box opens.
4. In the New Models dialog box, select **Create an entirely new model**.
5. From the **Model type** drop-down list, select **Exponential smoothing** as the model type. The New Exponential Smoothing Model window opens.

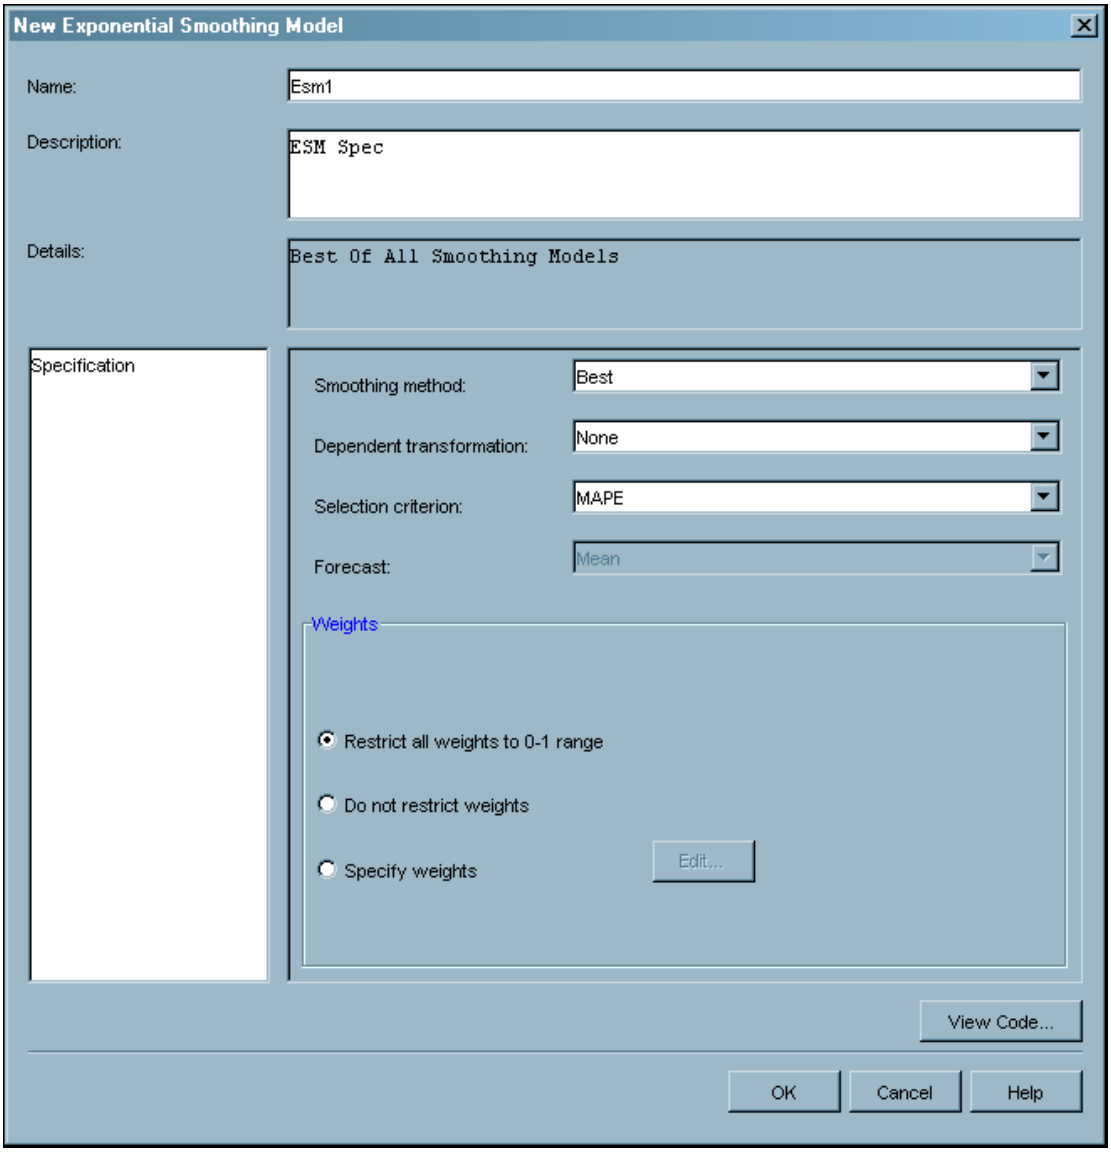

Figure 10.4. New Exponential Smoothing Model Dialog Box

Note: The Edit Exponential Smoothing Model window 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 207.

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 window. You cannot edit this text.

8. From the **Smoothing method** drop-down list, select the forecasting method to be used to forecast the time series. The following forecasting models are available:

   - **Best** - Best Candidate Smoothing model
   - **Bests** - Best Candidate Seasonal Smoothing model
   - **Seasonal** - Seasonal Smoothing model
   - **Winters** - Winters (Multiplicative) model
   - **Additive Winters** - Winters (Additive) model
   - **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

   Note: The **Winters** option is not available if the series that you selected contains negative values.

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

9. From the **Dependent 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. This is the default.
   - Logarithmic transformation.
   - Square-root transformation.
   - Logistic transformation.
   - Box-Cox transformation where the parameter number is between -5 and 5. Specify this parameter value in the Box-Cox Transformation dialog box. For more information, see "Specifying a Box-Cox Transformation" on page 210.
   - 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.

10. Select the selection criterion to use to select from several candidate models. The default selection criterion is the project default.

    Note: This option is available only when either of the following is true:

    - In the **Smoothing method** drop-down list, the smoothing method is **Best**, **BestN**, or **Bests**.
    - In the **Dependent transformation** drop-down list, you selected **Auto**, which means that the transformation for the dependent series is automatically chosen based on the model selection criteria.

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 edit the smoothing weights. For more information, see "Smoothing Weights" on page 214.

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** → **Project Models**. The Project Models dialog box opens. For more information, see "Overview of Project Models" on page 219.

   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 227.
2. Click the **Custom Models** tab in the Project Models dialog box.
3. On the **Custom Models** tab, click **New**. The New Models dialog box opens.
4. In the New Models dialog box, select **Create an entirely new model**.
5. From the **Model type** drop-down list, select **Intermittent demand** as the model type. The New Intermittent Demand Model window opens.

