THE POWER TO KNOW.

SAS® Forecast Studio 3.1

User's Guide

SAS® Documentation

The correct bibliographic citation for this manual is as follows: SAS Institute Inc. 2009. *SAS® Forecast Studio 3.1: User's Guide*. Cary, NC: SAS Institute Inc.

SAS® Forecast Studio 3.1: User's Guide

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

1st electronic book, March 2009

Contents

Part I

Introduction to SAS Forecast Studio

Chapter 1

What's New in SAS Forecast Server 3.1

Contents

Overview

SAS Forecast Server 3.1 has the following changes and enhancements:

- enhanced project management
- updates to project code
- new time intervals
- new scenario analysis
- new options for the system-generated ARIMA model
- report enhancements
- enhancements to the user interface

Enhanced Project Management

Environments for SAS Forecast Server

An environment in SAS Forecast Server is a virtual container of run-time settings for the client sessions. Environments specify what resources are available to users and where the user content is stored.

Environments provide the following benefits:

- You can group projects together in an environment to simplify the security and management of these projects.
- One SAS Workspace Server can host multiple environments, so work groups can maintain isolation while sharing resources.
- Administration can be performed by using the new SAS Forecast Server Plug-in for SAS Management Console.

SAS Forecast Server Plug-in for SAS Management Console

The SAS Forecast Server Plug-in for SAS Management Console is a Java application that provides a single point of control for managing resources that are used in SAS Forecast Server. You can use SAS Management Console to perform the administrative tasks that are required to create and maintain an integrated environment.

Enhancements to SAS Forecast Server Macros

New Macros

The following macros are new:

- %FSCLEAR - clears the project information that is currently stored in the global macro variables.
- %FSDELENV - deletes the specified environment.
- %FSGETENV - creates an output data set that contains the metadata for the project.
- %FSGETURP - prints to the log the names of the unregistered projects in the current environment. You must use the %FSREGPRJ macro to register these projects.

- %FSNEWENV - creates a new environment.
- %FSREGENV - creates a new environment and can register the projects in that environment.
- %FSREGPRJ - registers a project in the metadata.
- %FSSETOWN - assigns the owner of a project. This macro replaces the deprecated %FS-SETCRB macro, which was used to assign the name of the project's creator.
- %FSURGENV - unregisters the specified environment.

Deprecated Macros

The following table lists the macros that are deprecated in this release.

Table 1.1 Deprecated Macros

Macro Name	Replaced by
FSEXP14	
FSSETCRB	FSSETOWN
FSSETLOC	FSGETENV FSNEWENV FSREGENV

Deprecated Parameters for the FSCOPY and FSMOVE Macros

The following parameters are now obsolete. If a parameter was replaced, the table lists the new parameter.

Table 1.2 Deprecated Parameters

Parameter Name	Replaced by
description	
sourceuser destinationuser	user
sourcepassword destinationpassword	password
sourcehost destinationhost	host

New OWNEDBY Variable

The output data sets created by the FSEXPALL and FSGETPRJ macros now include the OWNEDBY variable. This variable lists the user ID of the project's owner.

Environments and the Workspace Parameter

In this release, environments replace workspace servers in the product macros. For example, in the FSCOPY macro, the SOURCEWORKSPACE and DESTINATIONWORKSPACE parameters are replaced with the SOURCEENVIRONMENT and DESTINATIONENVIRONMENT parameters. Similarly, the WORKSPACE parameter is replaced with the ENVIRONMENT parameter in other macros.

Default Host Port Value

The default value for the port in the HOST=host:port macro option depends on your configuration:

- A level-1 configuration uses a default value of 6411.
- A level-2 configuration uses a default value is 6412.

If you created scripts that use the HOST=host:port option (for example, the FSCREATE macro), then you must update the port value for that option in your script to reflect your system's configuration. The previous release of SAS Forecast Server used a default port value of 5099.

Using Macros to Update Projects

Before you can use macros to update projecs that you created with SAS Forecast Server 1.4, you must first upgrade these projects to SAS Forecast Server 2.1. You cannot go directly from SAS Forecast Server 1.4 to SAS Forecast Server 3.1. Before you upgrade any projects to this release, see "Upgrading from SAS Forecast Server 2.1 to SAS Forecast Server 3.1" in *SAS Forecast Server Administrator's Guide*.

Updates to the Project Code

You can now view the SAS code for the reconciliation steps. In the SAS Code dialog box, the following options are now available:

- **Reconcile Model Forecasts and Overrides** displays the code that SAS Forecast Studio uses to reconcile the model forecasts and overrides.
- **Reconcile Model Forecasts Only** displays the code that SAS Forecast Studio uses to reconcile model forecasts.
- **Reconcile Overrides Only** displays the code that SAS Forecast Studio uses to reconcile overrides.

Each of these options corresponds to a code file that you can use if you are working in batch mode. The following code files are now available:

- The RECONCILE_FORECASTS_AND_OVERRIDES_DO_NOT_IMPORT_DATA file contains the SAS code to reconcile model and forecast overrides.
- The RECONCILE_FORECASTS_DO_NOT_IMPORT_DATA file contains the SAS code to reconcile model forecasts.
- The RECONCILE_OVERRIDES_DO_NOT_IMPORT_DATA file contains the SAS code to reconcile overrides.

New Time Intervals

SAS Forecast Studio now supports the following time intervals:

- Retail 4-4-5 Month
- Retail 4-5-4 Month
- Retail 5-4-4 Month
- Retail 4-4-5 Quarter
- Retail 4-5-4 Quarter
- Retail 5-4-4 Quarter
- Retail 4-4-5 Year
- Retail 4-5-4 Year
- Retail 5-4-4 Year
- ISO 8601 Week
- ISO 8601 Year

SAS Forecast Server does not automatically detect these time intervals. If your data is recorded at these time intervals, then you must select this interval when you select the time ID variable.

New Scenario Analysis

SAS Forecast Studio enables you to quickly generate a large number of forecasts. However, you might want to see how the generated forecasts change when you manipulate the future values of the independent variables. For example, what happens to the forecasts when the price increases? This functionality is called scenario analysis and is available in SAS Forecast Studio from the Scenario Analysis View. In this view, you can create different scenarios without having to create multiple SAS Forecast Studio projects.

New Options for the System-Generated ARIMA Model

By default, SAS Forecast Studio creates an ARIMA model during model generation. Now, you can select what identification method should be used for model inclusion. In the Model Generation pane of the Forecasting Settings dialog box, you can now choose from the following options:

- **Create two models, each of which uses a different identification method for model inclusion**
- **Identify inputs and events for model inclusion before ARMA components**
- **Identify ARMA components for model inclusion before inputs and events**

Report Enhancements

The following enhancements have been made to reports:

- You can customize the report output. You can specify the style and format of the output, such as HTML, RTF, or PDF. Any valid SAS Output Delivery System (ODS) style can be used with reports.
- You can now export any report to an Excel file. You do this when you select CSV as the output format.
- You can view a history of the reports that you have run. The report history includes a status to indicate whether the report ran successfully or failed. You can also view the generated output and SAS logs for the selected report.
- The following reports have been added:
  - Level Model Statistic Volatility Plot
  - Node Reconciled Forecast Analysis Report
  - Level Reconciled Statistic Volatility Plot
  - Level Reconciled Status Table
- The reports have been reorganized, and all sample reports that ship with SAS Forecast Server have been updated.

Enhancements to the User Interface

The following enhancements have been made to the user interface:

- SAS Forecast Studio now ships with a default filter for each BY variable. You cannot remove these filters.
- You can now filter your data based on the value of the selection statistic of fit. The **Selection (Sel.) statistic of fit** category is now available when you create a filter.
- For the forecasts that appear in the Forecasting View, you can now specify how many numbers to show in the plot or data table. You can also specify the number of decimal places to display in the forecasted value.
- Graphs that appear in the SAS Forecast Studio workspace (such as the Forecasting View, Modeling View, or Series View) or in reports now use ODS styles by default. You can also now display your graphs in the Modeling View or Series View in high contrast.
- The Year over Year plot is now available from the Modeling View.
- When trying to determine the best-performing model, you can compare the statistics of fit for each model in the model selection list and you can specify the time series plots of multiple models. This functionality is available from the new **Plots** tab in the Compare Models dialog box.
- If you select the top-down reconciliation method for your hierarchy, then you can specify whether to ignore zero forecasts during reconciliation.
- When preparing your data in the New Project wizard or Forecasting Settings dialog box, you can now set missing values to zero.

Chapter 2

About SAS Forecast Studio

Contents

What Is SAS Forecast Studio?

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

Using this application, you can do the following tasks:

- Generate forecasts automatically by using models that ship with SAS Forecast Studio. You can generate forecasts by using a model selection list that you have selected.
- Create your own forecasting models.
- Perform top-down, bottom-up, and middle-out hierarchical forecasting.
- Visually analyze and diagnose time series data.
- Override forecasts and specify how conflicts should be resolved.
- Analyze the effect on the forecasts of various future values for the input series.
- Export projects as SAS code for processing in a batch environment.
- Generate reports to share forecasting results with other people at your site.

Benefits of Using SAS Forecast Studio

SAS Forecast Studio provides users with the following benefits:

Provides forecasts quickly through a user-friendly graphical interface
: Using SAS Forecast Studio, you can quickly produce high-quality forecasts without any SAS programming. From the user interface, you can interact and change forecasting models, add overrides to specific time periods, apply filters, and run reports.

Provides forecasts that reflect the realities of the business
: SAS Forecast Studio recognizes business drivers, holidays, or events in the input data source. Therefore, forecasts better reflect the business and require less overriding and fewer manual interventions. SAS Forecast Studio automatically builds and selects the most appropriate model for your data.

Improves forecasting performance across all products and locations, at any level of aggregation
: SAS has a complete array of advanced forecasting methods and can statistically estimate the effect of sales and marketing events. Graphical output quickly shows the effects of holidays, marketing events, sales promotions and unexpected events, such as weather. Users can use this information to improve their forecasts and plan future sales promotions and marketing events.

How SAS Forecast Studio Works

When you first open SAS Forecast Studio, you need to open a project. When a project is created, you can specify the following information:

- the input data set
- which variables contain the data for the time ID, dependent variable, and other roles
- whether the project is forecast hierarchically
- any events to include
- any customized forecast settings

After the project has been created, the historical and forecasted data appears in the main workspace. These forecasts are available in graphical and tabular output.

After analyzing the results, you can change your forecast by selecting a different forecasting model and specifying overrides for selected time periods.

How SAS Forecast Studio Fits into the SAS Forecast Server

The SAS Forecast Server is a large-scale automatic forecasting solution that enables organizations to produce huge quantities of high-quality forecasts quickly and automatically.

The SAS Forecast Server has three components:

- the SAS Analytics Platform
- SAS Forecast Server middle tier
- SAS Forecast Studio

For information about how to install each component, see the *SAS Forecast Server Administrator's Guide*.

SAS Forecast Studio is the client component of the SAS Forecast Server. SAS Forecast Studio provides a Java-based, graphical interface to the forecasting and time series analysis procedures contained in SAS High-Performance Forecasting and SAS/ETS software.

SAS High-Performance Forecasting automatically selects the appropriate model for each item being forecast based on user-defined criteria. Holdout samples can be specified so that models are selected not only by how well they fit past data but by how appropriate they are for predicting the future. If the best forecasting model for each item is unknown or the models are outdated, a maximum level of automation can be chosen in which all three forecasting steps (model selection, parameter estimation, and forecast generation) are performed. If suitable models have been determined, you can keep the current models and reestimate the model parameters and generate forecasts. For maximum processing speed, you can keep previously selected models and parameters and choose to simply generate the forecasts.

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*.

How to Get Help for SAS Forecast Studio

There are two ways to access Help from within SAS Forecast Studio:

- Select **Help** → **SAS Forecast Studio Help**.
- Click the **Help** button, which is available from any SAS Forecast Studio dialog box, window, or wizard pane.

Accessibility and Compatibility Features

SAS Forecast Studio 3.1 includes accessibility and compatibility features that improve the usability of the product for users with disabilities, with exceptions noted below. These features are related to accessibility standards for electronic information technology that were adopted by the U.S. Government under Section 508 of the U.S. Rehabilitation Act of 1973. If you have specific questions about the accessibility of SAS Forecast Studio, send them to accessibility@sas.com or call SAS Technical Support.

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

Table 2.1 Accessibility Exceptions

Section 508 Accessibility Criteria	**Support Status**	**Explanation**
(c) A well-defined on-screen indication of the current focus shall be provided that moves along interactive interface elements as the input focus changes. The focus shall be programmatically exposed so that assistive technology can track the focus and any focus changes.	Supported with minor exceptions	In the New Project wizard, the initial focus is not set correctly in the following steps: • In Steps 3, 5, 6, and 7, the initial focus is set to the **Next** button. • In Step 4, the initial focus is set to the **Cancel** button. • In Step 8, the initial focus is set to the **Next** button. However, this button is disabled. Because the initial focus is set on a button, the previous text in the step is not read. SAS currently plans to address this and other accessibility issues in the New Project wizard in a future release. In the Scenario Analysis View, you cannot use the keyboard to open the context menu for the graph. Therefore, you cannot change the graph properties, action mode, or data options.

Table 2.1 *continued*

Section 508 Accessibility Criteria	Support Status	Explanation
(d) Sufficient information about a user-interface element (including the identity, operation, and state of the element) shall be available to the assistive technology. When an image represents a program element, the information conveyed by the image must also be available in text.	Supported with exceptions	The following components could not be read by a JAWS screen reader: • Data in some tables might not be read properly. For example, the input table in the Scenario Analysis View cannot be read. This is a known issue in Java. • In the New Project wizard, the text for some of the options cannot be read. • The JAWS screen reader recognizes window items, such as buttons, drop-down lists, and text boxes. However, the reader does not explain how to access these window items by using the keyboard. SAS currently plans to address these and additional capability issues with screen readers in a future release.

Table 2.1 *continued*

Section 508 Accessibility Criteria	**Support Status**	**Explanation**
(g) Applications shall not override user-selected contrast and color selections and other individual display attributes.	Supported with exceptions	When your operating system is set to high contrast, the following icons and controls might be difficult to discern from the background: • In the hierarchy, the expand (+) and collapse (-) icons for each branch of the hierarchy appear black on black in high contrast, so it is difficult to tell if a branch is fully expanded or collapsed. • In Step 3 of the New Project wizard, the arrow buttons that you use to assign variables to roles appear black on black in high contrast. • In the Hierarchy and Variable Settings dialog box, the arrow buttons that you use to order variables in the hierarchy appear black on black in high contrast. • In the Modeling View and Series View, the minimize, maximize, and close buttons in the plot and table windows appear black on black in high contrast. • In the Scenario Analysis View, the colors in the grid table do not pass the foreground and background contrast ratio. Also, the **Set scenario forecast values as overrides** link cannot be read. Finally, you cannot change the colors in the graph.

Table 2.1 *continued*

Section 508 Accessibility Criteria	**Support Status**	**Explanation**
		SAS currently plans to address these and additional color and contrast issues in a future release. Until that time, low vision users who require high contrast might find the use of a screen magnifier with reverse video setting to be a sufficient accommodation. Graphs in the Modeling View and Series View and graphs that are in output from a report use styles from the SAS Output Delivery System(ODS). You can choose to display these graphs in high contrast by selecting **View** → **Graph Style** → **High Contrast**.
(l) When electronic forms are used, the form shall enable users of assistive technology to access the information, field elements, and functionality required for completion and submission of the form, including all directions and cues.	Supported with exceptions	To work with assistive technologies, software generally needs to support a visible focus indicator and work with a screen reader. SAS Forecast Studio has exceptions noted in both criteria (c) and (d) that limit its compliance with assistive technologies.

Chapter 3

Starting SAS Forecast Studio

Contents

How You Can Run SAS Forecast Studio

SAS Forecast Studio runs on Windows. Depending on your site, you can run SAS Forecast Studio in either of the following ways:

- from a local installation of SAS Forecast Server on your computer. If you installed SAS Forecast Studio on a computer in the Windows operating environment, select **Start** → **Programs** → **SAS** → **SAS Forecast Studio** → **SAS Forecast Studio 3.1**.
- by using Java Web Start. Contact your site administrator for the URL for your site.

Log On to SAS Forecast Studio

Whether you run a local installation of SAS Forecast Studio or you use Java Web Start, you need to log on to the application.

To log on to SAS Forecast Studio, complete the following steps:

1. Specify your user name and password for the SAS Forecast Server.
2. (Optional) Select the server that you want to use. The list of available servers is created by your site administrator. If only one SAS Forecast Server is available, then you cannot change this selection. For more information, see *SAS Forecast Server Administrator's Guide*.
3. Click **Log On** to start SAS Forecast Studio.

Chapter 4

Using the Workspace

Contents

Overview of the Workspace

The SAS Forecast Studio workspace consists of the following components:

Figure 4.1 SAS Forecast Studio workspace

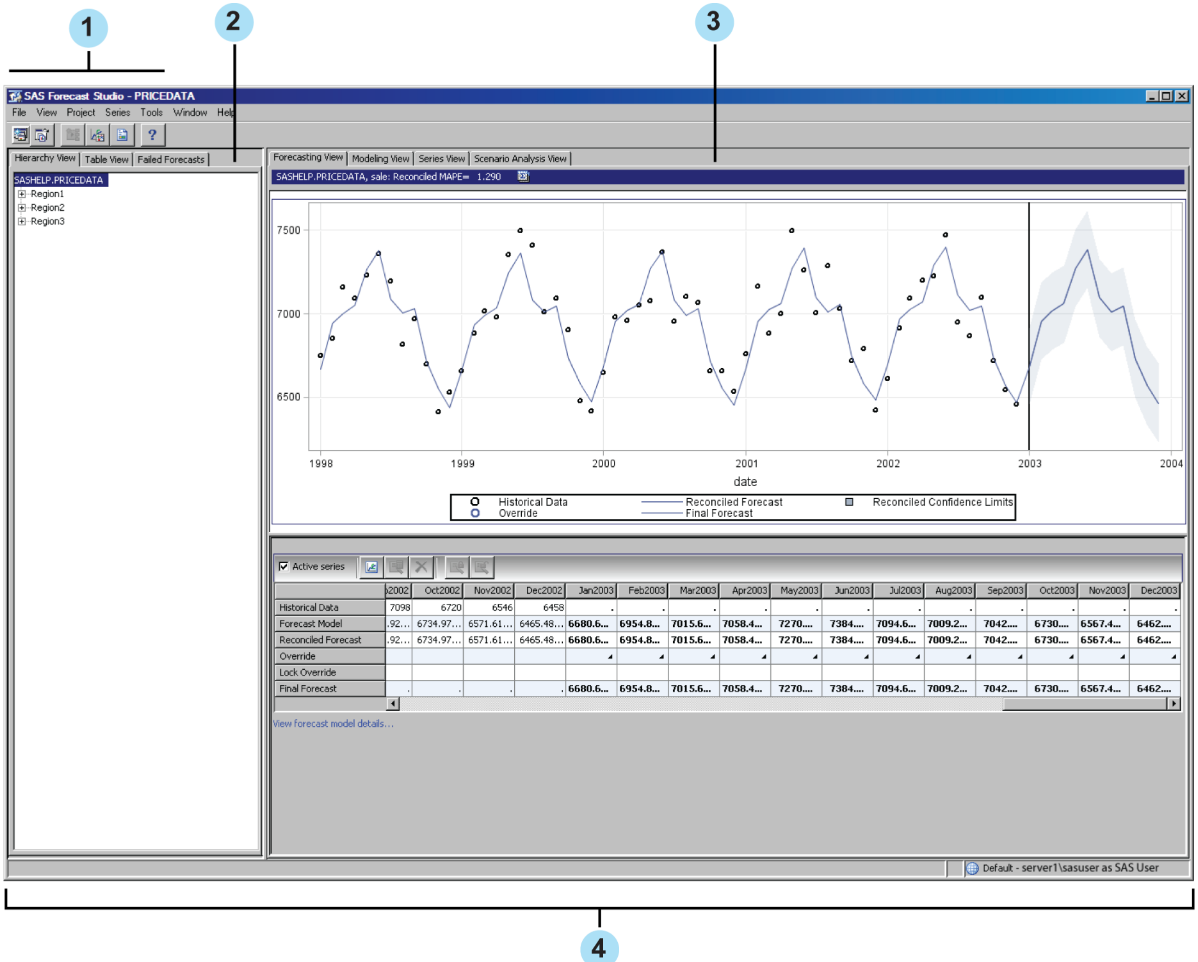

1. The toolbar displays some of the most commonly used SAS Forecast Studio options, so that you can quickly and easily manage your project.