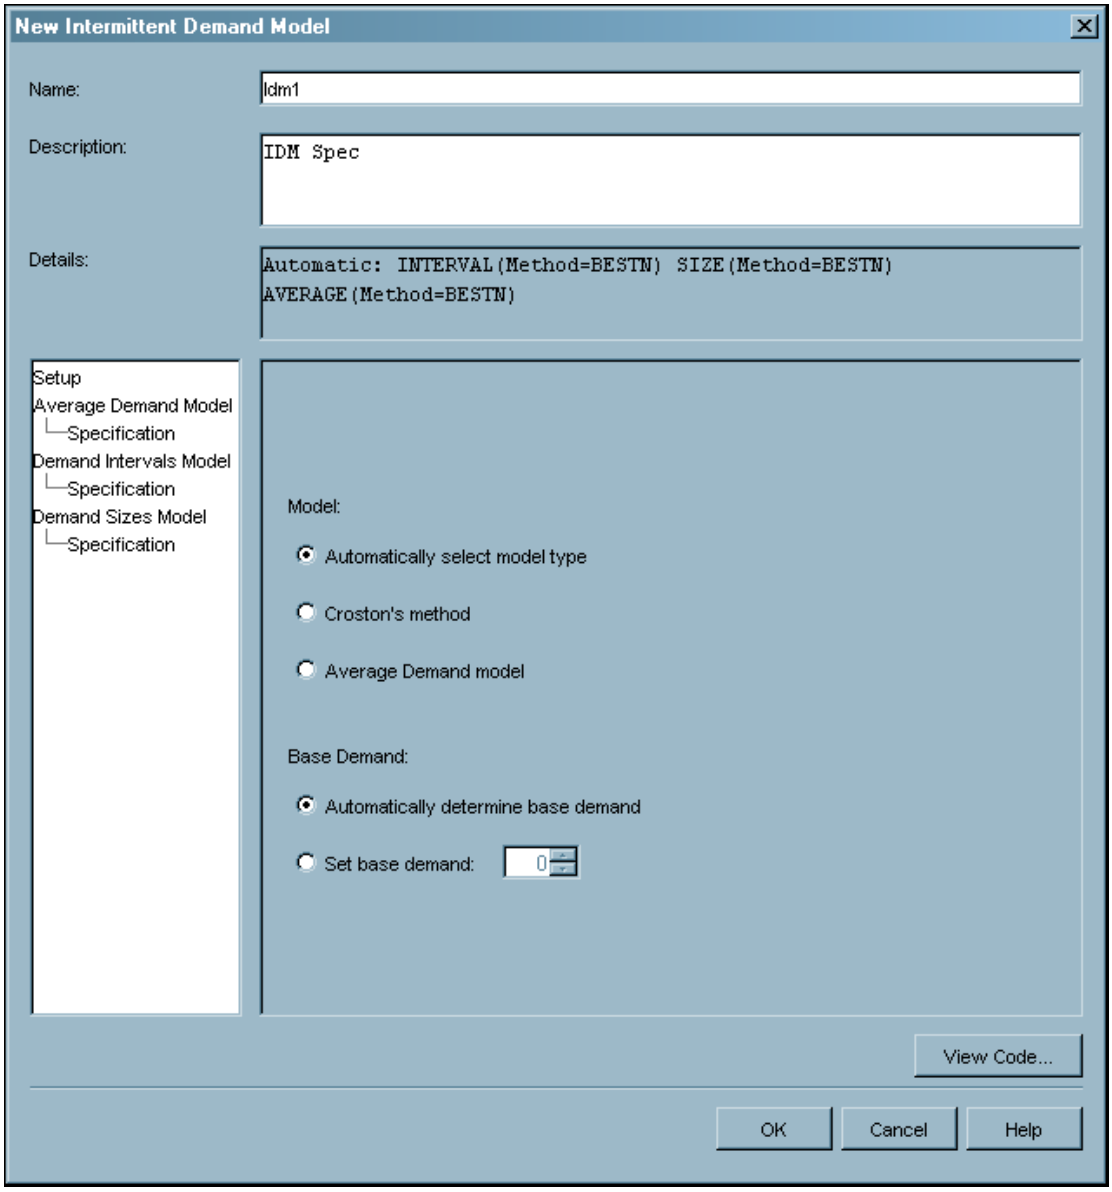

Figure 10.5. New Intermittent Demand Model Dialog Box

Note: The Edit Intermittent Demand Model window 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 207.

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 window. You cannot edit this text.

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

   - setup
   - specification

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

Setup for Intermittent Demand Models

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

- The model type. You can have SAS Forecast Studio automatically determine model type, or you can choose from the following model types:
  - **Croston's method** - models and forecasts each component independently and then combines the two forecasts.
  - **Average Demand model** - forecasts the intermittent time series by treating the average demand series as a time series based on the demand index. Optimal smoothing parameters can be estimated and predictions for average demand can be computed using nonseasonal exponential smoothing methods as well as their transformed versions.
- The base demand. The demand for the intermittent time series typically depends on a base value. You can have SAS Forecast Studio automatically determine the base demand based on the characteristics of the time series, or set the base demand yourself. The default base value is 0, but it can be any constant value.

Specification Options for Intermittent Demand Models

For an Intermittent Demand model, you can create an average demand model, a demand intervals model, and a demand sizes model. The demand intervals and demand sizes models use Croston's method.

For the model that you want to create, select **Specification** in the selection pane to access the following options:

- The forecasting method to be used to forecast the time series. The following forecasting models are available:

  - **Simple** - Simple (Single) Exponential Smoothing model (which is the default)
  - **Double** - Double (Brown) Exponential Smoothing model
  - **Linear** - Linear (Holt) Exponential Smoothing model
  - **Damp Trend** - Damped Trend Exponential Smoothing model
  - **BestN** - Best Candidate Nonseasonal Smoothing model

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

- The dependent transformation to be applied to the time series. The following transformation options are available:

  - 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 Transformation dialog box. For more information, see "Specifying a Box-Cox Transformation" on page 210.
  - An automatically chosen transformation based on the model selection criteria.

- The selection criterion to use to select from several candidate models. The default selection criterion is the project default. This option is available only when the method is **BestN** (Best Candidate Nonseasonal Smoothing model).
- 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 214.

Multiple Regression Models

Create a Multiple Regression Model

To create a multiple regression model, complete the following steps:

1. Select **Project** → **Project Models**. The Project Models dialog box opens. For more information, see "Overview of Project Models" on page 219.

   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 227.

2. Click the **Custom Models** tab in the Project Models dialog box.
3. On the **Custom Models** tab, click **New**. The New Models dialog box opens.
4. In the New Models dialog box, select **Create an entirely new model**.
5. From the **Model type** drop-down list, select **Multiple Regression** as the model type. The New Multiple Regression Model window opens.