2. The overview panel enables you to choose the series that you want to appear in the view area. The following tabs are available:

   - **Hierarchy View** - displays the project data as a tree when you forecast the data hierarchically. You can expand and collapse the branches of the tree.
   - **Table View** - displays the series that meet a selected filter criterion. You can select from predefined or user-defined filters.
   - **Failed Forecasts** - lists the series where the forecasts failed.

3. The view area displays the forecasting results for the selected series in the overview panel. The following tabs are available:

   - **Forecasting View** - displays the forecast plot and the data table for the selected series.
   - **Modeling 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 and compare models.

- **Series View** - enables you to view various plots and tables for the selected series. The plots in the Series View use the historical data of the time series.
- **Scenario Analysis View** - enables you to vary the future values for the input series to determine the impact on the generated forecasts.

4. The status bar can contain the following information:

   - the status of the current action in SAS Forecast Studio. For example, if you are generating forecasts, then the status bar displays the name of the SAS procedure that is currently executing.
   - A warning icon appears if there are errors in the log from the last code submission, and you have not yet opened the SAS log to view these errors. Click  to open the SAS log.
   - the environment, metadata account, and your user ID.

The Hierarchy View

The Hierarchy View displays a tree view of the hierarchy. This hierarchy is defined by the variables that you assign as the classification (BY) variable when you create the project. You can expand or collapse the branches in the tree.

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

NOTE: You can modify the configuration of the hierarchy from the Hierarchy and Variable Settings dialog box.

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

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

- If you add a note for a series, the  icon appears next to the series in the hierarchy.

Figure 4.2 Hierarchy View

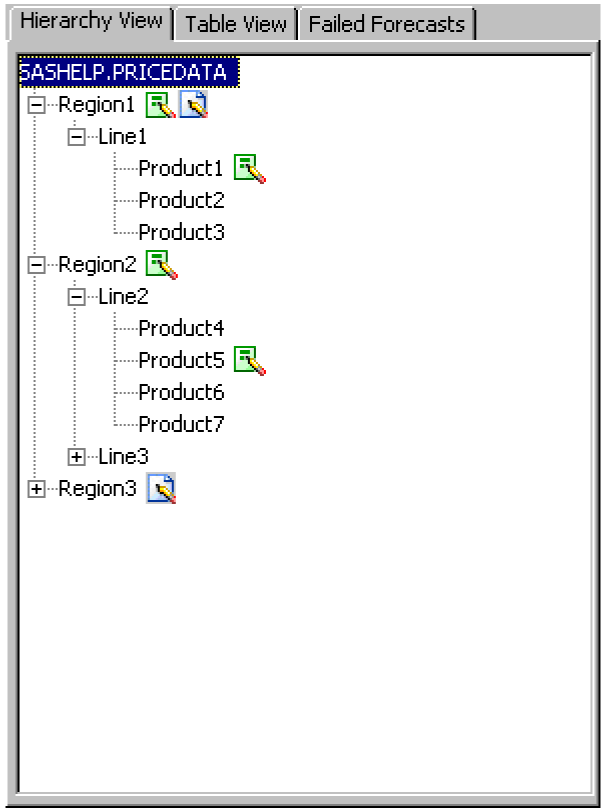

The Table View

By default, the Table View contains the following information:

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

  If you specified an out-of-sample range, then you can choose to filter on the statistic of fit for the holdout sample for the best-performing model in each series.

  An asterisk (*) in the table indicates aggregation. An asterisk in a hierarchy level (BY variable) column means that all of the series at that level have been aggregated to the next level of the hierarchy.

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

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

By default, the series are sorted in ascending order by the statistic of fit for model selection. 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 SAS Forecast Studio reforecasts the project, then the columns are again sorted by the statistic of fit for model selection. After a reforecast, a column remains in the table view if it meets one of the following criteria:

- The table view always includes columns for any BY variables, one or more dependent variables, the statistic of fit that is used in model selection, and the reconciled statistic of fit that is used in model selection.
- Columns that are part of the selected filter definition appear in the table view.
- You specifically add the column to the table view. For more information, see "Customize the Table View" on page 32.

Figure 4.3 Table View

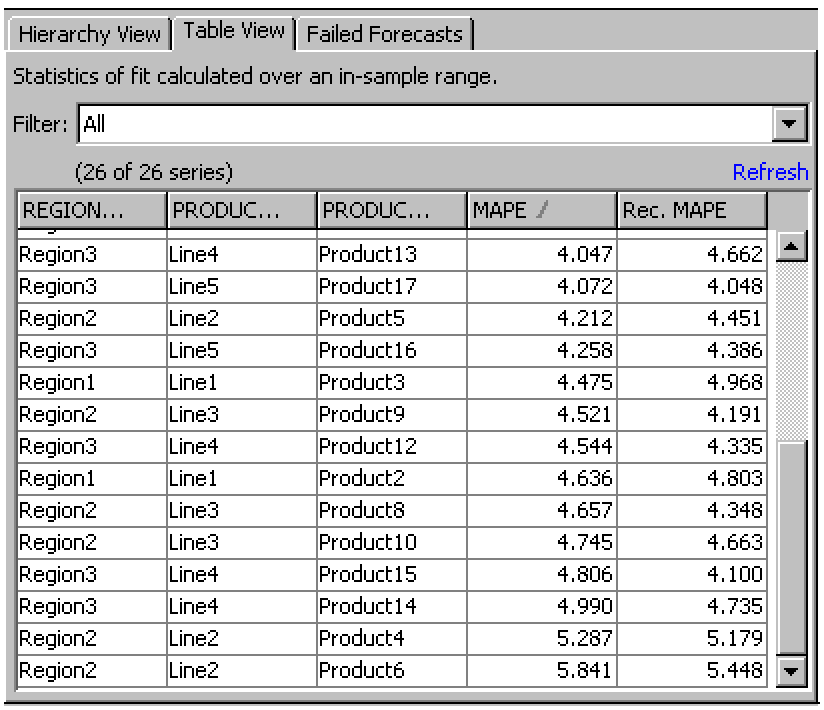

REGION...	PRODUC...	PRODUC...	MAPE	Rec. MAPE
Region3	Line4	Product13	4.047	4.662
Region3	Line5	Product17	4.072	4.048
Region2	Line2	Product5	4.212	4.451
Region3	Line5	Product16	4.258	4.386
Region1	Line1	Product3	4.475	4.968
Region2	Line3	Product9	4.521	4.191
Region3	Line4	Product12	4.544	4.335
Region1	Line1	Product2	4.636	4.803
Region2	Line3	Product8	4.657	4.348
Region2	Line3	Product10	4.745	4.663
Region3	Line4	Product15	4.806	4.100
Region3	Line4	Product14	4.990	4.735
Region2	Line2	Product4	5.287	5.179
Region2	Line2	Product6	5.841	5.448

Failed Forecasts

The **Failed Forecasts** tab displays the series where one of the following problems occurred:

- A forecast was not produced. In the **Problems** column, these series are identified by "Failed to forecast."
- SAS Forecast Studio could not fit a model from the model selection list, and, instead, the system used a model from the default model selection lists that ship with SAS Forecast Studio.

In the **Problems** column, these series are identified by "Model list problem". For more information about these model selection lists, see Appendix E, "Default Model Selection Lists."

An asterisk (*) in the table indicates aggregation. An asterisk in a hierarchy level (BY variable) column means that all of the series at that level have been aggregated to the next level of the hierarchy.

The Forecasting View

In the Forecasting View, you can view the forecast for the selected series in the hierarchy or the table view. If you open another view, you can return to the Forecasting View by selecting **View** → **Details** → **Forecasting View**.

The Forecasting View is divided into the following components:

Figure 4.4 Forecasting View

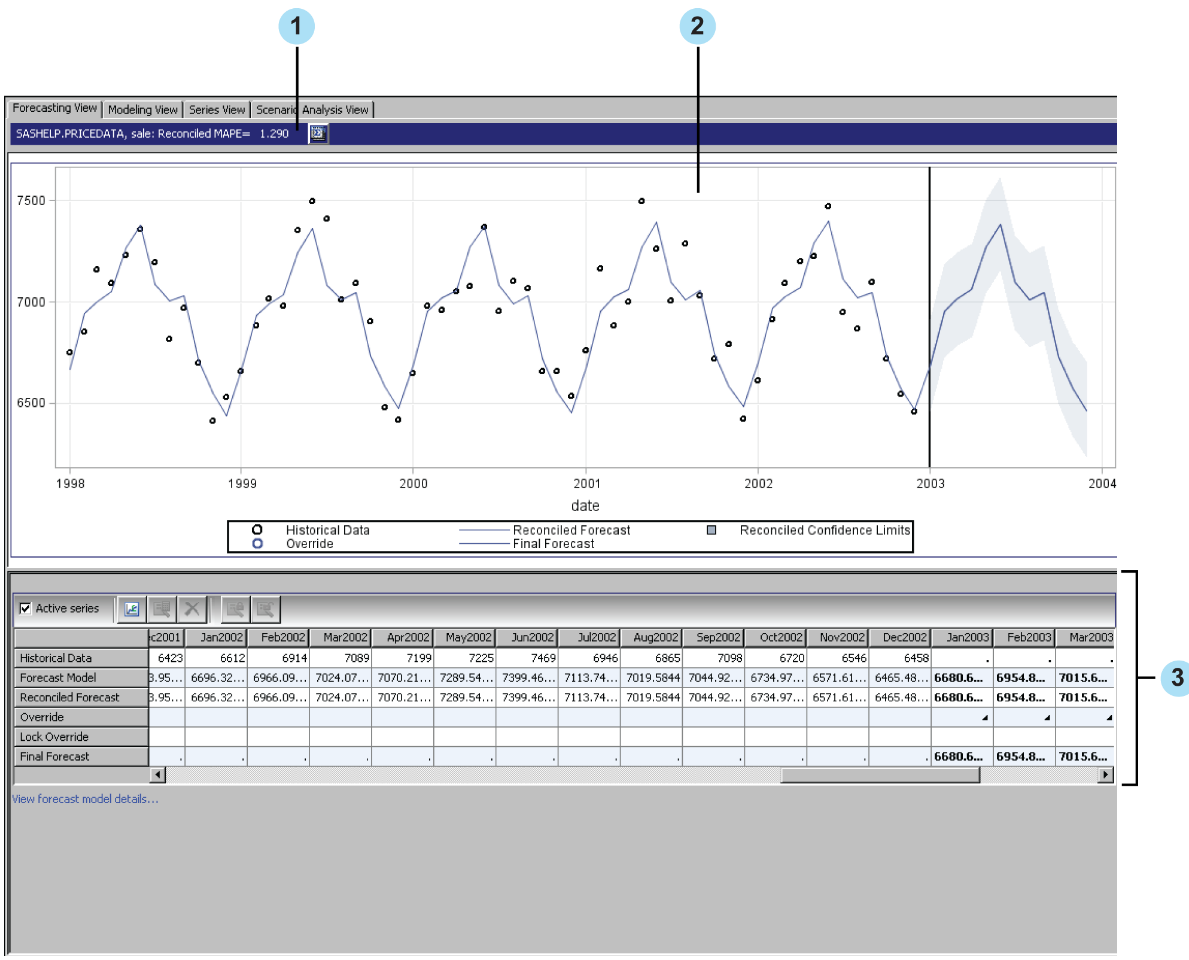

1. At the top of the Forecasting View, you can see the name of the series that you selected, the dependent variable for the project, and the statistics of fit for the series. If you specified an out-of-sample range, then you can click to view additional information about the statistics of fit.

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

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

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

   If you do not forecast your data hierarchically or you forecast your data hierarchically but do not select the **Reconcile hierarchy** check box, then the forecast plot shows statistical forecasts instead of the reconciled forecasts.

   You can select what to display in the forecast plot. For more information, see "Customize the Contents of the Forecast Plot or Data Table" on page 32.

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

   When the **Active series** 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 75.

   By default, the data table displays the following information:

   - historical data. Historical values are displayed in regular font. Future values are displayed in bold. Missing values appear as . (a period).
   - reconciled or statistical forecast.
   - overrides. You can specify overrides only for future time periods. The cells that you can edit have a white background.
   - final forecast.

   To customize the content of the data table, see "Customize the Contents of the Forecast Plot or Data Table" on page 32.

The Modeling View

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

To open the Modeling View, select **View** → **Details** → **Modeling View**.

Figure 4.5 Modeling View

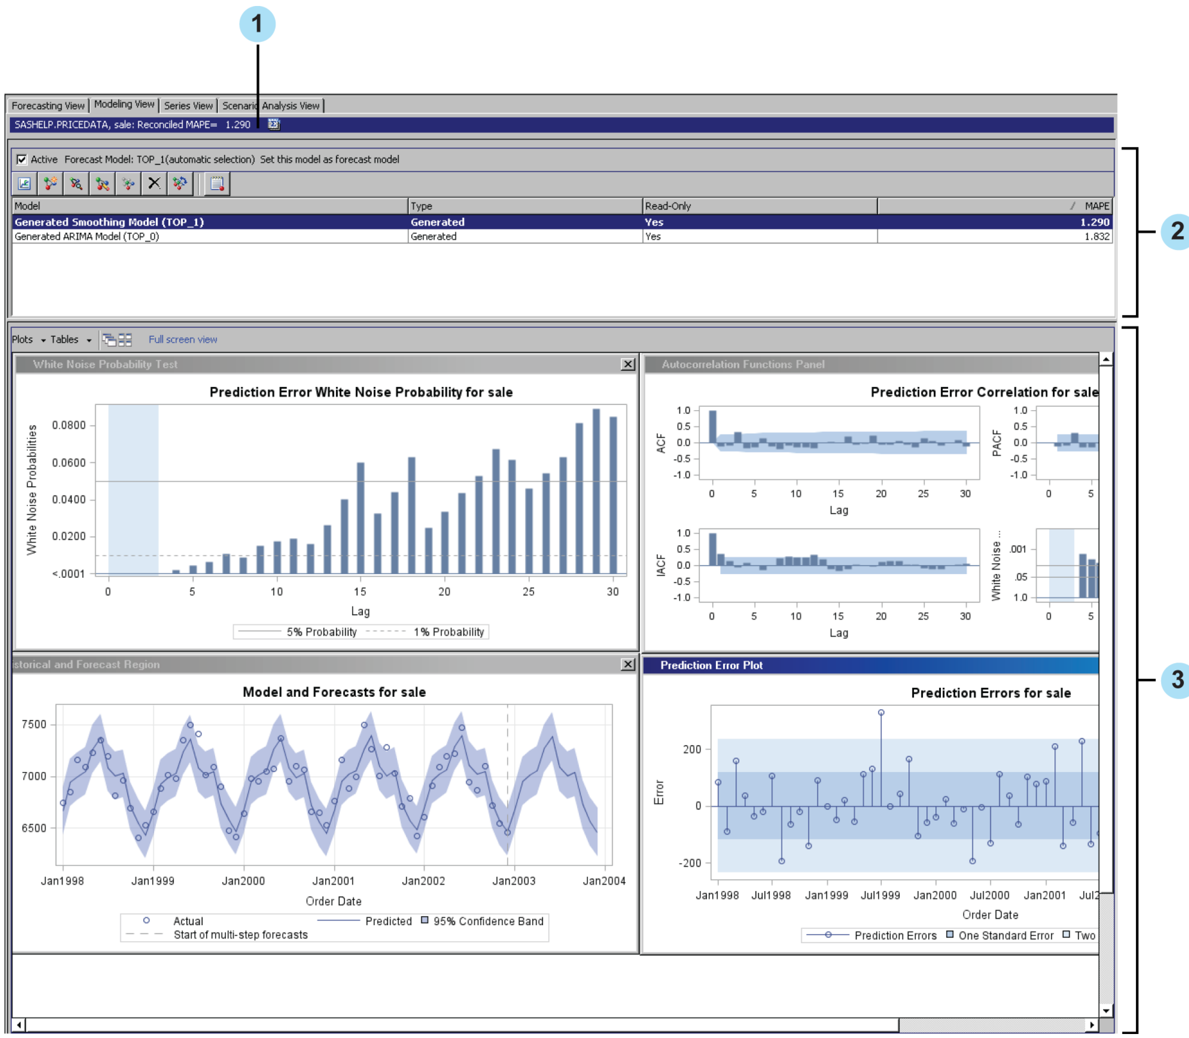

The Modeling View is divided into the following components:

1. At the top of the Modeling View, you can see the name of the series that you selected, the dependent variable for the project, and the statistics of fit for the series. If you specified an out-of-sample range, then you can click to view additional information about the statistics of fit.

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

3. By default, the following plots appear in the Modeling View:

    - a plot that includes the generated forecasts in the forecast horizon
    - a plot of the residuals for the predicted errors
    - a plot of the white noise probability test for the dependent variable

- plots of the autocorrelation function, partial autocorrelation function, inverse autocorrelation function, and white noise probability test for the predicted errors

You can open additional model plots or additional model tables. SAS Forecast Studio saves your final configuration and uses that same configuration the next time that you open the Modeling View while working in the current project. When you close the current project, then these settings are lost.

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

The Series View

The Series View provides diagnostics to help you identify model components that might help to improve the accuracy of your forecasts. From the Series View, you can choose the time series plots and tables that you want to display for a selected series. The plots in the Series View use the historical data for the time series. You can also transform a time series from this window.

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

To open the Series View, select **View** → **Details** → **Series View**.

Figure 4.6 Series View with the Plots Tiled in the Window

The Series View contains the following components:

1. At the top of the Series View, you can see the name of the series that you selected, the dependent variable for the project, and the statistics of fit for the series. If you specified an out-of-sample range, then you can click  to view additional information about the statistics of fit.

2. Use these options to transform the time series.

3. By default, the following graphs appear in the Series View:

   - a plot of the current time series.
   - a plot of the seasonal decomposition. This plot is not available for nonseasonal series, such as yearly data.
   - a plot in the log scale of the white noise probability test for the dependent variable.
   - plots of the autocorrelation function, partial autocorrelation function, inverse autocorrelation function, and white noise probability test.

You can open time series plots or tables to help you analyze the time series. SAS Forecast Studio saves your final configuration and uses that same configuration the next time that you open the Series View while working in the current project. When you close the current project, then these settings are lost.

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

The Scenario Analysis View

The Scenario Analysis View enables you to vary the future values for the input series to determine the impact on the generated forecasts. You can create and edit scenarios from this view.

To open the Scenario Analysis View, select **View** → **Details** → **Scenario Analysis View**.

Figure 4.7 Scenario Analysis View

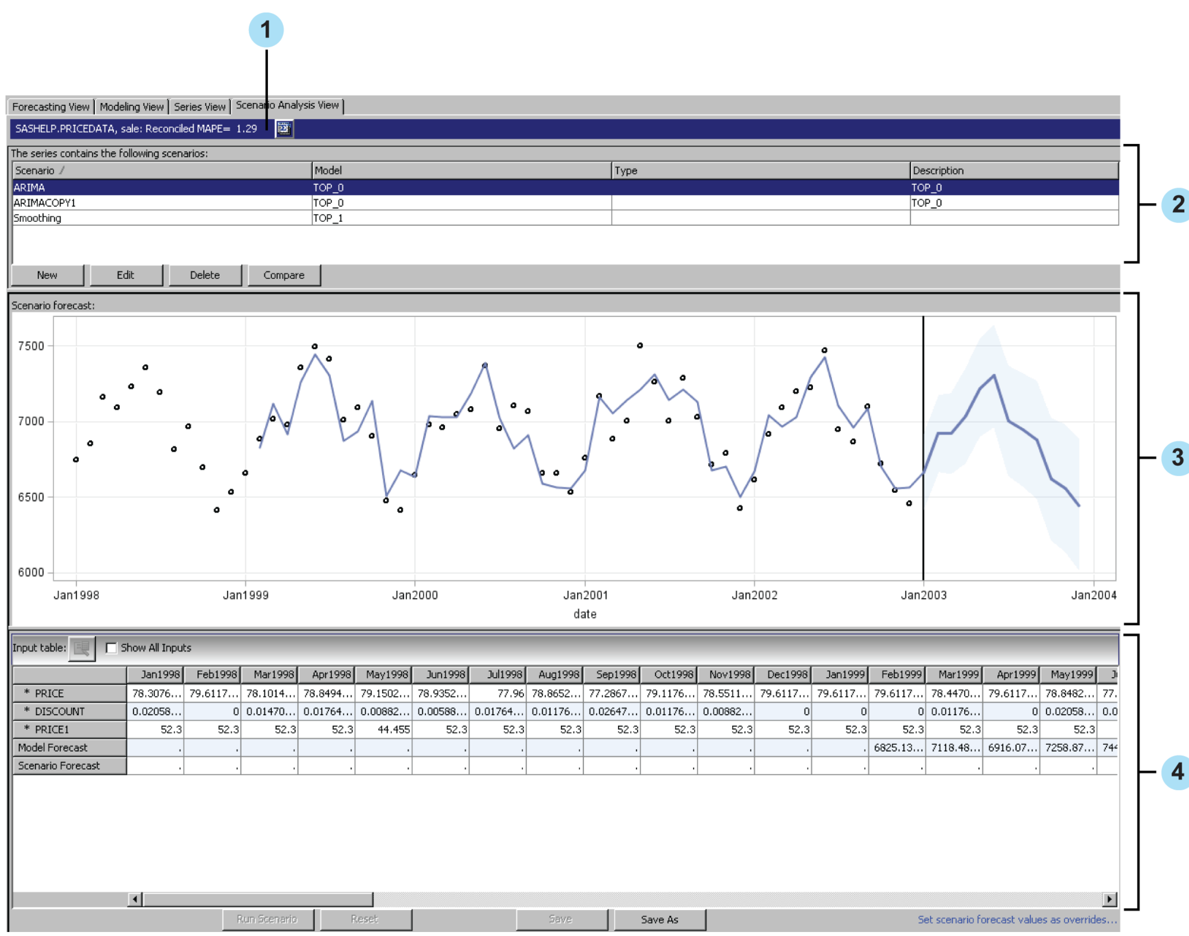

The Scenario Analysis View contains the following components:

1. At the top of the Scenario Analysis View, you can see the name of the series that you selected, the dependent variable for the project, and the statistics of fit for the series. If you specified an out-of-sample range, then you can click to view additional information about the statistics of fit.
2. This table lists the scenarios for the project. You can use the options at the bottom of this section to add new scenarios, edit existing scenarios, delete scenarios, or compare scenarios.
3. The graph shows the historical time series data and the baseline forecast for the selected scenario.
4. The input table displays the current values for the independent variables, the forecasts that were generated by the model, and the forecasts that were generated by the scenario. You can change the input values that influence the forecasts. Any cells that you can edit have a black triangle in the corner.

Customizing Your Workspace

Customize the Table View

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

1. Select **View** → **Select Table View Columns**. The Table View Columns dialog box opens.
2. Select a category from the **Available** drop-down list. Then select the statistics of fit or property to display in the table and click .
3. Select whether to include a column for overrides or a column for notes in the table.
4. Click **OK**.

Customize the Contents of the Forecast Plot or Data Table

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

1. Select **View** → **Edit Forecasting View Properties**. The Forecasting View Properties dialog box opens.
2. Select what properties to display on the forecasting plot and what properties to display in the data table. Click **OK**.

Part II

Getting Started in SAS Forecast Studio

Chapter 5

Preparing an Input Data Set for SAS Forecast Studio

Contents

Understanding Time Series Data

Requirements for Time Series Data

To generate forecasts in SAS Forecast Studio, you need time series data. You might already have this time series data, or you might have transactional data. If you have transactional data, you can use the accumulation options in SAS Forecast Studio to convert the transactional data into a time series.

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 36.

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

1. The data is accumulated to the appropriate time interval if the input is one of the following:
   - timestamped data that is recorded at no particular frequency (also called transactional data)
   - data recorded at a smaller time interval than needed for forecasting
2. Any gaps in the data are filled in. Gaps appear when there is not an observation for each time period or when the data is not equally spaced. The added observations have the required values of the time ID variable and the value that you specified for missing values. For more information, see “Working with Missing Values” on page 38.
3. The data is sorted by the BY variables and the time ID variable.

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

Selecting the Input Data Set

Required Format

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

You can have the following variables in the input data set:

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

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

NOTE: The names of the variables cannot match any of the reserved variable names that are used in the output data set. For more information, see Appendix B, "Reserved Names in SAS Forecast Studio."

For more information about variable roles, see Chapter 6, "Understanding Variable Roles."

Examples

Here are two examples of input data sets for SAS Forecast Studio:

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

Date	Revenue	Avg. Price	Holiday
JAN2008	18817	26.3	0
FEB2008	52573	25.3	0
...	...	...	...
DEC2008	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
JAN2008	355	25.3	0	Region 1	Product 1
FEB2008	398	25.3	0	Region 1	Product 1
...	...	...	...	...	...
JAN2008	555	19.8	0	Region 1	Product 2
FEB2008	390	25.3	0	Region 1	Product 2
...	...	...	...	...	...
JAN2008	301	27.1	0	Region 2	Product 1
FEB2008	350	25.3	0	Region 2	Product 1
...	...	...	...	...	...
JAN2008	314	27.2	0	Region 2	Product 2
FEB2008	388	25.3	0	Region 2	Product 2
...	...	...	...	...	...
DEC2008	518	20.3	1	Region 2	Product 2

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

Working with Missing Values

If your data contains missing values in variables other than the time ID variable (such as the dependent, independent, and external forecast variables), you can specify how to interpret missing values (regardless of the variable roles) when you create a project.

Additional Information

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

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

Chapter 6

Understanding Variable Roles

Contents

Overview of the Types of Roles

When you create a project in SAS Forecast Studio, you can assign the following roles to the variables in the input data set:

- time ID variable
- BY variable
- dependent variable
- independent variable
- reporting variable
- adjustment variable

Time ID Variables

What Is the Time ID Variable?

You specify the time ID variable when you create the project using the New Project Wizard. After the project has been created, you cannot change the time ID variable. The time ID variable is a variable in the input data set that contains the SAS date or datetime 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 does not support time values. When you assign a variable to the time ID role, SAS Forecast Studio recognizes the data's format. If your data uses a time format, then SAS Forecast Studio generates a message stating that time data is not valid for the time ID variable. If your data has no format, then SAS Forecast Studio assumes that the values are dates, and errors will occur when your project is created.

Time Intervals in SAS Forecast Studio

SAS Forecast Studio supports the following time intervals:

Day
: specifies daily intervals.

Hour
: specifies hourly intervals.

ISO 8601 year
: specifies ISO 8601 yearly intervals. The ISO 8601 year starts on the Monday on or immediately preceding January 4th. Note that it is possible for the ISO 8601 year to start in December of the preceding year. Also, some ISO 8601 years contain a leap week.

ISO 8601 week
: specifies ISO 8601 weekly intervals of seven days. Each week starts on Monday. The starting subperiod s is in days (DAY). Note that WEEKV differs from WEEK in that WEEKV.1 starts on Monday, WEEKV.2 starts on Tuesday, and so forth.

Minute
: specifies minute intervals.

Month
: specifies monthly intervals.

Quarter
: specifies quarterly intervals (every three months). The starting subperiod is in months.

Retail 4-4-5 Year
: specifies ISO 8601 weekly interval, except that the starting subperiod s is in retail 4-4-5 months.

Retail 4-5-4 Year
: specifies ISO 8601 weekly interval, except that the starting subperiod s is in retail 4-5-4 months.

Retail 5-4-4 Year
: specifies ISO 8601 weekly interval, except that the starting subperiod s is in retail 5-4-4 months.

Retail 4-4-5 Month
: specifies retail 4-4-5 monthly intervals. The 3rd, 6th, 9th, and 12th months are five ISO 8601 weeks long with the exception that some 12th months contain leap weeks. All other months are four ISO 8601 weeks long. R445MON intervals begin with the 1st, 5th, 9th, 14th, 18th, 22nd, 27th, 31st, 35th, 40th, 44th, and 48th weeks of the ISO year

Retail 4-5-4 Month
: specifies retail 4-5-4 monthly intervals. The 2nd, 5th, 8th, and 11th months are five ISO 8601 weeks long. All other months are four ISO 8601 weeks long with the exception that some th months contain leap weeks. R454MON intervals begin with the 1st, 5th, 10th, 14th, 18th, 23rd, 27th, 31st, 36th, 40th, 44th, and 49th weeks of the ISO year.

Retail 5-4-4 Month
: specifies retail 5-4-4 monthly intervals. The 1st, 4th, 7th, and 10th months are five ISO 8601 weeks long. All other months are four ISO 8601 weeks long with the exception that some th months contain leap weeks. R544MON intervals begin with the 1st, 6th, 10th, 14th, 19th, 23rd, 27th, 32nd, 36th, 40th, 45th, and 49th weeks of the ISO year.

Retail 4-4-5 Quarter
: specifies retail 4-4-5 quarterly intervals (every 13 ISO 8601 weeks). Some fourth quarters will contain a leap week. The starting subperiod s is in retail 4-4-5 months.

Retail 4-5-4 Quarter
: specifies retail 4-5-4 quarterly intervals (every 13 ISO 8601 weeks). Some fourth quarters will contain a leap week. The starting subperiod s is in retail 4-5-4 months.

Retail 5-4-4 Quarter
: specifies retail 5-4-4 quarterly intervals (every 13 ISO 8601 weeks). Some fourth quarters will contain a leap week.

Second
: specifies second intervals.

Semimonth
: specifies semimonthly intervals. Each month consists of two periods. The first period starts on the first and the second period starts on the 16th.

Semiyear
: specifies intervals every six months. The starting subperiod is in months.

Ten-day
: specifies ten-day intervals. Each month consists of three periods. The first period is the first through the tenth day of the month. The second period is the eleventh through the twentieth day of the month. The third period is the twenty-first through the end of the month.

Week
: specifies weekly intervals of seven days. The days of the week are numbered as follows:

Value of the Shift	Day of the Week
1	Sunday
2	Monday
3	Tuesday
4	Wednesday
5	Thursday
6	Friday
7	Saturday

Weekday
: specifies daily intervals with weekend days included in the preceding weekday. The weekday interval is the same as the day interval, except that the weekend days are absorbed into the preceding weekday. The default weekend days are Saturday and Sunday, but you can specify the days to include in the weekend. If you use the default weekend, then there are five weekday intervals in a calendar week: Monday, Tuesday, Wednesday, Thursday, and the three-day period Friday, Saturday, and Sunday.

Year
: specifies yearly intervals. The starting subperiod is in months.

Understanding SAS Time Intervals

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 datetimes before January 1, 1960.

For many businesses, their time series data is equally spaced, or any two consecutive indices have the same difference between the time intervals. The following table shows an equally spaced time series with a one-year interval.

Year	Number of Sales
2005	42,100
2006	45,000
2007	47,000
2008	50,000

If the time interval cannot be detected from the variable that you assign, then you need to specify the interval and seasonal cycle length. For example, the following table shows an unequally spaced time series.

Year	Number of Sales
2003	32,100
2004	45,000
2007	47,000
2008	50,000

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

In SAS Forecast Studio, the interval determines the frequency of the output. You can modify the time interval. You can change the interval from a higher frequency to a lower frequency or from a lower frequency to a higher frequency. Time intervals are specified in SAS by using character strings. Each of these strings is formed according to a set of rules that enables you to create an almost infinite set of attributes. For each time interval, you can specify the type (such as monthly or weekly), a multiplier, and a shift (the offset for the interval). 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.

NOTE: If you change the time interval after adding overrides to your project, then any overrides, custom models, and notes are removed.

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. However, you can specify a seasonal cycle length other than the default if you want to model a cycle in the data. For example, your data might contain a 13-week cycle, so you need to specify a 13-week seasonal cycle length in SAS Forecast Studio.

The form of an interval is *name<multiple><.starting point>* where the following conditions obtain:

name
: is the name of the interval.

multiple
: specifies the multiple of the interval. This value can be any positive number. By default, the multiplier is 1. For example, YEAR2 indicates a two-year interval.

.starting-point
: specifies the starting point for the interval. By default, this value is one. A value greater than 1 shifts the start to a later point within the interval. The unit for the shift depends on the interval. For example, YEAR.4 specifies a shift of three months, so the year is from April 1 through March 31 of the following year.

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

Interval Name (in SAS code format)	**Default Starting Point**	**Shift Period**	**Example**
YEARm.s	January 1	Months	YEAR2.7 specifies an interval of every two years. Because the value for the shift is 7, the first month in the year is July.
SEMIYEARm.s	January 1 July 1	Months	SEMIYEAR.3 - six-month intervals, March-August and September-February.
QTRm.s	January 1 April 1 July 1 October 1	Months	QTR3.2 - three-month intervals starting on April 1, July 1, October 1, and January 1.
SEMIMONTHm.s	First and sixteenth of each month	Semi-monthly periods	SEMIMONTH2.2 - intervals from the sixteenth of one month through the fifteenth of the next month.
MONTHm.s	First of each month	Months	MONTH2.2 - February-March, April-May, June-July, August-September, October-November, and December-January of the following year.
TENDAYm.s	First, eleventh, and twenty-first of each month	Ten-day periods	TENDAY4.2 - Four ten-day periods starting at the second ten-day period
WEEKm.s	Each Sunday	Days (1=Sunday . . . 7=Saturday)	WEEK6.3 specifies six-week intervals starting on Wednesdays.
DAYm.s	Each day	Days	DAY3 - three-day intervals starting on Sunday
HOURm.s	Start of the day (midnight)	Hours	HOUR8.7 specifies eight-hour intervals starting at 6:00 a.m., 2:00 p.m., and 10:00 p.m.

Classification (BY) Variables

SAS Forecast Studio 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 specify the order of the hierarchy when you create the project. An example of a hierarchy is Region > Product Line > Product Name. When you preview the hierarchy in SAS Forecast Studio, you see the following:

Figure 6.1 Example of Hierarchy

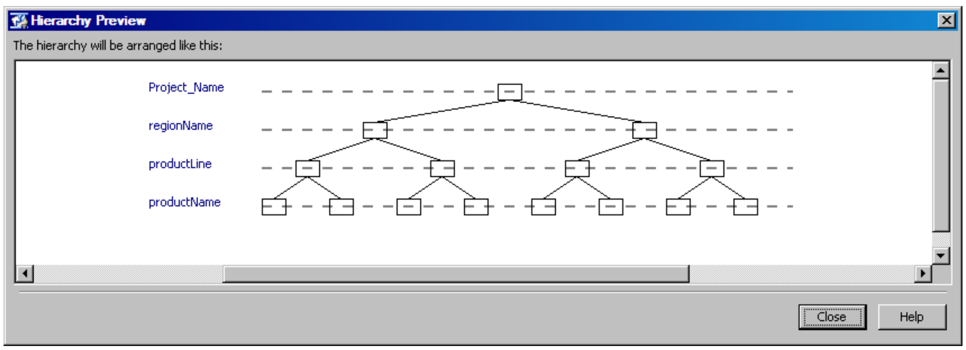

NOTE: : If you have a hierarchy and get a message that the time interval of the data cannot be detected, verify that the BY variables are specified in the same order that the data set is sorted.

For more information about hierarchies, see Chapter 7, "Forecasting Hierarchically."

Dependent Variables

Dependent variables are the variables that you want to model and forecast. You must assign at least one numeric variable to this role. If you forecast your data hierarchically, then only one dependent variable can be forecast.

For example, you want to forecast the sales for each product. When you create your project in SAS Forecast Studio, you assign the Number of Sales variable to the dependent variable role. When the project is created, the forecasting plot shows the sales forecast, as shown below.