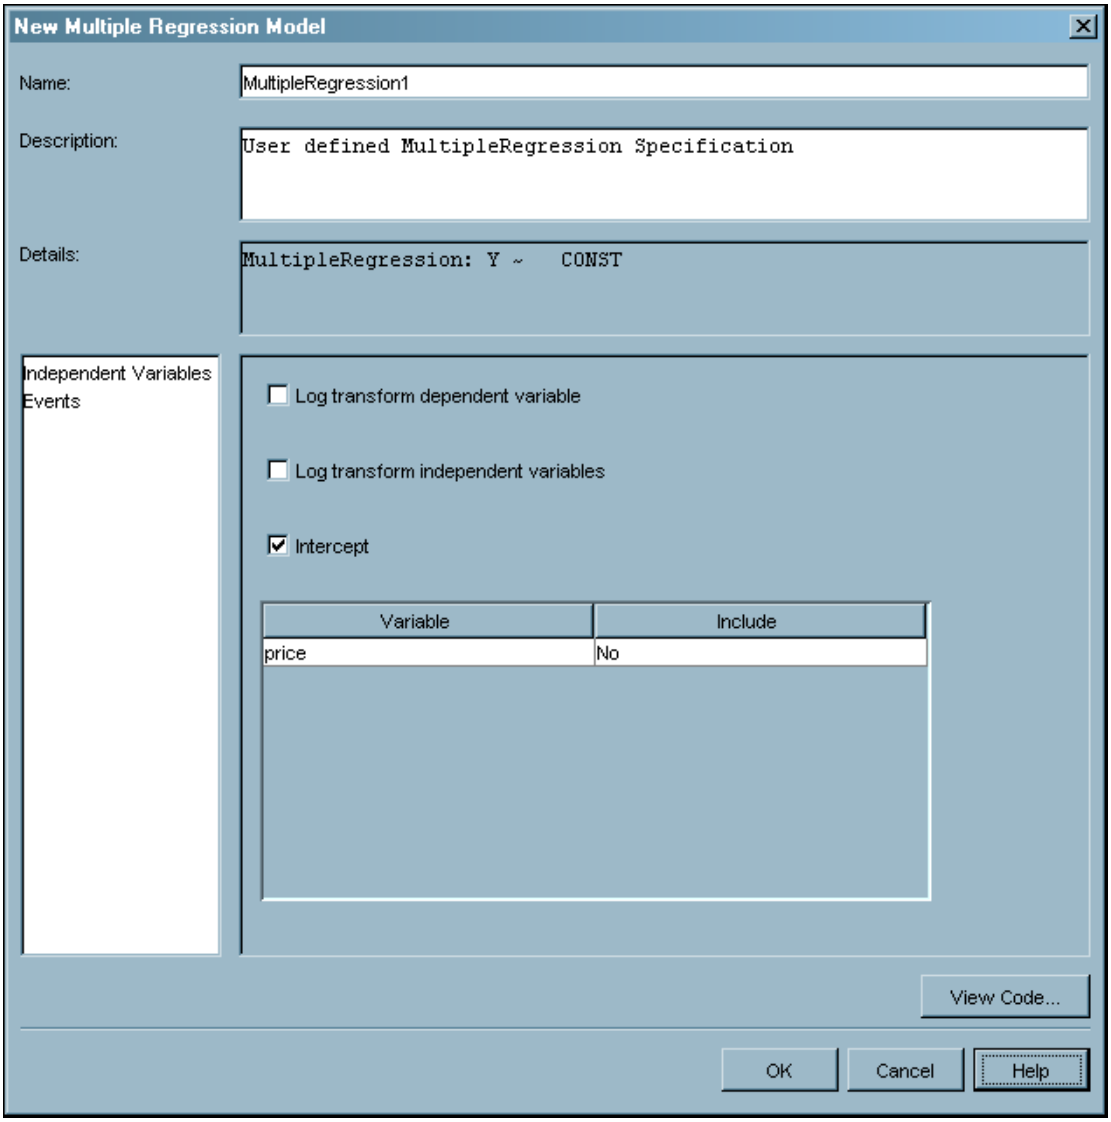

Figure 10.6. New Multiple Regression Model Dialog Box

Note: The Edit Multiple Regression Model window 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 207.

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 window. You cannot edit this text.

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

   - independent variables
   - events

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

Independent Variables for Multiple Regression Models

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

1. In the Multiple Regression Model window, select **Independent Variables** in the selection pane.
2. Select the **Log transform the dependent variable** check box to log transform the dependent variable.
3. Select the **Log transform the independent variables** check box to log transform the independent variables.
4. Select the **Intercept** check box to include the intercept in the model. By default, the intercept is included.
5. Use the drop-down list in the Include column to specify which independent variables should be included in the model. By default, all of the independent variables are ignored.

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 window, 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 211.

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** → **Project Models**. The Project Models dialog box opens. For more information, see "Overview of Project Models" on page 219.

   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 227.
2. Click the **Custom Models** tab in the Project Models dialog box.
3. On the **Custom Models** tab, click **New**. The New Models dialog box opens.
4. In the New Models dialog box, select **Create an entirely new model**.
5. From the **Model type** drop-down list, select **Curve Fitting** as the model type. The New Curve Fitting Model window opens.

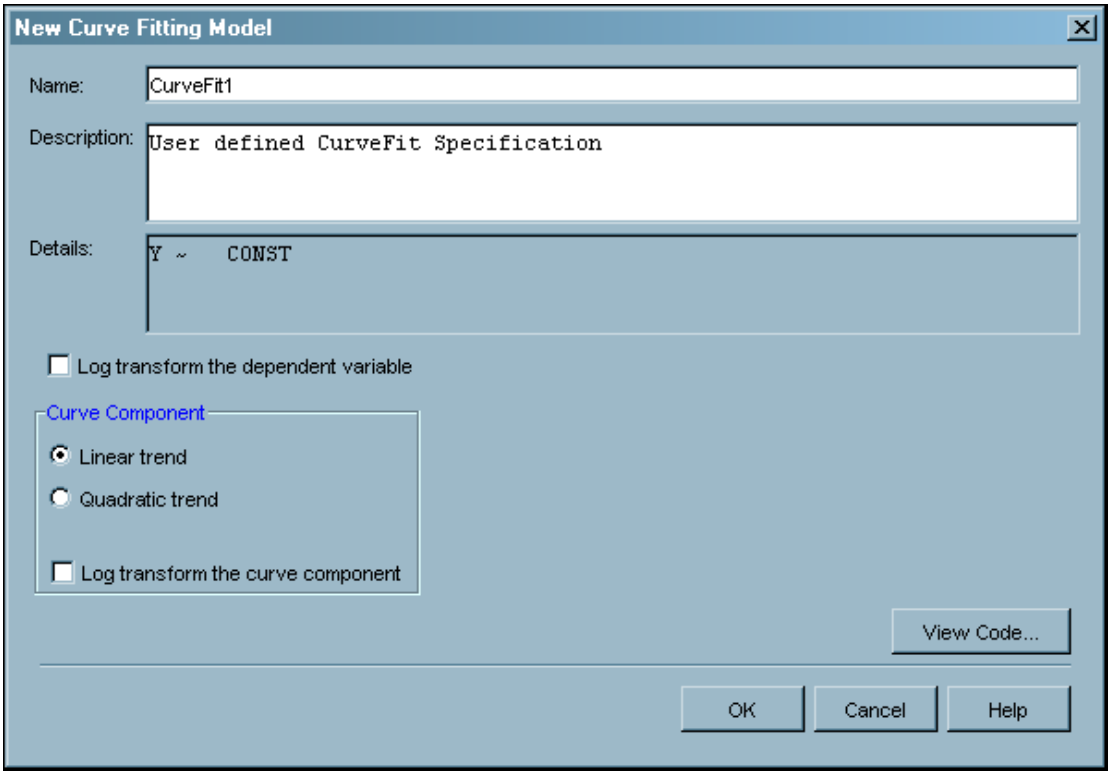

Figure 10.7. New Curve Fitting Model Dialog Box

Note: The Edit Curve Fitting Model window 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 207.