Figure 6.2 Forecast Values for Sales Dependent Variable

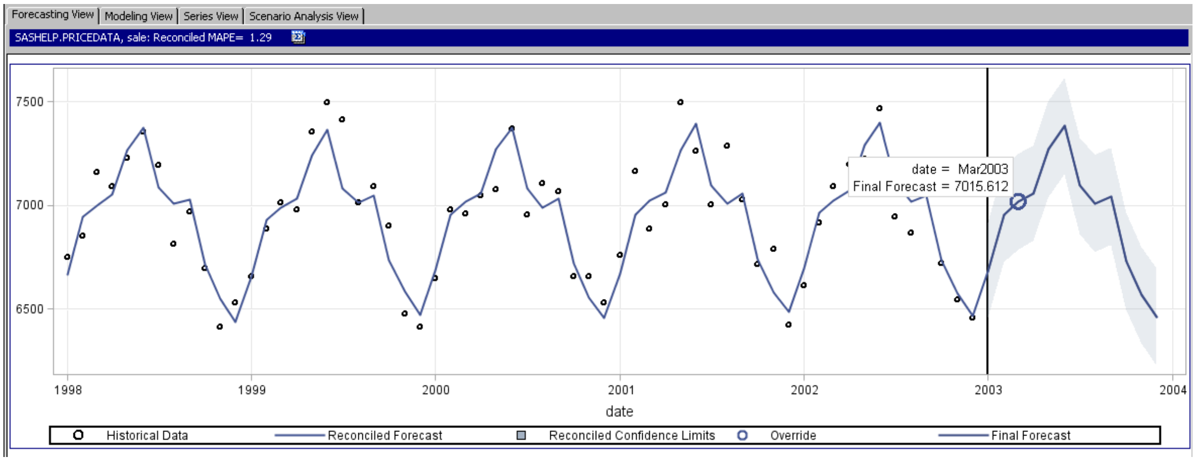

Independent Variables

Independent variables are any explanatory, input, predictor, or causal factor variables. You can assign only numeric variables to this role. When creating the system-generated models, SAS Forecast Studio tries to use the independent variables in the model generation.

For example, in your initial forecasts, you assign the Number of Sales variable to the dependent variable role, and you do not assign an independent variable. After studying your forecasts, you want to see what impact the Discount variable has on your forecasts, so you assign the Discount variable to the independent variable role. SAS Forecast Studio uses this independent variable when generating the models, and as a result, the forecasts change slightly.

Figure 6.3 Forecast Values When Discount Is an Independent Variable

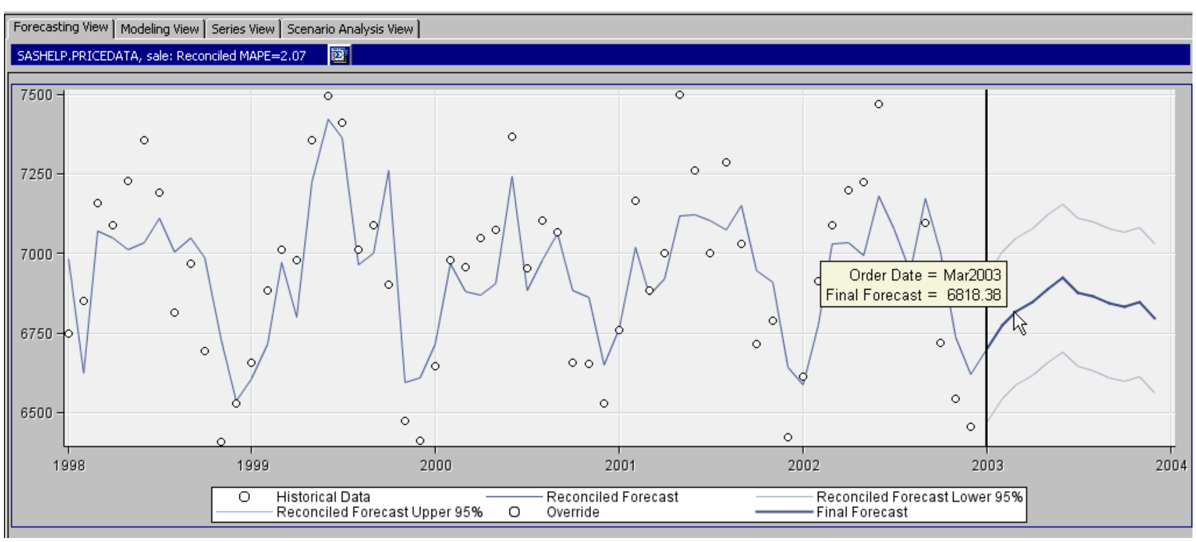

Reporting Variables

Reporting variables are the additional variables that you want to include in the reports. Unlike independent variables, reporting variables are not used to generate the forecasts.

For example, you might want to run a report that compares the unit cost to the sales for a given month. In this example, Sales is the dependent variable, and Cost is a reporting variable. The values of cost do not influence the forecasts.

Figure 6.4 Cost as a Reporting Variable

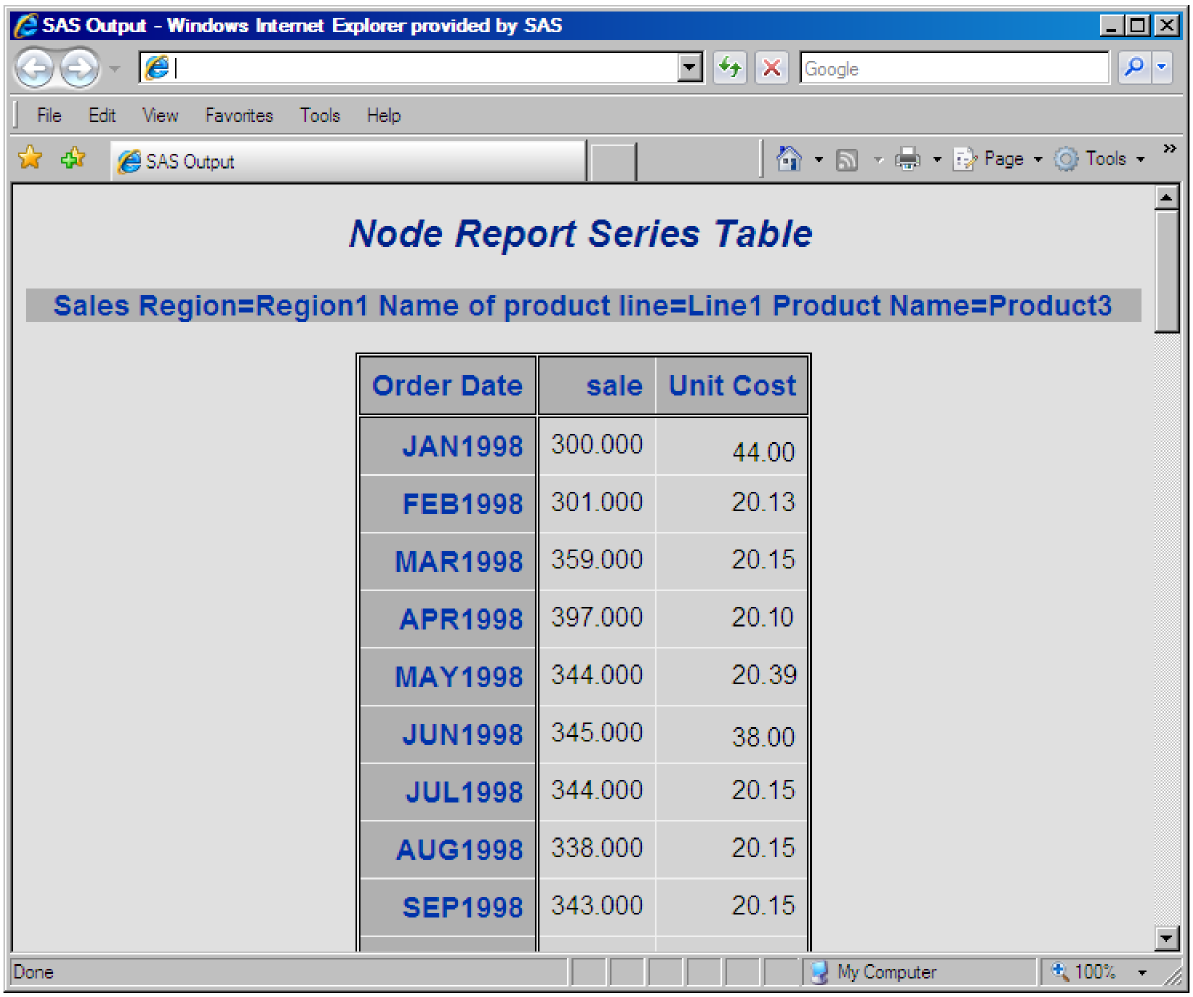

Adjustment Variables

What Is an Adjustment Variable?

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

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

- before generating the forecasts. After the time-stamped data has been accumulated and interpreted, you might need to adjust the time series before generating forecasts. By adjusting the time series for any known systematic variations or deterministic components, the underlying time series process can be more easily identified and modeled.

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

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

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

Add an Adjustment Variable

To add an adjustment variable, complete the following steps:

1. Select **Project** → **Hierarchy and Variable Settings**. The Hierarchy and Variable Settings dialog box opens.
2. Click **Adjustments**. The Adjustments dialog box opens.
3. Click **New**. The Adjustment Properties dialog box opens.
4. Select the dependent variable that you want to adjust.
5. Select the variable that you want to use in the adjustment. The variable that you assign to this role cannot have another role in the project.
6. (Optional) Specify a description of the adjustment variable.
7. Specify the operations to perform before and after forecasting. You can choose from minimum, maximum, add, subtract, multiply, and divide.
8. Click **OK**.

Examples of Adjustment Variables

Example 1: Adjusting Sales for the Exchange Rate

In the following example, you need to adjust your sales data for the exchange rate. Before generating the forecasts, SAS Forecast Studio divides the sales value by the exchange rate. After the forecasts are generated, SAS Forecast Studio adjusts the forecasts by multiplying the sales values by the exchange rate.

To create this adjustment variable, complete the following steps:

1. Select **Project** → **Hierarchy and Variables Settings**. The Hierarchy and Variables Settings dialog box opens.
2. Click **Adjustments**. The Adjustments dialog box opens.
3. Click **New**. The Adjustment Properties dialog box opens.
4. In the **Adjusted variable** list box, select the **Sales** dependent variable.
5. From the **Adjusting variable** drop-down list, select **Exchange Rate**.
6. In the **Description** text box, type `Adjusting a time series for exchange rate`.
7. From the **Pre-operation** drop-down list, select **Divide**.
8. From the **Post-operation** drop-down list, select **Multiple**.

Figure 6.5 Adjustment Properties Dialog Box for the Exchange Rate

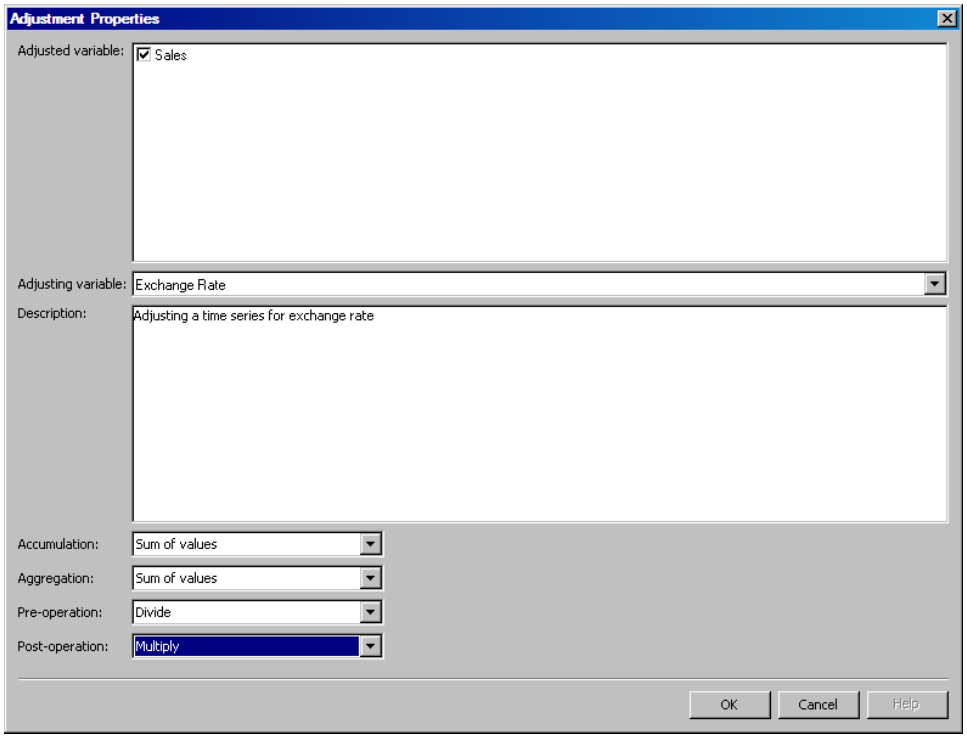

9. Click **OK**.

Example 2: Adjusting Sales for Advanced Bookings

In the following example, you need to adjust your sales data for any advanced bookings. Before generating the forecasts, SAS Forecast Studio subtracts the advanced bookings from the time series. After the forecasts are generated, SAS Forecast Studio adjusts the forecasts by adding the advanced bookings.

To create this adjustment variable, complete the following steps:

1. Select **Project** → **Hierarchy and Variable Settings**. The Hierarchy and Variable Settings dialog box opens.
2. Click **Adjustments**. The Adjustments dialog box opens.
3. Click **New**. The Adjustment Properties dialog box opens.
4. In the **Adjusted variable** list box, select the **Sales** dependent variable.
5. From the **Adjusting variable** drop-down list, select **Advanced Bookings**.
6. In the **Description** text box, type `Adjusting a time series for advanced bookings`.
7. From the **Pre-operation** drop-down list, select **Subtract**.
8. From the **Post-operation** drop-down list, select **Add**.

Figure 6.6 Adjustment Properties Dialog Box for the Advanced Bookings

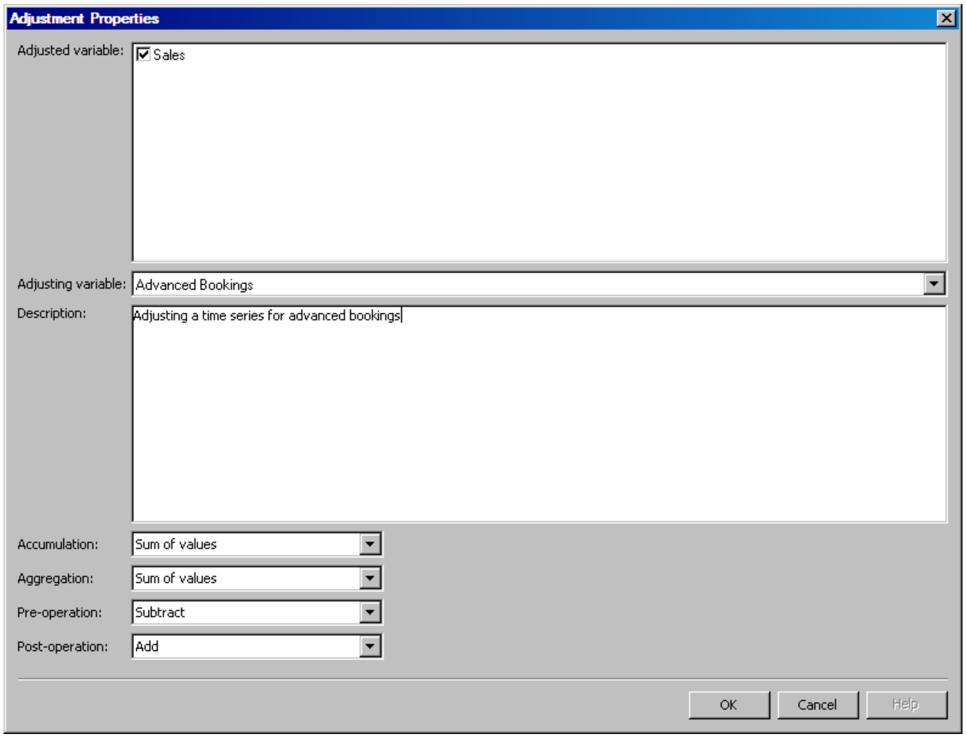

9. Click **OK**.

Chapter 7

Forecasting Hierarchically

Contents

Understanding Hierarchies

To forecast hierarchically in SAS Forecast Studio, you must assign at least one classification (BY) variable. When you assign a classification variable, then you can specify whether SAS Forecast Studio should forecast your data hierarchically. If you choose to forecast your data hierarchically, then SAS Forecast Studio sorts the dependent variable into individual series using the hierarchy specified by the BY variables. If you choose to forecast hierarchically, then only one dependent variable can be forecast.

For dependent, independent, adjustment, and reporting variables, you can specify how to aggregate the data across the levels of the hierarchy. You can also specify whether SAS Forecast Studio should reconcile the hierarchy and the reconciliation method to use.

Understanding Aggregation and Accumulation

Difference between Aggregation and Accumulation

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

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

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

Accumulation can be either of the following:

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

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

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

The Accumulation and Aggregation Methods

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

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

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

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

None
: does not accumulate the vector values.

Sum

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

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

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

Average

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

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

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

Minimum

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

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

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

Maximum

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

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

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

Median

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

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

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

Number of Nonmissing Observations

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

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

Number of Observations

accumulates the vector values based on the number of values.

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

Number of Missing Observations

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

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

First Occurrence

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

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

Last Occurrence

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

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

Standard Deviation

accumulates the vector values based on their standard deviation.

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

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

Uncorrected Sum of Squares

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

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

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

Corrected Sum of Squares

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

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

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

Understanding Reconciliation Methods

What Are the Reconciliation Methods?

When forecasting hierarchically in SAS Forecast Studio, you can choose from the following reconciliation methods.

Top-down

aggregates the data from the lowest levels in the forecast and then uses these values to generate the forecasts at the highest level. SAS Forecast Studio then uses this forecast and the disaggregation method that you specified to reconcile the forecasts for lower levels in the hierarchy.

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

Bottom-up

uses the data at the lowest level of the hierarchy to generate the forecasts. These forecasts are then used to reconcile the forecasts for the higher levels in the hierarchy.

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

Middle-out

aggregates the data from the lower levels and then uses these values to generate the forecasts for the middle level. SAS Forecast Studio uses the forecasts at the middle level to reconcile the forecasts for both the higher and lower levels. Some hierarchies have more than one middle level, so in SAS Forecast Studio, 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***.

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 7.1 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

Chapter 8

Working with Projects

Contents

What Is a Project?

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

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

Windows operating environments
!ROOT\SAS\Forecast Server\3.1*default-environment*\Projects*project-name*

UNIX operating environments

!ROOT/SAS/ForecastSever.3.1/*default-environment*/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*.

Creating a New Project

Overview of the New Project Wizard

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

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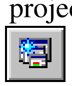 in the toolbar.

Step 1: Name Your Project

1. Specify a valid SAS name for the project. By default, the project name is Project*n*, where *n* is the lowest available integer value. The name must start with a letter and can contain an underscore. For more information about SAS naming conventions, see *SAS Language Reference: Concepts*.
2. (Optional) Provide a brief description of the project.
3. (Optional) To share this project, select the **Allow others to view and edit this project** check box. This option is not selected by default. If this check box is not selected, only the creator of the project can open it.
4. Click **Next**.

Step 2: Select the Data Set That You Want to Forecast

NOTE: Before you can use a data set in SAS Forecast Studio, it must be in the required format. For more information, see Chapter 5, "Preparing an Input Data Set for SAS Forecast Studio."

1. Double-click the name of the library to view the data sets within that library.
2. Select the data set that you want to use.

   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.
3. Click **Next**

Step 3: Specify Classification (BY) Variables and Whether to Forecast Hierarchically

1. Assign one or more variables to the BY variables role.
2. To forecast the data hierarchically, select the **Forecast a hierarchy using the above classification (BY) variables** check box. The hierarchy is created based on the order that the variables appear in the **Classification (BY) variables selected** box.

   Click **Preview** to preview the hierarchy.
3. To have SAS Forecast Studio reconcile the hierarchy after any changes, select the **Reconcile the hierarchy** check box and then select the reconciliation method to use.
4. To specify the advanced reconciliation settings for the hierarchy, click **Advanced**.
5. Click **Next**.

Step 4: Specify the Properties of the Time Dimension of Your Data

1. From the **Time ID variable** drop-down list, select the variable in the data set that contains the time dimension of your data.
2. (Optional) Specify the interval, multiplier, shift, seasonal cycle length, and format of the time ID variable.

By default, SAS Forecast Studio determines the time interval automatically. For more information about when SAS Forecast Server cannot detect the time interval, see "Troubleshooting the Time ID Variable" on page 249.

3. Click **Next**.

Step 5: Assign Roles to Variables in Your Data

1. Specify the roles of the remaining variables in your data set. If you are forecasting hierarchically, only one variable can be assigned to the dependent variable role.

2. Select the accumulation and aggregation methods for each variable. To use the same method for accumulation and aggregation, select the **Set accumulation to the value used for aggregation** check box. This option is available only if you are forecasting hierarchically.

3. Specify how the independent variables should be used in a system-generated model.

4. (Optional) To specify any adjustment variables, click **Adjustments**.

5. Click **Next**.

   NOTE: If you assign a variable to a role and the variable name is reserved, then an error message will appear. For more information about variable names, see Appendix B, "Reserved Names in SAS Forecast Studio."

Step 6: Specify How to Prepare the Data for Each Forecast

1. Specify how SAS Forecast Studio should interpret embedded missing values in the data.

2. Select which leading/trailing missing values SAS Forecast Studio should remove.

3. Select which leading/trailing zero values SAS Forecast Studio should interpret as missing.

4. (Optional) To specify a starting date for your forecasts, select the **Ignore data points earlier than the following date** check box. Then specify the starting date.

5. Click **Next**.

Step 7: Finish Creating Your Project

1. Specify the number of periods to forecast (also called the horizon). By default, the horizon is 12.

2. To change the diagnostics or model settings, click **Change other forecasting settings**. Specify the options that you want to change in the Forecasting Settings dialog box, and click **OK**.

3. To include an event, click **Create and import events**.

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

Step 8: Optional Project Actions

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 finish the wizard.

2. Click **Finish**.

You can monitor the progress of the forecasts from the status bar. The Forecast Summary dialog box opens after SAS Forecast Studio is finished. For more information, see "Reviewing a Summary of the Forecast Run" on page 74.

Open an Existing Project

You can open a project in the following ways:

- When you first invoke SAS Forecast Studio, you can open a project from the Projects dialog box.

- After SAS Forecast Studio is running, select **File** → **Projects** to open the Projects dialog box.

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

Migrating SAS Forecast Studio Projects

Update the Project Version

When you open a SAS Forecast Studio 2.1 project in SAS Forecast Studio 3.1, the Update Project Version dialog box opens.

When updating a project, follow these steps:

1. (Optional) Specify a name for the project. By default, SAS Forecast Studio uses the current name.

2. Specify how to update the statistical models for each series.
3. (Optional) Select the **Reset all series to "automatic selection"** check box.
4. (Optional) Specify whether the project should use any updated data, if it is available. If you select the **Use updated data if available** check box, then SAS Forecast Studio determines whether the data source has been updated, and, if so, updates the project data and reforecasts the existing series based on the option that you selected.
5. Click **OK**. SAS Forecast Studio automatically converts any 2.1 projects to be compatible with 3.1. If you try to open this project again in SAS Forecast Studio 2.1, the project might not work as expected.

Overview of the Changes to Projects in SAS Forecast Studio 3.1

You can open a SAS Forecast Studio 2.1 project in SAS Forecast Studio 3.1. When you open a 2.1 project in SAS Forecast Studio 3.1, you will see the following changes:

- When you migrate a project, most of the data sets in the SAS Forecast Studio project must be re-created using SAS High-Performance Forecasting 3.1. Consequently, a SAS Forecast Studio 2.1 project might take longer to open in SAS Forecast Studio 3.1. When the 2.1 project is opened in 3.1, the project is converted to the 3.1 format, so subsequent openings of the project in SAS Forecast Studio 3.1 will be faster.
- SAS Forecast Studio 3.1 enables you to filter on the sum of the series. This property was not available in previous releases of SAS Forecast Studio.

Updating Your Project Data

The data for the project might have changed since the last time that you forecast your data. For example, the data might have included a BY group, such as REGION=1 > LINE=1 > PRODUCT = 1. However, in the latest version of the data, this BY group does not exist.

When you reforecast a project, you can choose whether to use the updated version of the data. In the previous example, if you use the latest version of the data, then you do not see this BY group in the output data sets for the project.

To keep track of your forecasts and any BY groups that have been removed from the data, it is recommended that you archive your project before reforecasting the project using any updated data. For more information about archiving a project, see *SAS Forecast Server Administrator's Guide*.

For more information about how to update your project data, see "Reforecast a Project" on page 76.

Viewing Properties

View Project Properties

To view the properties for a project, complete the following steps:

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

View Series Properties

To view the properties for a series, complete the following steps:

1. Select the series in the **Hierarchy View**, **Table View**, or **Failed Forecasts** tab.
2. Select **View** → **Series Properties**. The series properties appear at the bottom of the tab.

For example, the following figure shows the series properties for the Region1:Line1:Product1 node in the hierarchy.

Figure 8.1 Series Properties for Node in Hierarchy View

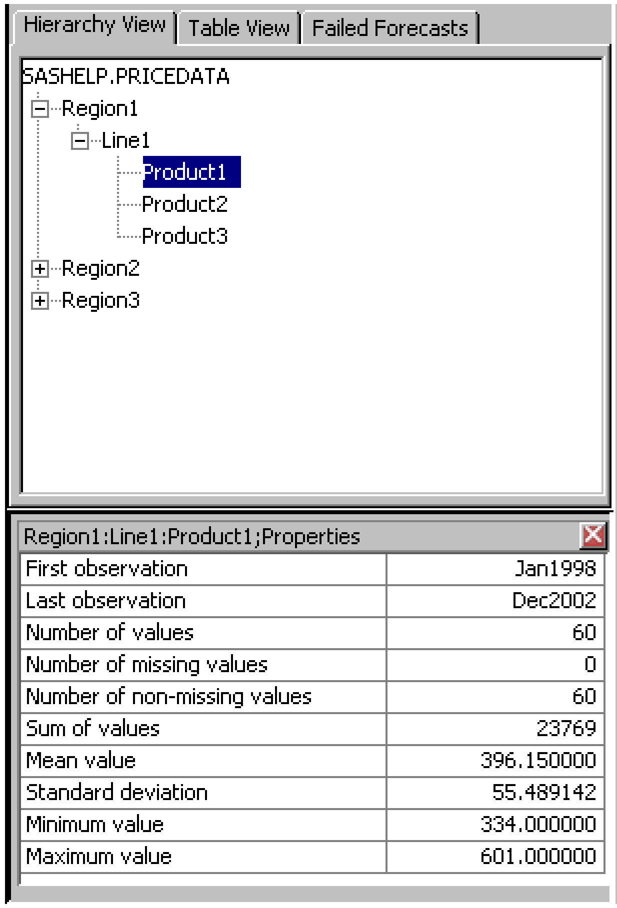

Save a Copy of the Project

To save a copy of the project, complete the following steps:

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

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

The new project is saved to the SAS Forecast Studio directory. The current project is closed, and the new project is automatically opened in SAS Forecast Studio.

Part III

Analyzing Your Results

Chapter 9

Generating Forecasts

Contents

Process for Generating a Forecast

SAS Forecast Studio uses the following forecasting processes:

1. Accumulation: The time-stamped data is accumulated into time series data based on a particular frequency.

2. Interpretation: Missing values and zero values are interpreted before the time series is analyzed.

3. Adjustment: If required, the time series is adjusted before the analysis. After the analysis, the forecast might require a post-forecast adjustment to return forecasts to the original metric.

4. Diagnose: The series are analyzed for trends, seasonality, and cyclicity. The series diagnostics provide recommendations about whether the forecasting model should contain a log transform, trend terms, or seasonal terms. For example, identifying trends in the series can help you determine whether a dependent series requires simple and seasonal differencing. These diagnostics also help determine whether the time series is intermittent.

   SAS Forecast Studio performs the following diagnostic tests:

   a) The intermittent test computes the average time interval between demands. If the average time interval is greater than the specified limit, then an intermittent forecasting model is used.

b) A seasonality test is also performed. A seasonal dummy model with AR(1) errors is fit, and the joint significance of the seasonal dummy estimates is tested. If the seasonal dummies are significant, the AIC statistic for this model is compared to the AIC statistic for the AR(1) model without seasonal dummies.

For more information about these tests, see "Set the Diagnose Properties for Each Series" on page 72.

Any system-generated models that are appropriate for the series are created. Any independent variables or events are included in the system-generated ARIMA and unobserved components models. These variables and events are used in the models if they improve performance.

5. Select: All models are evaluated against the data. Using the model selection criterion, SAS Forecast Studio selects the best-performing model.
6. Fit: The best-performing model is fit to the data to estimate a final set of parameters.
7. Forecast: SAS Forecast Studio uses the model parameter estimates, the model residual variance, and the full range of data to generate the forecasts (predictions, prediction standard errors, prediction errors, and confidence limits). After forecasts are produced for each series, the hierarchy is reconciled.
8. Evaluation: The in-sample statistics of fit are computed-based on the one-step-ahead forecasts and the actual data. These statistics help identify poorly fitting models.

When you change the forecast settings or properties of a project, SAS Forecast Studio automatically regenerates the forecasts. When you choose to reforecast a series or project or you choose to update the project data, you can select where to start in the forecasting process. If SAS Forecast Studio performs a nondestructive diagnose, then any user-created models are preserved, and any system-generated models are re-created. A destructive diagnose re-creates all models.

Set the Diagnose Properties for Each Series

To set the diagnose properties for each series, complete the following steps:

1. Select **Project** → **Forecasting Settings**. The Forecasting Settings dialog box opens.
2. In the selection pane, click **Diagnostics**.
3. In the **Perform intermittency test** box, specify a number greater than 1 that is used to determine whether a time series is intermittent. This option is selected by default, and the default value is 2. To specify a different sensitivity, type a new value in the **Sensitivity** box.
4. Specify the options for performing 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.

5. Specify the minimum number of seasonal cycles for a seasonal model. No seasonal model is fit to a series that contains fewer observations than *n* multiplied by the seasonal length.
6. Specify the minimum number of observations for a trend model. No trend model is fit to a series that contains fewer than *n* nonmissing values. Incorporation of a trend is checked only for the following models:
   - simple smoothing models
   - unobserved component models that do not include a slope component
   - ARIMA models with no differencing of the dependent variable
7. Specify the minimum number of observations for a non-mean model. Any series with fewer than *n* nonmissing values are not fit using the models in the selection list. Instead, these models are forecast as the mean of the observations in the series.
8. Select the functional transformation to apply to the dependent variable.
9. Specify how SAS Forecast Studio should diagnose the independent variables.
10. Specify how to calculate the forecasts for the system-generated models. Forecasts can be based on the mean or median.
11. Specify the criteria for detecting outliers in ARIMA models.
12. Click **OK**.

Specify the Number of Periods to Forecast

To specify how to fit the forecast model for each series, complete the following steps:

1. Select **Project** → **Forecasting Settings**. The Forecasting Settings dialog box opens.
2. In the selection pane, click **Forecast**.
3. Specify 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.
4. Specify whether to calculate the statistics of fit over an out-of-sample range.
5. Specify the confidence level for the series. By default, this confidence level is 0.05, which is a 95% confidence limit.
6. To allow negative forecasts, select the **Allow negative forecasts** check box.

   NOTE: If you clear the **Allow negative forecasts** check box, then any negative values in the forecast model are set to 0, and any overrides with negative values are removed.

Reviewing a Summary of the Forecast Run

After you create a project, the forecast summary is available by default. When you reforecast your project, this summary does not automatically open. To view the forecast summary, select **Project** → **Forecast Summary**.

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

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

To view more information about a graph or table, double-click the graph or table to open it in a separate window. From this dialog box, you can also view the failed forecasts or forecast the project.

Here is an example of the Forecast Summary dialog box:

Figure 9.1 Forecast Summary Dialog Box

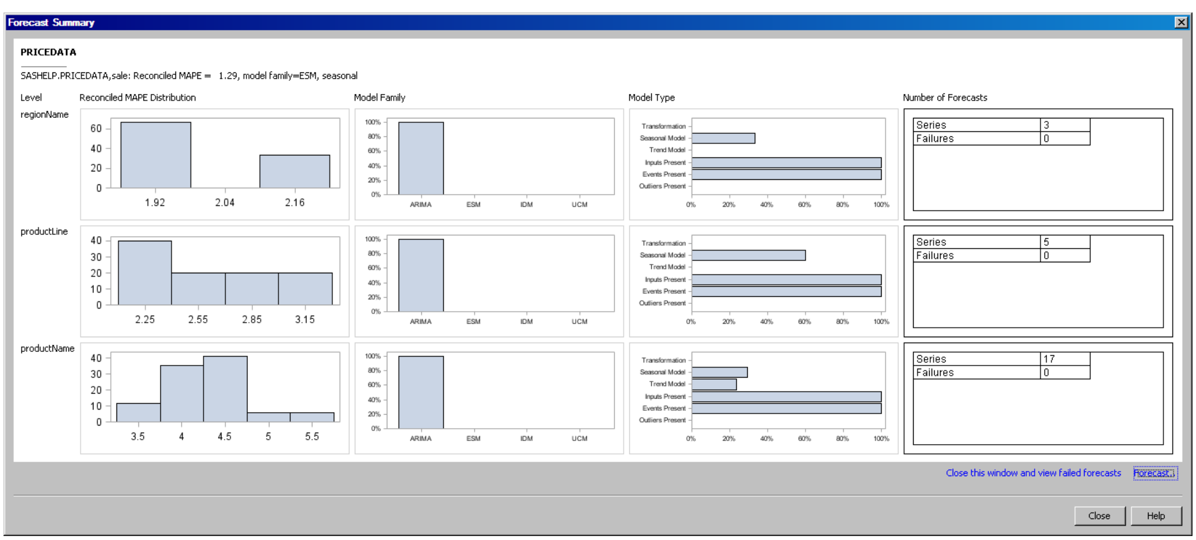

Specifying an Inactive Series

What Is an Inactive Series?

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

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

- Any child series must also be inactive. SAS Forecast Studio prompts you to confirm making these child series inactive when you clear the **Active series** check box in the Forecasting View. (You can also make a series inactive by clearing the **Active** check box in the Modeling View.)
- Any overrides that you have specified for this series are removed, and you cannot add new overrides to an inactive series.
- A "Reconciliation needed" message appears above the Forecasting View. You can click **Reconcile** to reconcile the hierarchy.
- In the Modeling View, no forecast model is selected.

Set a Series as Inactive

To specify an inactive series, complete the following steps:

1. In the hierarchy or table view, select the series that you want to make inactive.
2. Clear the **Active series** check box in the data table of the Forecasting View.

A series remains inactive even after SAS Forecast Studio rediagnoses the project. To make the series (and any child series) active again, select the **Active series** check box.

Reforecasting a Series or Project

Reforecast a Series

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

To reforecast the selected series, complete the following steps:

1. Select the series in the hierarchy or table view and select **Series** → **Forecast Series**. The Forecast Series dialog box opens.
2. Specify how to forecast the series.
3. (Optional) To have SAS Forecast Studio select the best-performing model, select the **Reset the series to "automatic selection"** check box. If you select this option, any models that you might have selected for a series are ignored. For more information about how you select a model, see "Specifying the Forecast Model" on page 181.
4. Click **OK** to reforecast the series.

Reforecast a Project

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

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

1. Select **Project** → **Forecast Project**. The Forecast Project dialog box opens.
2. Specify how to forecast the project.

   NOTE: If you choose any option other than **Diagnose: create new generated models, refit all models, and select a model**, then any new series are forecast using the model selection list that ships with SAS Forecast Studio. For more information, see Appendix E, "Default Model Selection Lists."
3. (Optional) To have SAS Forecast Studio select the best-performing model, select the **Reset all series to "automatic selection"** check box. If you select this option, any models that you might have selected for a series are ignored. For more information about how you select a model, see "Specifying the Forecast Model" on page 181.
4. (Optional) Select the **Use updated data if available** check box. SAS Forecast Studio determines whether the data source has been updated, and if so, updates the project data and reforecasts the existing series based on the new data.
5. Click **OK** to reforecast the project.

Chapter 10

Working with Filters

Contents

What Is a Filter?

Filters help you identify forecasts that meet a specific criterion. If you forecasted hierarchically, then SAS Forecast Server automatically generates a filter for each level in the hierarchy. For example, if the hierarchy for your data is Region > Product Line > Product, then the system-generated filters are the following:

- Region_level - displays all series at the Region level
- Productline_level - displays all series at the Product Line level
- Product_level - displays all series at the Product level

You cannot edit or delete these system-generated filters.

In SAS Forecast Studio, you can create additional filters. For an example of how to create a filter, see "Examples of Filters" on page 79.

To view the results from a filter, select the **Table View** tab in the overview panel.

Create a Filter

To create a filter, complete the following steps:

1. Select **Tools** → **Filters**. The Filters dialog box opens.
2. Click **New**. The Filter Properties dialog box opens.
3. Specify a valid SAS name. You can also specify a description for the filter.
4. Specify the hierarchy level for the filter. By default, the filter is applied to all the levels in the hierarchy. If you select a specific level in the hierarchy, then the results from the filter contain only series in that level of the hierarchy.

   NOTE: The **Hierarchy level** drop-down list is available only if you forecast your data hierarchically.
5. Specify the filter by using the columns in the table.
6. (Optional) To add additional filter expressions, click .
7. Specify whether the series must have overrides, notes, or both.
8. Click **OK**. The new filter now appears in the Filters dialog box.

Edit a Filter

To edit an existing filter, complete the following steps:

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

Delete a Filter

To delete a filter, complete the following steps:

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

View the Results from a Filter

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

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

Examples of Filters

Example 1: Creating a Simple Filter

In this example, you want to know which series have a Mean Absolute Percent Error (MAPE) greater than 3.

To create this filter, complete the following steps:

1. Select **Tools** → **Filters**. The Filters dialog box opens.
2. Click **New**. The Filter Properties dialog box opens.
3. In the **Name** box, type `MAPE`.
4. In the **Description** box, type `Filter to identify which series have a MAPE greater than 3`.
5. From the **Hierarchy level** drop-down list, select **All**.
6. In the table, specify the following values for each column:
   - In the **Category** column, select **Reconciled (Rec.) statistic of fit**.
   - In the **Property** column, select **Mean absolute percent error**.
   - In the **Condition** column, select **>**.
   - In the **Value** column, type `3` and press ENTER.