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 window. 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**.

Moving Average Models

Overview of Moving Average Models

Using SAS Forecast Studio, you can create a moving average model that you choose. The formula for the Moving Average model with width *k* is the following:

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

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

Create a Moving Average Model

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

1. Select **Project** → **Project Models**. The Project Models dialog box opens. For more information, see "Overview of Project Models" on page 219.

   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 227.

2. Click the **Custom Models** tab in the Project Models dialog box.
3. On the **Custom Models** tab, click **New**. The New Models dialog box opens.
4. In the New Models dialog box, select **Create an entirely new model**.
5. From the **Model type** drop-down list, select **Moving Average** as the model type. The New Moving Average Model window opens.

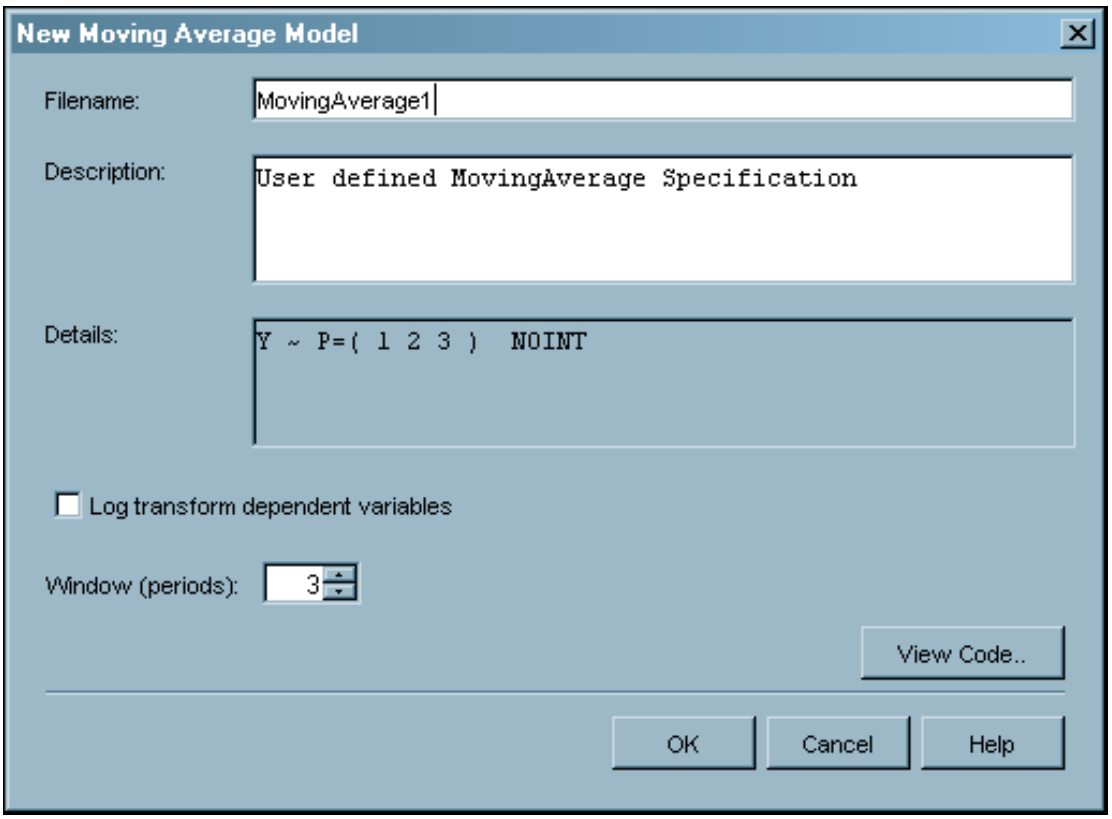

Figure 10.8. New Moving Average Model Dialog Box

Note: The Edit Moving Average Model window 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 207.

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 window. You cannot edit this text.

8. Select the **Log transform the dependent variable** 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**.

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** → **Project Models**. The Project Models dialog box opens. For more information, see "Overview of Project Models" on page 219.

   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 227.
2. Click the **Custom Models** tab in the Project Models dialog box.
3. On the **Custom Models** tab, click **New**. The New Models dialog box opens.
4. In the New Models dialog box, select **Create an entirely new model**.
5. From the **Model type** drop-down list, select **Random walk** as the model type. The New Random Walk Model window opens.