The final window should look like the following:

Figure 10.1 Completed Filter Properties Dialog Box for Filter Example

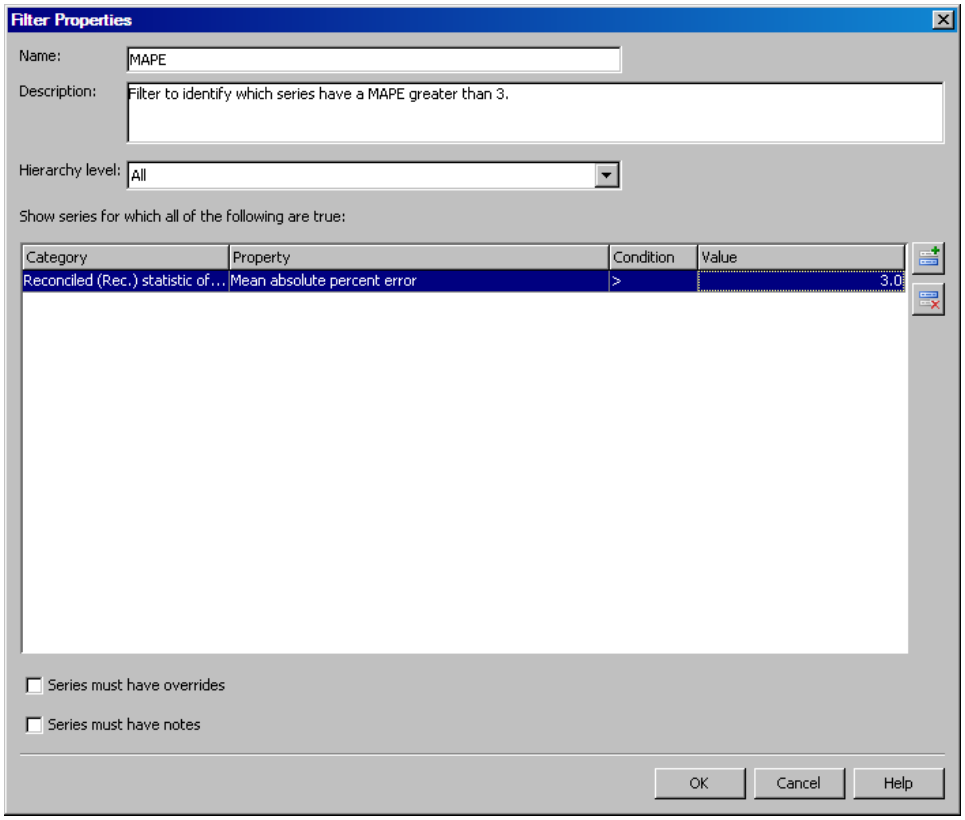

7. Click **OK** to add the filter to the list in the Filters dialog box.

Figure 10.2 New Filter in Filters Dialog Box

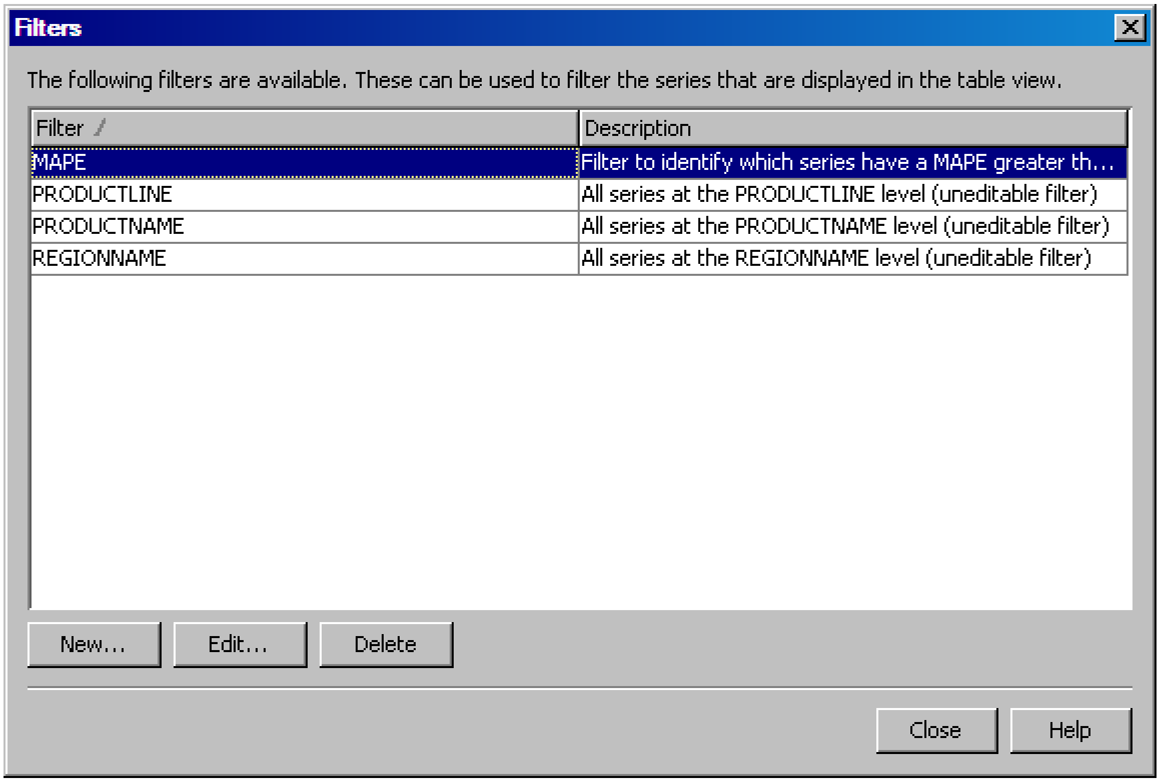

8. Click **Close** in the Filters dialog box to return to the project.

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

1. In the overview panel, select the **Table View** tab.
2. From the **Filter** drop-down list, select **MAPE**.

Here are the results:

Figure 10.3 Results in the Table View after Applying the MAPE Filter

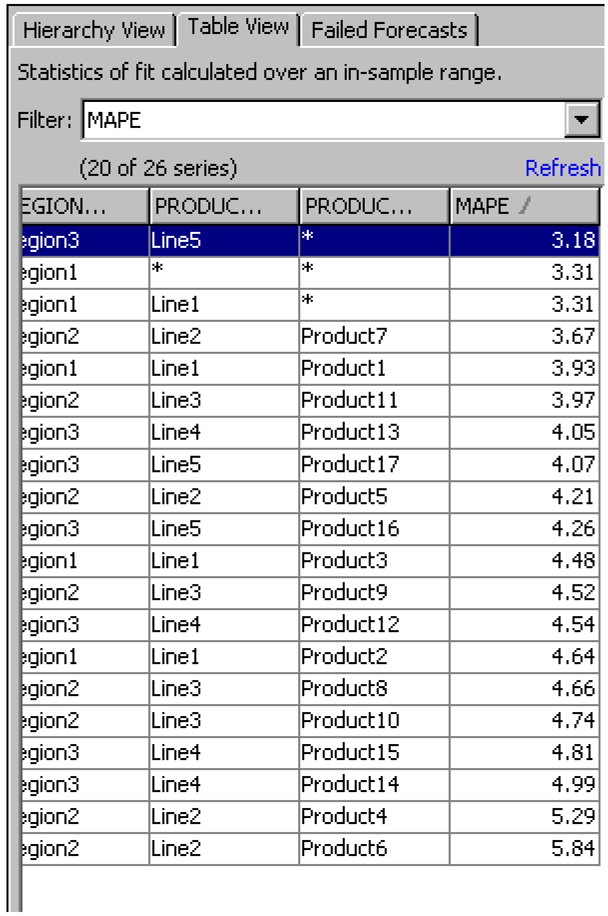

EGION...	PRODUC...	PRODUC...	MAPE
gion3	Line5	*	3.18
gion1	*	*	3.31
gion1	Line1	*	3.31
gion2	Line2	Product7	3.67
gion1	Line1	Product1	3.93
gion2	Line3	Product11	3.97
gion3	Line4	Product13	4.05
gion3	Line5	Product17	4.07
gion2	Line2	Product5	4.21
gion3	Line5	Product16	4.26
gion1	Line1	Product3	4.48
gion2	Line3	Product9	4.52
gion3	Line4	Product12	4.54
gion1	Line1	Product2	4.64
gion2	Line3	Product8	4.66
gion2	Line3	Product10	4.74
gion3	Line4	Product15	4.81
gion3	Line4	Product14	4.99
gion2	Line2	Product4	5.29
gion2	Line2	Product6	5.84

Example 2: Creating a Filter with Multiple Expressions

In this example, you want to know which series have a Mean Absolute Percent Error (MAPE) greater than 3 and a Minimum value (Min) less than 300.

To create this filter, complete the following steps:

1. Select **Tools** → **Filters**. The Filters dialog box opens.
2. Click **New**. The Filter Properties dialog box opens.
3. In the **Name** box, type `MAPE7MIN300`.
4. In the **Description** box, type `Filter to identify which series have a MAPE greater than 3 and a Minimum value less than 300`.
5. From the **Hierarchy level** drop-down list, select **All**.
6. In the table, specify the following values for each column:

- In the **Category** column, select **Reconciled (Rec.) statistic of fit**.
- In the **Property** column, select **Mean absolute percent error**.
- In the **Condition** column, select **>**.
- In the **Value** column, type `3` and press ENTER.

7. Click . A new row appears in the Filter Properties dialog box. SAS Forecast Studio automatically joins the filter conditions by AND.

8. In the new row, specify the following values for each column:

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

   The final window should look like the following:

Figure 10.4 Completed Filter Properties Dialog Box for Filter Example

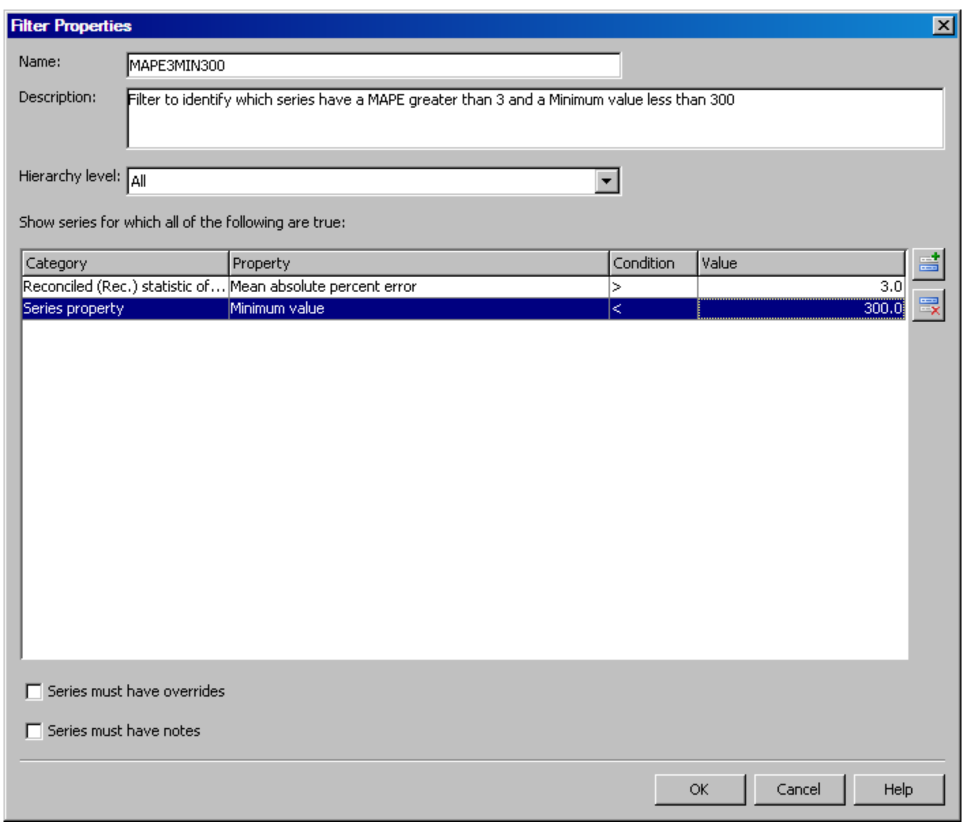

9. Click **OK** to add the filter to the list in the Filters dialog box.

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

The filter that you created is now available from the table view in the overview panel. From the **Filter** drop-down list, select **MAPE3MIN300** to view the series that meet these filter criteria.

Figure 10.5 Results in the Table View after Applying the MAPE3MIN300 Filter

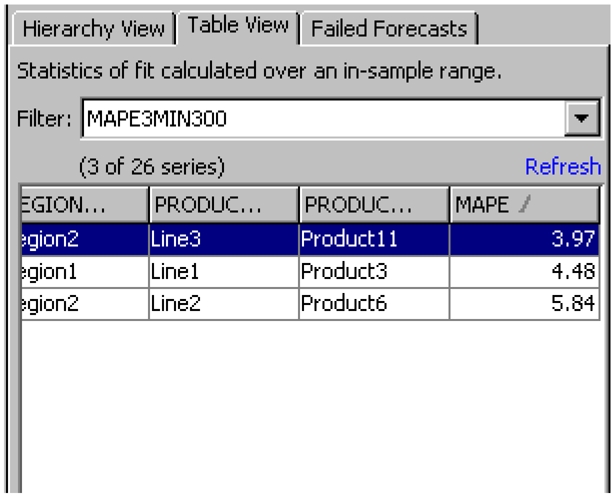

Example 3: Creating a Filter for Missing Values

In this example, you want to know which series have missing values. In SAS Forecast Studio, missing values are represented by a . (a period).

To create this filter, complete the following steps:

1. Select **Tools** → **Filters**. The Filters dialog box opens.
2. Click **New**. The Filter Properties dialog box opens.
3. In the **Name** box, type `Missing_Values`.
4. In the **Description** box, type `Filter to identify which series have missing values`.
5. From the **Hierarchy level** drop-down list, select **All**.
6. In the table, specify the following values for each column:
   - In the **Category** column, select **Reconciled (Rec.) statistics of fit**.
   - In the **Property** column, select **Mean absolute percent error**.
   - In the **Condition** column, select **=**.
   - In the **Value** column, type `.` and press ENTER.

The final window should look like the following:

Figure 10.6 Completed Filter Properties Dialog Box for Filter Example

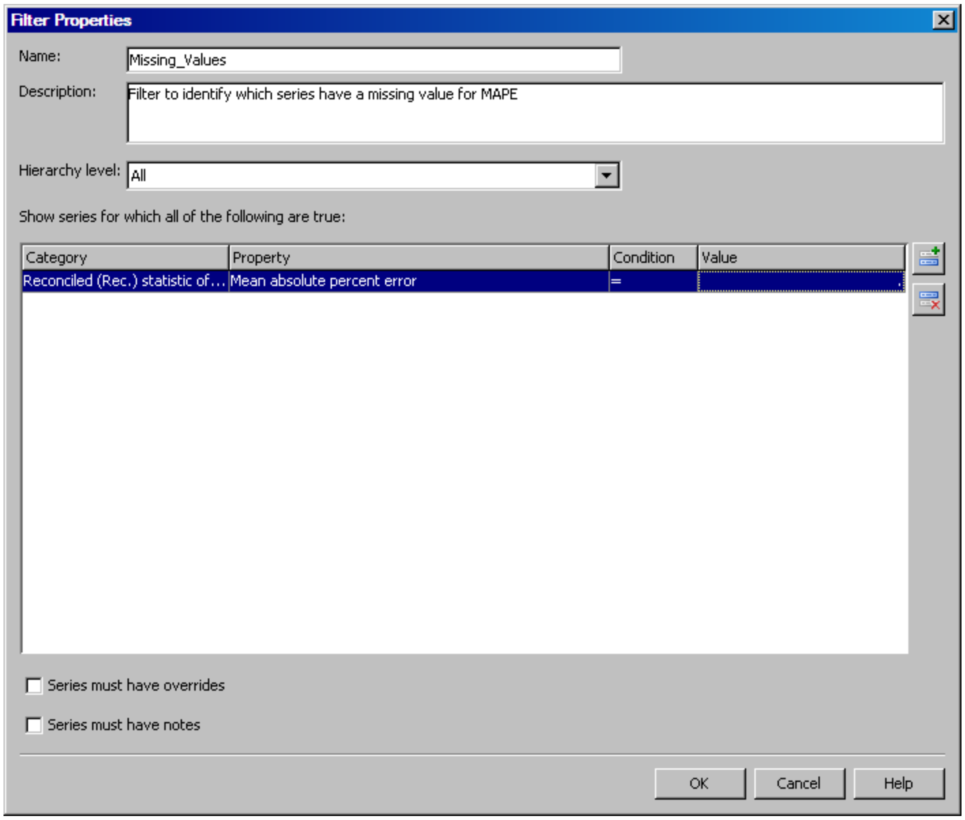

7. Click **OK** to add the filter to the list in the Filters dialog box.
8. Click **Close** in the Filters dialog box to return to the project.

Chapter 11

Running Reports

Contents

Overview of Reports

You can use reports to run custom SAS code against your project data. 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.

When you run a report, you can specify the output format as HTML, PDF, RTF, PS, CSV, or XML. You can also select a style for the output. By default, SAS Forecast Studio lists several available styles. However, you can specify any style in the SAS Output Delivery System (ODS).

Sample Reports in SAS Forecast Studio

SAS Forecast Studio ships with a collection of 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 G, "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 88.

NOTE: Some sample reports, included for illustration purposes, require SAS products (such as SAS OLAP Server) that are not directly licensed with SAS Forecast Server. If your site has licenses for these additional SAS products, then you can run these sample reports. If your site does not have licenses for these procedures, then these reports will not run properly. Warnings or errors appear in the SAS log when a procedure for an unlicensed SAS product is required. For more information about what is available at your site, contact your site administrator.

Create a Report

Although SAS Forecast Studio ships with sample reports, you can create your own reports. To create a report, you must write the SAS code for the report and then register the code as a stored process in SAS Management Console.

For more information about how to write a stored process and make it available in SAS Forecast Studio, see "Additional Administration Tasks" in the *SAS Forecast Server Administrator’s Guide*.

Generating a Report

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. Select the format and style of the output.

4. Click **Run**. If the stored process code that generates the report includes any parameters that need to be customized at run time, a dialog box opens and prompts the user. 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: If the report you select requires information that the current project does not have, then an error message appears and the report does not run. For example, you cannot run a report that requires a reporting variable if you have not assigned any variables to the reporting role.

Example: Generating a Report on the Final Forecasts for a Selected Node

To improve the accuracy of your forecasts, you added overrides to the Jan2003, Jun2003, Dec2003 time periods for the top node in the hierarchy. You want to generate a report that shows the impact of these overrides on the final forecasts.

To generate this report, you complete the following steps:

1. Select **Tools** → **Reports and Stored Processes**. The Reports and Stored Processes dialog box opens.

2. Expand the Reports, Samples, and Final Forecast nodes. In the Final Forecast node, select **Node Final Forecast Table** and click **Run**. By default, reports are generated in HTML.

Figure 11.1 Selected Report in Reports and Stored Processes Dialog Box

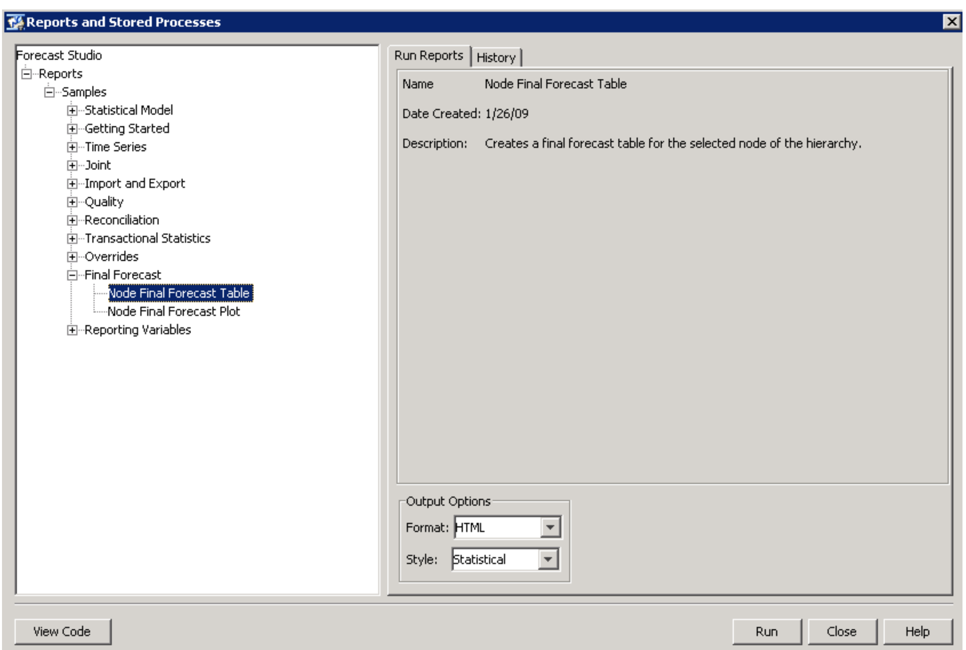

3. In the prompt dialog box, click **OK** to accept the default values for the prompts.

Figure 11.2 Example of a Prompt Dialog Box

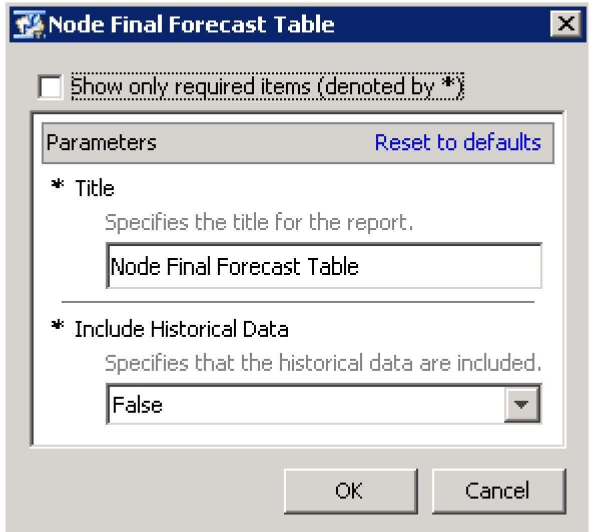

A report (similar to the following) automatically opens in your Web browser.

Figure 11.3 Example of the Output from a Report

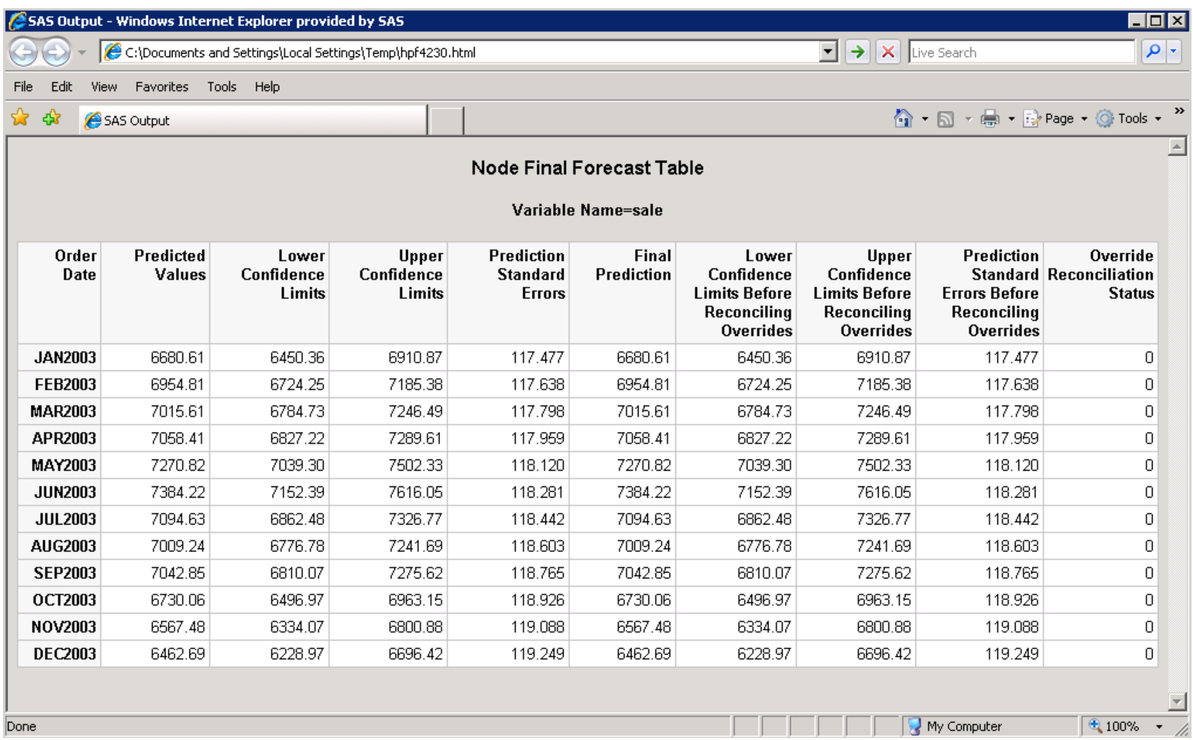

Node Final Forecast Table

Variable Name=sale

Order Date	Predicted Values	Lower Confidence Limits	Upper Confidence Limits	Prediction Standard Errors	Final Prediction	Lower Confidence Limits Before Reconciling Overrides	Upper Confidence Limits Before Reconciling Overrides	Prediction Standard Errors Before Reconciling Overrides	Override Reconciliation Status
JAN2003	6680.61	6450.36	6910.87	117.477	6680.61	6450.36	6910.87	117.477	0
FEB2003	6954.81	6724.25	7185.38	117.638	6954.81	6724.25	7185.38	117.638	0
MAR2003	7015.61	6784.73	7246.49	117.798	7015.61	6784.73	7246.49	117.798	0
APR2003	7058.41	6827.22	7289.61	117.959	7058.41	6827.22	7289.61	117.959	0
MAY2003	7270.82	7039.30	7502.33	118.120	7270.82	7039.30	7502.33	118.120	0
JUN2003	7384.22	7152.39	7616.05	118.281	7384.22	7152.39	7616.05	118.281	0
JUL2003	7094.63	6862.48	7326.77	118.442	7094.63	6862.48	7326.77	118.442	0
AUG2003	7009.24	6776.78	7241.69	118.603	7009.24	6776.78	7241.69	118.603	0
SEP2003	7042.85	6810.07	7275.62	118.765	7042.85	6810.07	7275.62	118.765	0
OCT2003	6730.06	6496.97	6963.15	118.926	6730.06	6496.97	6963.15	118.926	0
NOV2003	6567.48	6334.07	6800.88	119.088	6567.48	6334.07	6800.88	119.088	0
DEC2003	6462.69	6228.97	6696.42	119.249	6462.69	6228.97	6696.42	119.249	0

Part IV

Improving Your Forecasts

Chapter 12

Working with Events

Contents

Understanding Events

What Is an Event?

The most likely problem with your initial forecast model is that the model ignores a known event. You can improve the accuracy of your model by adding events. When you add or delete an event in a project, SAS Forecast Studio must reforecast the project in order for the models to reflect these changes.

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

For example, retail sales data follows a fairly steady pattern depending only on the day of the week. However, when you include events such as a one-time New Year's Day promotion, the forecasting model can predict a temporary increase in sales and then return to normal sales levels after the event. 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.

You can use SAS Forecast Studio to predict the effect of a future event. Before the forecasting model can predict the effects of a future event, your data must include the effects of past occurrences of the event. For example, after reviewing your December sales data in the United States, you decide to create an event for next Christmas. SAS Forecast Studio must have data that explains how retail sales were impacted by previous Christmas events. Without this data, SAS Forecast Studio cannot predict the effect of the next Christmas event on retail sales.

In SAS Forecast Studio, the following components can define an event:

- the event type
- the occurrence
- the recurrence

Event Types

In SAS Forecast Studio, you can select from four types of events.

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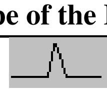	Temporary change in the magnitude of a time series process. The magnitude returns to the former level immediately after the change.
level shift	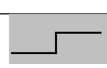	Persistent change in the level of a time series process.
ramp	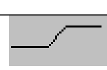	A sudden growth or decay in the trend or slope of a time series process. This change in growth could persist or end when the time series reaches a new level.

Event Type	Shape of the Event	Description
temporary change	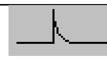	Temporary change in the magnitude. The magnitude decays to the former level after the change.

Event Occurrence

The occurrence is the date of the event. For each occurrence, you can specify the duration and the time period of the occurrence. For example, you want to create an event to reflect sales of candy for Halloween. In the United States, Halloween is celebrated on October 31, so you specify the date of the event as October 31. However, you anticipate higher than normal sales in the days before October 31. Therefore, you specify that the time of the occurrence is seven days before October 31.

Event Recurrence

The recurrence is how frequently the event occurs and how long the recurrence should last. The event can recur at the same interval or at a longer interval as the interval for the time ID variable. You cannot specify a recurrence if you have selected two or more occurrences or if you select a holiday as the occurrence.

For example, your company has an annual 25% sale during the month of June. You create an occurrence for the month of June. Your forecasting model needs to reflect this yearly increase in sales. You create an indefinite recurrence that occurs before and after the selected occurrence date. Now, your model can account for the increase in sales for your historical data and for any future data that you record.

Creating New Events

Create an Event

To create a new event, complete the following steps:

1. Select **Project** → **Event Repository**. The Event Repository dialog box opens.
2. Click **New**. The Event Properties dialog box opens.
3. Specify a valid SAS name for the event. You can also specify a description for the event.

4. Click **Change** to specify an event type. For more information about these event types, see "What Is an Event?" on page 93.

5. From the **Usage in system-generated models** drop-down list, specify whether an event is used for all of the models that SAS Forecast Studio generates. For more information, see "How Events Are Used in Model Generation" on page 106.

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

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

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

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

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

9. Click **OK** to save the event. The new event appears in the Event Repository dialog box.

Example 1: Creating an Event for Annual Promotion

Your company holds its annual sale during the month of June. Historical data shows that sales increase significantly during this month, so when forecasting your data, you want SAS Forecast Studio to try to include this event in model generation. The recurrence option in SAS Forecast Studio enables you to create an event and specify how frequently this event will occur. The following example shows how to create a pulse event for this annual promotion.

To create an event for the June promotion, you complete the following steps:

1. Select **Project** → **Event Repository**. The Event Repository dialog box opens.

2. Click **New**. The Event Properties dialog box opens.

3. In the **Name** box, type `Annual_Sale`.

4. In the **Description** box, type `Event for annual sale in June`.

5. To add an occurrence, click **Add** to add an occurrence. The Select Occurrences dialog box opens.

6. From the **Month** drop-down list, select **June** and click  to select this event.

Figure 12.1 Select Occurrences Dialog Box

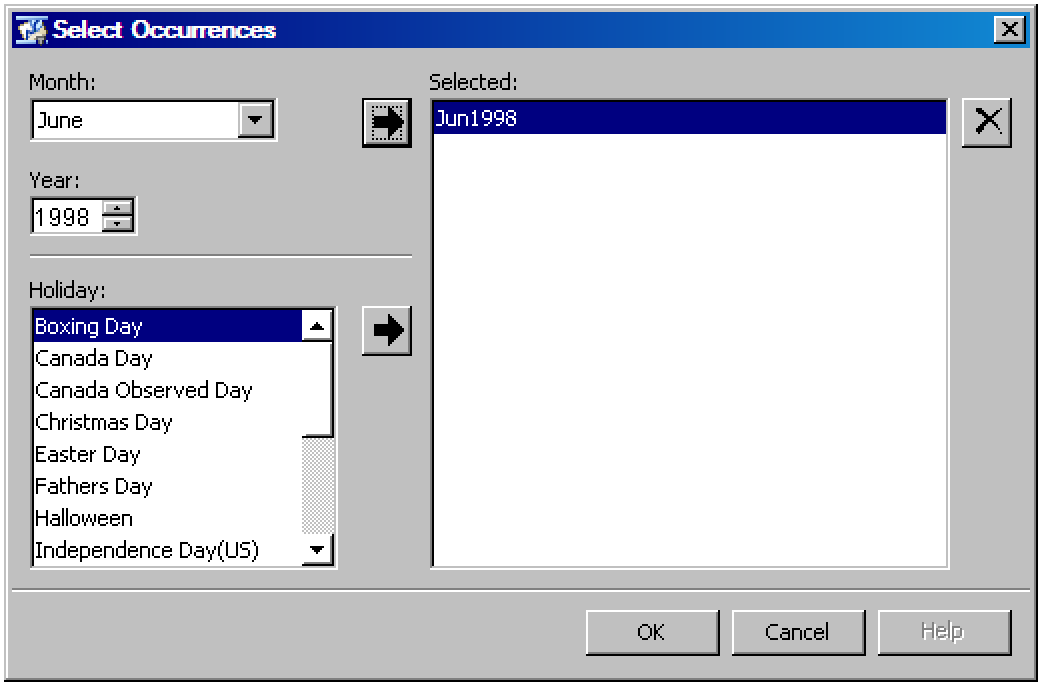

Click **OK**. The **Occurrence** box in the Event Properties dialog box now lists Jun1998 as the occurrence of this event.

7. To specify the recurrence, click **Edit**. The Event Recurrence dialog box opens.

8. From the **Select how frequently the event will recur** drop-down list, select **Year**. To specify how long the recurrence will last, select **Indefinitely starting with and after the selected occurrence date**.

Figure 12.2 Event Recurrence Dialog Box

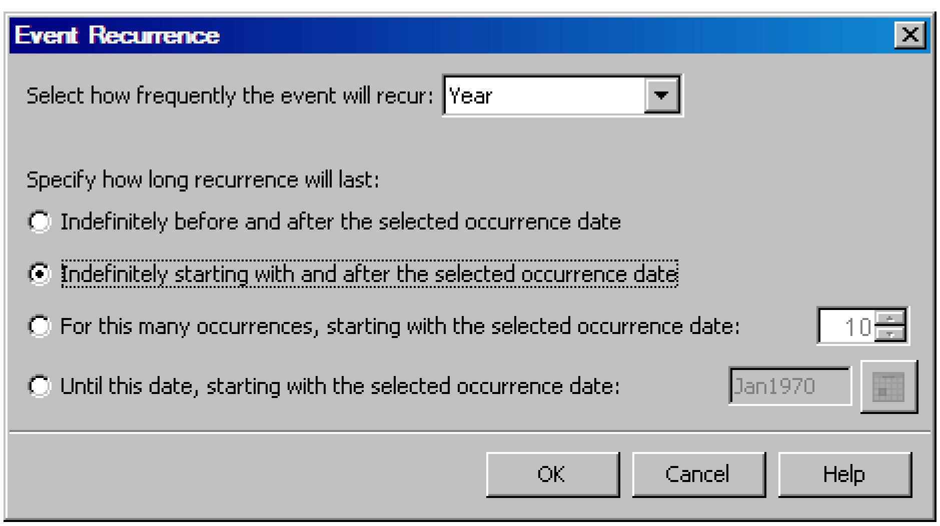

Click **OK**. The **Recurrence** box in the Event Properties dialog box now lists Year as the recurrence of the event.

9. From the **Usage in system-generated models** drop-down list, select **Try to use**, so SAS Forecast Studio will try to use this event in all system-generated models.

   The Event Properties dialog box looks like the following:

Figure 12.3 Event Properties for an Annual Event

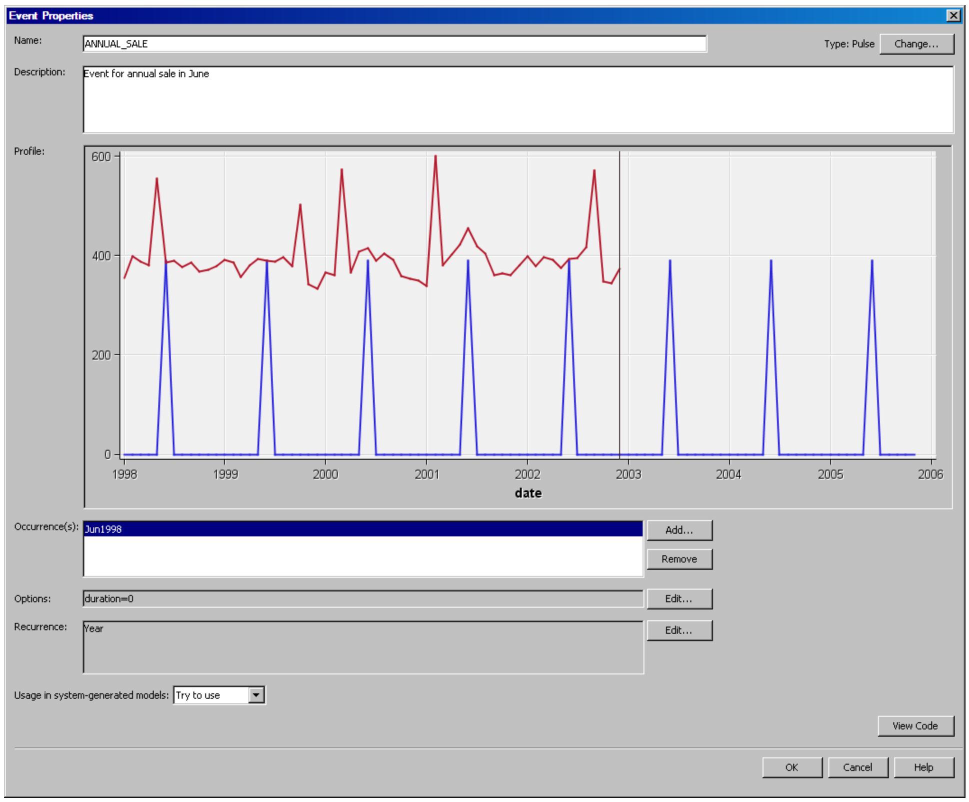

10. Click **OK** to create the event.

Example 2: Creating an Event for a Holiday

In the United States and Europe, your company has an increase in sales before the Christmas holiday. The following example shows how to create a pulse event for a holiday.