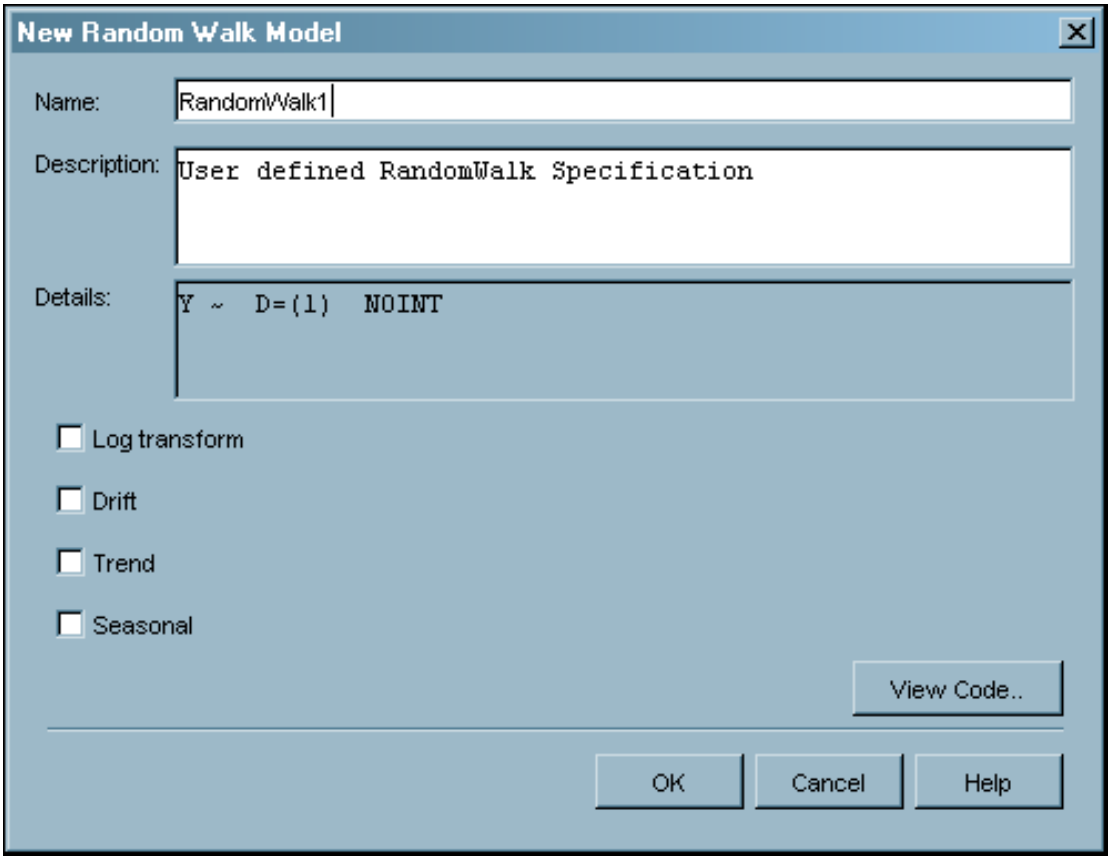

Figure 10.9. New Random Walk Model Dialog Box

Note: The Edit Random Walk Model window 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 207.

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 window. You cannot edit this text.

8. Select the **Log transform** check box to log transform the dependent variable. When you select this option, the model uses the data for the log transformed series.

9. Select the **Drift** check box to include a drift term in the model.

10. Select the **Trend** check box to include a trend term in the model.

11. Select the **Seasonal** check box to include 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 name of the system-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 221.

Autoregressive, Differencing, and Moving Average Orders

What Models Use Autogressive, 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 PROC HPFARIMASPEC 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= option	Result in Details pane
Autoregressive model with degrees from 1 to an integer value	5	P = (1 2 3 4 5)
Autoregressive model with parameters at specified lags	(5)	P = (5)
Autoregressive model with a seasonal component with "s" as a placeholder	(1 2 3)(2 7)s	P = (1 2 3)(2 7)s
Autoregressive model with a seasonal component of 12	(1 2 3)(2 7)12	P = (1 2 3)(24 84)
Autoregressive model with three factors	(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= option	Result in Details pane
Single differencing order	5	Incorrect syntax; no order specified
Single differencing order	(5)	D = (5)
Multiple differencing orders	(1 1 3 5)	D = (1 1 3 5)
Differencing order with a seasonal component	(1 s s)	D = (1 s s)

Moving Average Orders

The **Moving average (q)** option specifies the moving-average part of the model. The manner of specification of the moving-average part is identical to the specification of the autoregressive part described in the p= option.

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= option	Result in Details pane
Moving average model with degrees from 1 to an integer value	5	Q = (1 2 3 4 5)
Moving average model with parameters at specified lags	(5)	Q = (5)
Moving average model with a seasonal component with "s" as a placeholder	(1 2 3)(2 7)s	Q = (1 2 3)(2 7)s
Moving average model with a seasonal component of 12	(1 2 3)(2 7)12	Q=(1 2 3)(24 84)
Moving average model with three factors	(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= option	Value of q= option	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)

Specifying a Box-Cox Transformation

For the ARIMA, Subset (factored) ARIMA, Exponential Smoothing, Intermittent Demand, and Unobserved Components models, you can specify Box-Cox as the transformation to apply to the dependent variable.

When you select the Box-Cox as the transformation to apply to the dependent series for the model, the Box-Cox Transformation dialog box opens. In the **Set parameter** box, specify the Box-Cox transformation where the parameter number is between -5 and 5. The default value is 0. Click **OK** to continue creating or editing the model.

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 by using the Variables dialog box. For more information, see "Modify Variable Assignments" on page 66.

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 window, select **Independent Variables**.
2. For each independent variable that you want to include in the model, change the value in the **Include** column to **Yes**.

You can also specify the input series options for each independent variable. To specify these series options, select the variable from the table and click **Options**. The Input Series Options dialog box opens. For more information, see "Specifying the Input Series Options" on page 212.

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 window, select **Predefined Variables**.
2. Using the options in the **Transform** column, specify whether the variable will be transformed. The options that are available depend on the variable, and some variables cannot be transformed.
3. For each predefined variable that you want to include in the model, change the value of the **Include** column to **Yes**.

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 window, select **Events**.
2. For each event variable that you want to include in the model, change the value of the **Include** column to **Yes**.

On the Events pane, you can also perform the following tasks:

- To create a new event, click **New** on the Events pane. The New Event dialog box opens. After the event is created, it is added to the list of events, and the event is automatically included in the model. The event that you created also appears in the list of events for the project, but it is not used at the project level. For more information, see "Create an Event" on page 97.
- To remove an event from the model, set the value of the **Include** column to **No**.
- To edit an event, select the event that you want to edit and click **Edit** on the Events pane. For more information, see "Edit an Event" on page 101.
- To specify the input series options for an event, select the event from the list and click **Options**. The Input Series Options dialog box opens. For more information, see "Specifying the Input Series Options" on page 212.

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 window, 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 panel in the ARIMA Model or Subset ARIMA Model window the outlier that you deleted does not appear.

Specifying the Input Series Options

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

From this dialog box, the following options are available:

- **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 Transformation dialog box. For more information, see "Specifying a Box-Cox Transformation" on page 210.

  Note: The transformation is applicable only if the input variable is positive.
- **Specify lagging order** - specifies the delay (or lag) for the input series.
- **Specify differencing orders** - specifies the differencing orders for the input series. For more information about the syntax that you must use, see "Autoregressive, Differencing, and Moving Average Orders" on page 207.

  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 window or the Subset ARIMA Model window.
- 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.

  If you specify a dynamic regression for an ARIMA or subset ARIMA model, then you can specify the numerator (NUM) and denominator (DEN) polynomial of the transfer function.

  The syntax to use for these polynomials is the following:

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

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

  For information about how to specify the polynomial order, see the autoregressive (p) option in "Autoregressive Orders" on page 208.

Once 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 window, 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

In the Unobserved Components and Intermittent Demand models, you can specify the smoothing weights to use.

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

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

Note: Only the parameters that are relevant for a model are available.

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

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

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

Chapter 11
Working with Models

Chapter Contents

Chapter 11
Working with Models

Managing Models

Overview of Project Models

To view all of the models that are available for the current project, select **Project** → **Project Models**.

The Project Models dialog box shows all of the user-defined and system-generated models. These models are stored in separate model repositories. For more information, see "How Models Are Stored" on page 220.

Note: The list of models shown in the Project Models 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 222.

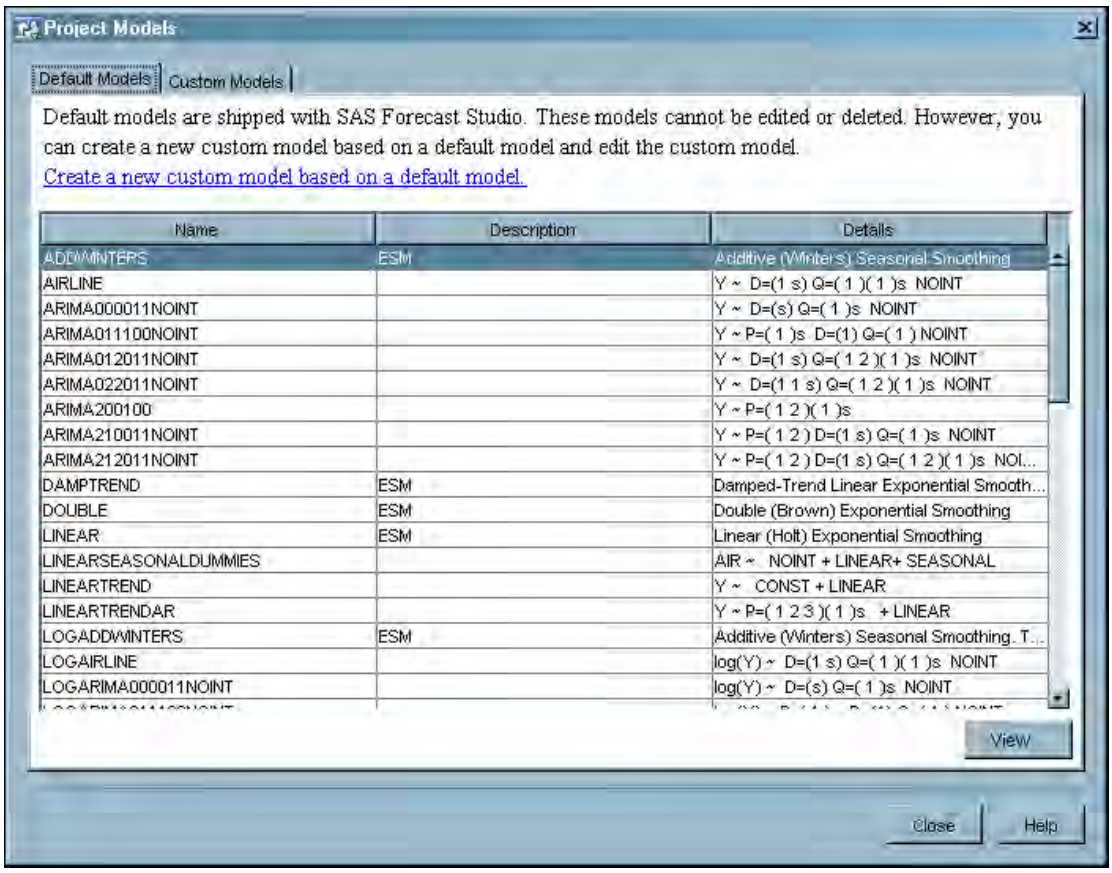

Figure 11.1. Project Models Dialog Box

The Project Models dialog box contains the following tabs:

Default Models
lists the default 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 221 and Appendix C, "Default Model Selection Lists."

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 222.

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 207.

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** → **Project Models**. The Project Models dialog box opens.
2. Select the 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."
3. When you have finished reviewing the options in the model specification dialog box, click **OK** to close this dialog box.

Note: You can also view the model specification for a default model from the model selection list. To view the model specification for a default model, double-click the model from the model selection list. In the dialog box that opens, click **View** to view a read-only version of the model specification. To edit the model specification, you must click **Copy**. For more information about copying models, see "Copy Models" on page 229.

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** → **Project Models**. The Project Models 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.

   For more information about the options that are available for each model, see Chapter 10, "Creating User-Defined Models."
5. Edit the model options and click **OK**. SAS Forecast Studio saves the model and displays it on the **Custom Models** tab in the Project Models dialog box.

Working with Custom Models

From the **Custom Models** tab in the Project Models 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 an Entirely New Model" on page 175.
- 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 228.
- 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 228.
- 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 part of 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.

The models in this list are the ones that SAS Forecast Studio has selected as the best models to use to generate forecasts. The best model is the forecast model, and it is this result that is displayed in the forecast plot and data table in the Forecasting View. In the model selection list, the name of the forecast model is in bold and has **Forecast Model** prepended to the name.

The following statistics are displayed for the forecast model:

- The performance statistic for the model. This statistic shows how well the model is predicting the future. This statistic is available only when you specify a value greater than zero for the **Out of Sample range** option in the Options dialog box. When available, this statistic is labeled Performance *Statistic of Fit*. In Figure 11.2, the performance statistic is Performance MAPE=0.21.
- The fit statistic that you selected for the project. This statistic shows how well the model is fitting the data. You can change the fit statistic by using the **Selection Criterion** drop-down list in the Options dialog box. This statistic is labeled Fit Range *Statistic of Fit*. In Figure 11.2, the fit statistic is Fit Range MAPE=1.31.

- The selection statistic that was used to choose the forecast model. The selection statistic appears in the last column of the model list and is labeled Selection *Statistic of Fit*. In Figure 11.2, the fit statistic is Selection MAPE=1.31.

For more information about statistics of fit, see Appendix B, "Statistics of Fit."

Note: If SAS Forecast Studio fails to fit a model, then no statistics of fit are displayed for that model. Instead, the last column says "Failed."

The following is an example of a model selection list:

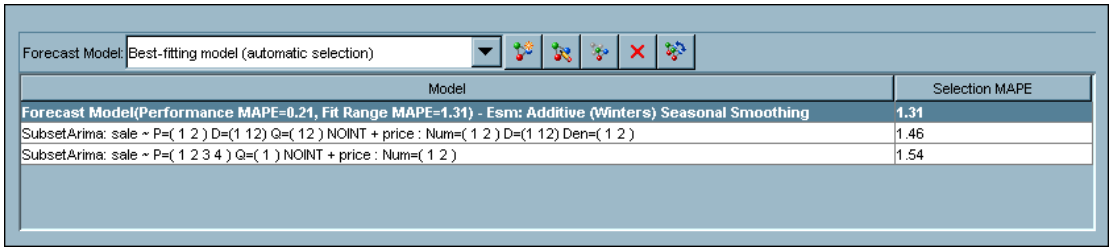

Figure 11.2. Model Selection List in Model View

Note: The status of a model can change depending on the other models in the model selection list. 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 does not fit models that are inappropriate for the data. However, if none of the models fit, then SAS Forecast Studio attempts to force a model to fit.

The following example explains when the status of an intermittent demand model can change. When you initially create a project, you only have one model in the model selection list, so SAS Forecast Studio selects this intermittent demand model as the forecast (or best) model. You add an ARIMA model to the model selection list, and when SAS Forecast Studio rediagnoses the project, the intermittent demand model now fails.

In this example, the intermittent demand model fails because SAS Forecast Studio has determined that the ARIMA model is a better fit. SAS Forecast Studio initially forced the intermittent demand model to fit because it was the only model in the model selection list. Now that an ARIMA model has been added, the intermittent demand model fails because after running the series of diagnostics, SAS Forecast Studio determines that it is inappropriate for the data.

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
- add a model to a series

- edit a model
- copy a model
- delete a model
- compare models

Note: 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 Options dialog box, then the seasonality test that is used to recreate 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 "Specify the Model Selection Options" on page 120.

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:

   - **Autocorrelation Function** - displays the autocorrelation function and standardized autocorrelation function of the prediction errors.
   - **Model Components** - displays plots of each of the components that are used by the model.
   - **Forecast Only** - displays the forecast in the forecast horizon only.
   - **Forecast** - displays a plot of the model and the forecast that covers the time range of the actual data and the forecast horizon.
   - **Inverse Autocorrelation Function** - displays the inverse autocorrelation function and standardized inverse autocorrelation function of the prediction errors.
   - **Partial Autocorrelation Function** - displays the partial autocorrelation function and standardized partial autocorrelation function of the prediction errors.
   - **Model** - displays a plot of the model with predictions and confidence limits that cover the time range of the actual data.
   - **Prediction Errors** - displays a needle plot and histogram of the prediction errors.

- **White Noise Probabilities** - displays the white noise probability plot and log white noise probability plot of the prediction errors.

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

Select the Forecast Model

You can select the forecast model (also called the baseline model) from the model selection list in the Model View. 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 **Forecast 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 a specific model as the forecast model. When you reforecast the project, SAS Forecast Studio keeps the forecast model that you selected, even if this is not the best-fitting model.

Note: If you select an automatically generated model as the forecast model and then reforecast the project to recreate the automatically generated models, then the selection in the **Forecast Model** drop-down list reverts to the **Best-fitting model (automatic selection)** option. For more information, see "Reforecast the Data" on page 122.

To select a specific model as the forecast model, complete the following steps:

1. Select **View** → **Model View** to open the Model View.
2. From the **Forecast 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.

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 135.

Add a Model to a Series

To add a model to a series, complete the following steps:

1. You can add a model in either of the following ways:

   - To add a model to an individual series, click  in the model selection list in the Model View.
   - To add a model to all series in the project, select **Project** → **Add Model to All Series**. The Add Model dialog box opens.

   The Add Model dialog box opens.
2. Select how to create the new model that you are adding to the series. You can choose from the following options:

   - **Select an existing model** - enables you to select the default or custom model to add to all of the series. For more information, see "Overview of Project Models" on page 219.
   - **Create a new model** - enables you to create an entirely new model. For more information, see "Create an Entirely New Model" on page 175.
3. Click **OK**.

If you selected **Project** → **Add Model to All Series**, SAS Forecast Studio adds the model to all of the series in the project. If you are working in the model selection list, then SAS Forecast Studio adds the model to the current series.

Import Models

You might want to use a model that is external to SAS Forecast Studio. You can do this by importing the model selection list that contains the model that you want to use.

To import a model, complete the following steps:

1. Select **Project** → **Options**. The Options dialog box opens.
2. In the selection pane, select **Model Generation**.
3. Select the **Fit 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 "Time Series Forecasting Models" on page 254.

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 Options dialog box to save your changes.

SAS Forecast Studio reforecasts the data and includes the imported models.

Edit a Model

About Editing Models

SAS Forecast Studio provides a variety of ways to edit a model. You can edit a model from the model selection list in the Model View or from the list of project models.

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

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

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

Edit a Model from the Model View

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

1. Select **View** → **Model View** to open the Model View in the main SAS Forecast Studio workspace. The model selection list appears at the top of this view. For more information, see "Using the Model Selection List" on page 222.

2. Select the model that you want to edit and click  . The model specification dialog box for that model opens.

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 221.
- models that SAS Forecast Studio created using the HPFDIAGNOSE procedure. You select the models that SAS Forecast Studio generates from the Model Generation pane in the Options dialog box. For more information, see "Specify the Model Generation Options" on page 118.
- imported models. For more information, see "Import Models" on page 227.

Edit a Model from the List of Project Models

To edit a model from the list of project models, complete the following steps:

1. Select the series in the **Hierarchy** tab or **Table** tab and select **Project** → **Project Models**. The Project Models dialog box opens. For more information about this dialog box, see "Overview of Project Models" on page 219.
2. Select the model that you want to edit from the list and click **Edit**. The model specification dialog box for that model opens.

Additional Information

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 176
- "Create a Curve Fitting Model" on page 202
- "Create an Exponential Smoothing Model" on page 193
- "Create an Intermittent Demand Model" on page 196
- "Create a Moving Average Model" on page 203
- "Create a Multiple Regression Model" on page 199
- "Create a Random Walk Model" on page 205
- "Create a Subset (Factored) ARIMA Model" on page 180
- "Create an Unobserved Components Model" on page 183

Copy Models

To copy a model, select the model name in the model selection list and click  . The new model is added to the list with the name *model-name*Copy*n* where *model-name* is the original model name and *n* is the lowest integer value available.

You can then edit the model from the model specification dialog box that opens. For more information, see "Edit a Model" on page 228.

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 is removed from 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. Through this comparison, you can determine which model that you want to use.

From the model selection list in the Model View, you can compare all of the models for a series by clicking  . The Compare Models dialog box opens.

If you are not working in the Model View, then you can compare models for the currently selected series by completing the following steps:

1. In the **Hierarchy** or **Table** tab, select the series that you want to use.
2. Select **Series** → **Compare Models**. The Compare Models dialog box opens.

In the Compare Models dialog box, each model is in a row in the table and the statistics of fit are in columns. For each statistic of fit, the cell for the best model is shaded in green. For more information about the available statistics, see Appendix B, "Statistics of Fit."

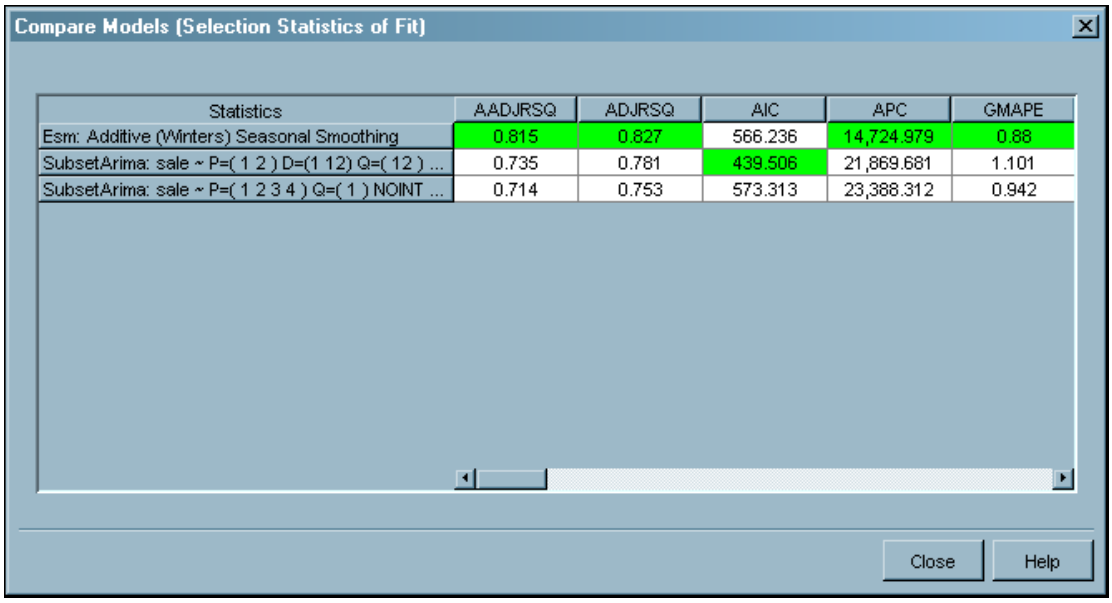

Figure 11.3. Compare Models Dialog Box

Chapter 12
Generating Reports

Chapter Contents

Chapter 12
Generating Reports

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 D, "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 236.

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 66.

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 the "Additional Administration Tasks" in the *SAS Forecast Server Administrator's Guide*.

Part 3
Appendixes

Contents

Appendix A
SAS Forecast Studio Variables

Appendix Contents

Appendix A
SAS Forecast Studio Variables

Reserved Variable Names

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

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

Variable Name	Description
_*VariableName*	Any variable name that begins with an underscore
ACTUAL	Dependent series value
COMP	Name of the component
COMPONENT	Model component
CROSS	Cross-variable name
DSVAR	Data set variable mapping
EST	Parameter estimate
FACTOR	Model factor
LABEL	Parameter or statistic label
LAG	Lag of input
LOWER	Lower confidence limit
MODE	Mode of decomposition
MODEL	Name of model
MODELVAR	Model variable mapping
NAME	Variable name
PARM	Parameter name
PREDICT	Component forecast
PVALUE	Parameter estimate *p*-value
SEASON	Seasonal index
SELECT	Name of selection list
SHIFT	Shift
STAT	Statistic name
STATUS	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

Variable Name	Description
TRANSFORM	Transformation applied
UPPER	Upper confidence limit
AADJRSE	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
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
CCFLPROB	Cross-correlation log probabilities
CCFPROB	Cross-correlation probabilities
CCFSTD	Cross-correlation standard errors
CCOV	Cross-covariances
CSS	Corrected sum of squares
ERROR	Prediction errors
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
MAXERR	Maximum error
MAXIMUM	Maximum value
MAXPE	Maximum percent error
ME	Mean error
MEAN	Mean value
MEDIAN	Median value
MINERR	Minimum error
MINIMUM	Minimum value
MINPE	Minimum percent error
MPE	Mean percent error
MSE	Mean square error

Variable Name	Description
N	Number of nonmissing observations or number of variance products
NAME	Variable name
NMISS	Number of missing observations
NOBS	Number of observations
ORIGINAL	Original series index
PACF	Partial autocorrelations
PACF2STD	PACF beyond two standard errors
PACFLPROB	Partial autocorrelation log probabilities
PACFNORM	Partial normalized autocorrelations
PACFPROB	Partial autocorrelation probabilities
PACFSTD	Partial autocorrelations standard errors
PCSA	Percent change seasonal adjusted component
PERIOD*t*	Decomposition component value or trend statistic for time period *t*
PREDICT	Predicted values
RANGE	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*
SSE	Sum of squares error
STD	Prediction standard errors
STDDEV	Standard deviation
SUM	Summation value
TC	Trend component
TCC	Trend-cycle component
TCS	Trend-cycle-seasonal component
UMSE	Unbiased mean square error
URMSE	Unbiased root mean square error
UPPER	Upper confidence limits
USS	Uncorrected sum of squares
WN	White noise test statistics
WNLPROB	White noise test log probabilities
WNPROB	White noise test probabilities

Appendix B
Statistics of Fit

Appendix Contents

Appendix B
Statistics of Fit

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 are 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 Percent Error (GMAPE)
: The geometric mean percent error.

Geometric Mean Predictive Percent Error (GMAPPE)
: The geometric mean absolute predictive percent prediction error.

Geometric Mean Symmetric Percent Error (GMASPE)
: The geometric mean of the absolute symmetric percent errors.

Geometric Mean Relative Absolute Error (GMRAE)
: The geometric mean of the relative absolute errors.

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.

Median Percent Error (MDAPE)
: The median of the percent errors.

Median Predictive Percent Error (MDAPPE)
: The median of the predictive percent errors.

Median Symmetric Percent Error (MDASPE)
: The median of the symmetric percent errors.

Median Relative Absolute Error (MDRAE)
: The median of the relative absolute errors.

Mean Error (ME)
: The mean prediction error, $\frac{1}{n}\sum_{t=1}^{n}(y_t - \hat{y}_t)$.

Mean Absolute Error (MAE)
: The mean absolute prediction error, $\frac{1}{n}\sum_{t=1}^{n}|y_t - \hat{y}_t|$.

Mean Absolute Percent Error (MAPE)
: The mean of the absolute percent errors.

Symmetric Mean Absolute Predictive Percent Error (MAPPE)
: The mean of the absolute symmetric predictive percent error.

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.

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.

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.

Total Sum of Squares (SST)
: The total sum of squares for the series, uncorrected for the mean: $\sum_{t=1}^{n} y_t^2$.

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.

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 Details" section of *SAS High-Performance Forecasting User's Guide*.

Appendix C
Default Model Selection Lists

Appendix Contents

Appendix C

Default Model Selection Lists

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** → **Project Models**.

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 D
Sample Reports in SAS Forecast Studio

Appendix Contents

Appendix D

Sample Reports in SAS Forecast Studio

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 260.

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.
Node Macro Definition Export Table	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 for Getting Started Reports** 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 66.

Name of Report	Description
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 66.
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.

Name of Report	Description
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
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

Name of Report	Description
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
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

Name of Report	Description
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
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

Name of Report	Description
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
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.
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.
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.
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.

Name of Report	Description
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.
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.

Name of Report	Description
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 66.

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.
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.

statistic of fit
: a statistical value that is used to evaluate how well a forecasting model performs 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

N

O

P

R

S

T

U

V

W

Your Turn

If you have comments or suggestions about *SAS® Forecast Studio 1.4: User's Guide,* please send them to us on a photocopy of this page or send us electronic mail.

For comments about this book, please return the photocopy to

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

For suggestions about the software, please return the photocopy to

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