To create an event for a holiday, you complete the following steps:

1. Select **Project** → **Event Repository**. The Event Repository dialog box opens.
2. Click **New**. The Event Properties dialog box opens.
3. In the **Name** box, type `Holiday_Sale`.

4. To add an occurrence, click **Add** to add an occurrence. The Select Occurrences dialog box opens.

5. From the list of predefined holidays, select **Christmas** and click  to select this event.

   Click **OK**. The **Occurrence** box in the Event Properties dialog box now lists Christmas Day as the occurrence of this event.

6. To specify options for the event, click **Edit** next to the **Options** box. The Event Options dialog box opens.

7. To specify that sales will increase before the event, select **Relative to the specified period** and select **Earlier** from the **Direction of offset** drop-down list.

Figure 12.4 Event Options Dialog Box

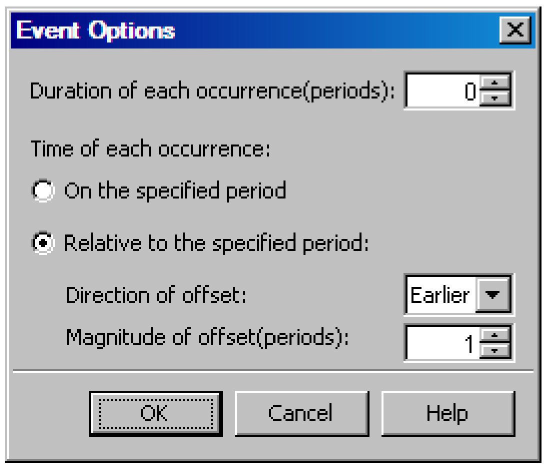

Click **OK**. The **Options** box in the Event Properties dialog box now lists the timing and duration of the event.

The Event Properties dialog box looks like the following:

Figure 12.5 Event Properties for a Holiday Event

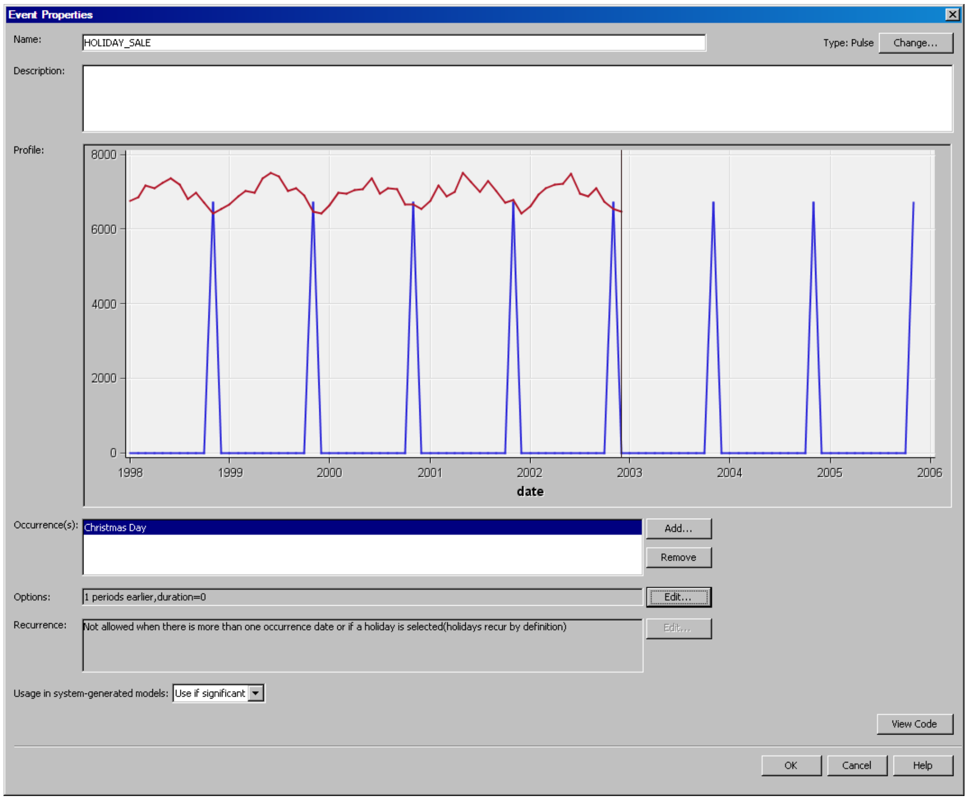

8. Click **OK** to create the event.

Understanding Combination Events

Combine Events

To combine events, complete the following steps:

1. Select **Project** → **Event Repository**. The Event Repository dialog box opens.
2. Click **Combine**. The Select Events dialog box opens.
3. Select the events that you want to combine and click **OK**. The Combination Event Properties dialog box opens.

Figure 12.6 Combination Event Properties Dialog Box

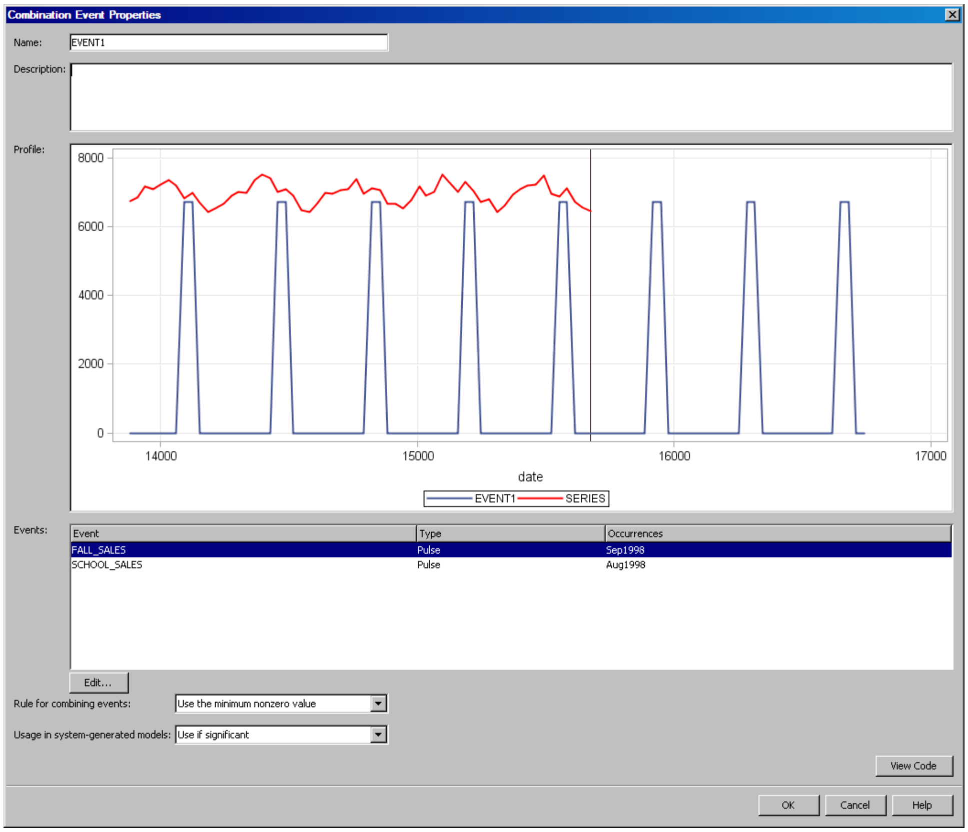

4. Specify a name and a description for the combined event.

5. Specify the events to include in the combination event. By default, the **Events** box lists the events in the combination event. To add or remove an event, click **Edit**.

6. From the **Rule for combining events**, specify what value to use in the combination event.

7. Click **OK**. The combined event appears in the Event Repository dialog box.

Example 3: Creating an Event for Third Quarter Sales

In your project, you have an event for back-to-school sales and an event for fall sales. You want to add the values in these two events to get a single event for third quarter sales.

To create an event for third quarter sales, complete the following steps:

1. Select **Project** → **Event Repository**. The Event Repository dialog box opens.

Figure 12.7 Events in the Events Repository

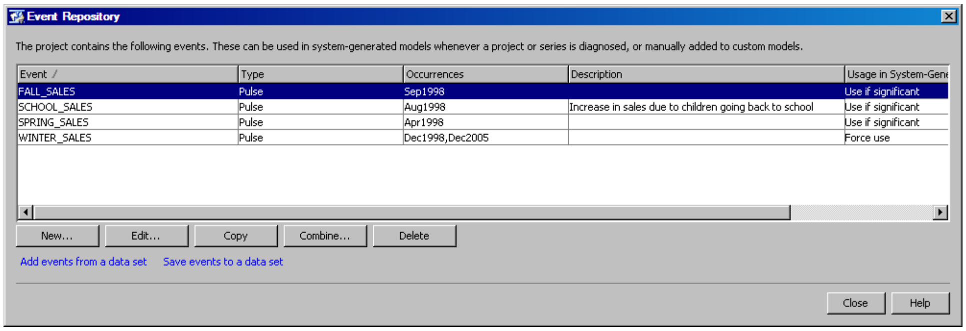

2. Click **Combine**. The Select Events dialog box opens.

3. Select the check boxes for the FALL_SALES and SCHOOL_SALES events and click **OK**. The Combination Event Properties dialog box opens.

Figure 12.8 Select the Events

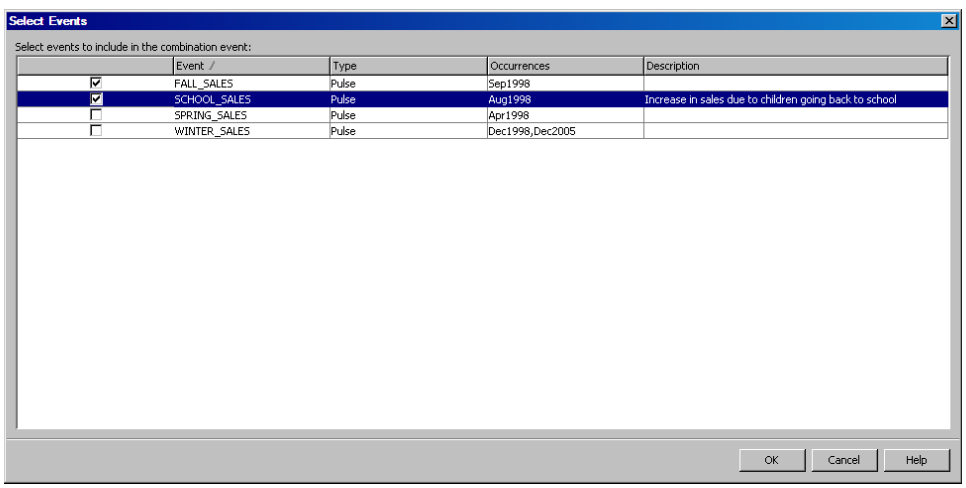

4. In the **Name** box, type `Quarter3` and in the **Description** box, type `A combined event for third quarter sales. This event combines back-to-school sales and fall sales.`

5. From the **Rule for combining events**, select **Add the values**.

Figure 12.9 Combination Event Properties Dialog Box for the Quarter3 Event

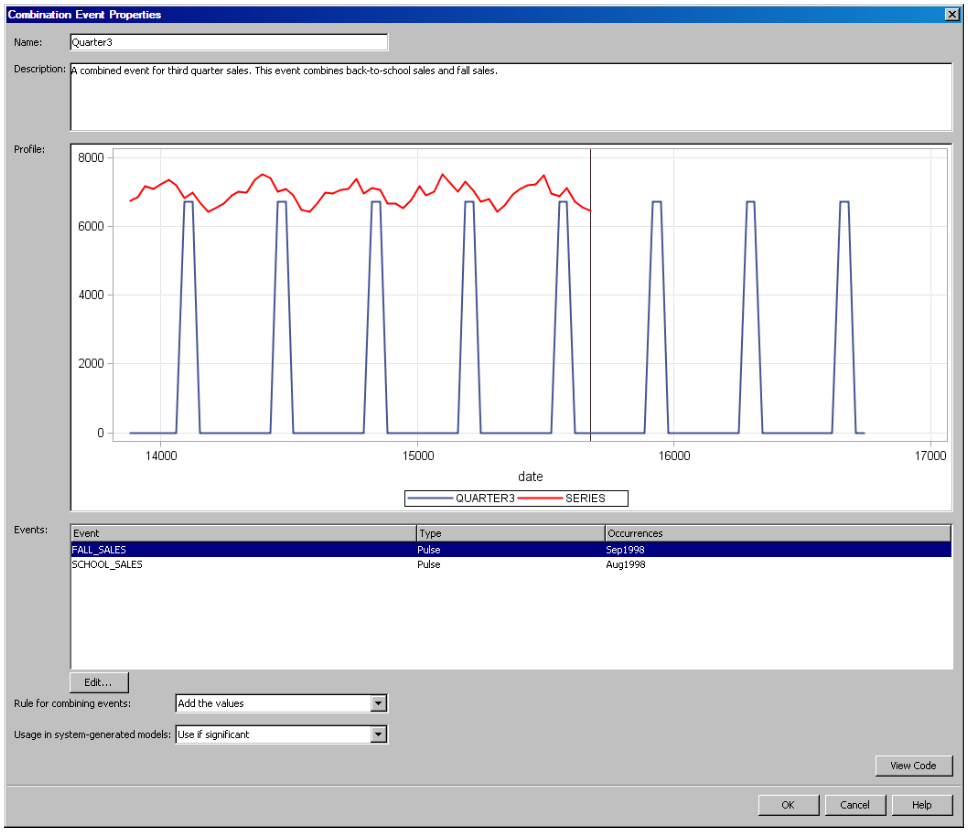

6. Click **OK**. The Quarter3 event now appears in the Event Repository dialog box.

You must reforecast the project to incorporate the new event in the project. SAS Forecast Studio displays the Event Repository Changed dialog box, so you can quickly reforecast the project.

Figure 12.10 Event Repository Changed Dialog Box

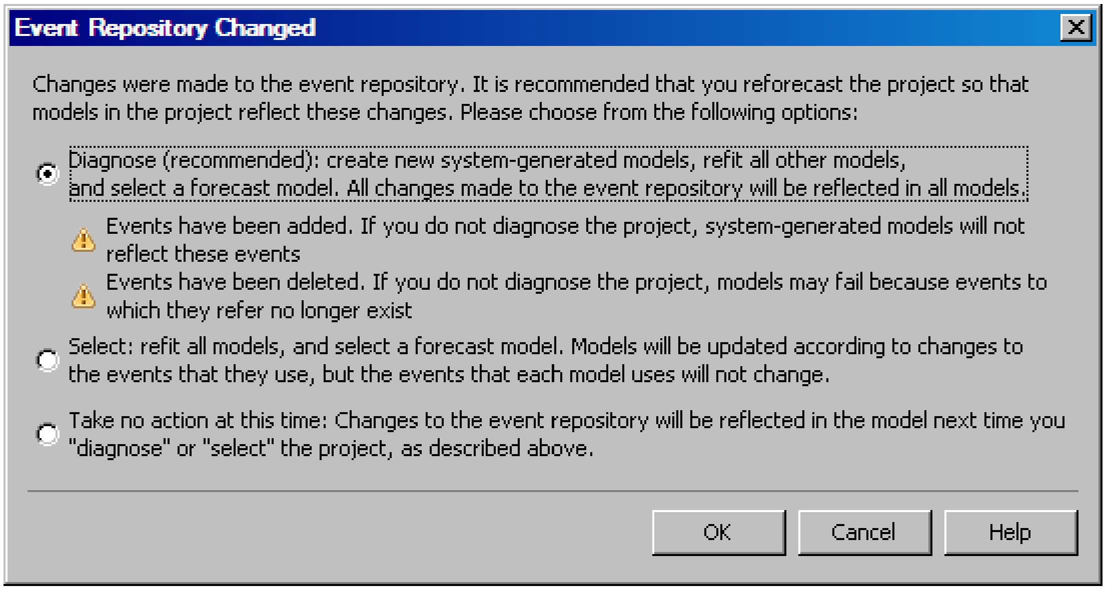

Edit an Event

To edit an event, complete the following steps:

1. Select **Project** → **Event Repository**. The Event Repository dialog box opens.
2. Select the event to modify and click **Edit**. The Event Properties dialog box opens. The options in this dialog box are the same as when you created the event.
3. Edit the event properties that you want to change and click **OK**.

Copy an Event

To copy an event, complete the following steps:

1. Select **Project** → **Event Repository**. The Event Repository dialog box opens.
2. Select an event and click **Copy**. The Event Properties dialog box opens. By default, the copied event is named *event-name*COPY*n* where *n* is the lowest integer value available.
3. Edit the event properties that you want to change and click **OK**. The new event appears in the Event Repository dialog box.

Sharing Events

Add Events from a SAS Data Set

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

To import an event, complete the following steps:

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

   If an event with the same name already exists in your project, you are prompted whether to overwrite the event with the same name.
4. Click **OK**. The imported events appear in the list of events for the project.

Save Events to a SAS Data Set

By saving events to a SAS data set, you can share events with other SAS Forecast Studio users at your site. This SAS data set 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 save an event to a data set, complete the following steps:

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

NOTE: When you export an event, two SAS data sets are created. The first data set (called *data-set-name*) contains the variables for your event. The second data set (called *data-set-name*2) contains the information about whether the event should be included in the model generation. Only the *data-set-name* data set is visible in SAS Forecast Studio when you import an event. However, both data sets are automatically imported when you select the *data-set-name* data set. For more information, see “Add Events from a SAS Data Set” on page 105.

Delete an Event

1. Select **Project** → **Event Repository**. The Event Repository dialog box opens.
2. Select the event to delete and click **Delete**.

Because you deleted an event from the event repository, SAS Forecast Studio must reforecast the project to reflect this change in the models. When you close the Event Repository dialog box, the Event Repository Changed dialog box prompts you to reforecast the project.

How Events Are Used in Model Generation

Understanding How SAS Forecast Studio Interprets Events

For each event, you can specify how to include the event when SAS Forecast Studio generates the models. You can choose from the following options:

Do not use
: does not include the event in the model.

Try to use
: includes the event in the model if feasible and if the resulting model is stable and an improvement over the simple model.

Use if significant
: includes the event in the model if feasible, if the resulting model is stable, an improvement over the simple model, and if the event is significant.

Force use
: includes the event in the model if feasible.

This status is used for all the models that SAS Forecast Studio generates. However, you can also choose to include an event only in an ARIMA, Subset ARIMA, Unobserved Components, or Multiple Regression model. You add this event as part of the model specification. For more information about creating models, see “Adding an Event to a Custom Model” on page 107.

Specify the Event Status

To specify the status of an event for all the models that SAS Forecast Studio generates in the project, complete the following steps:

1. Select **Project** → **Event Repository**. The Event Repository dialog box opens.
2. Select the event(s) whose status you want to change.
3. From the drop-down list in the **Usage in System-Generated Models** column, specify how the event should be used in model generation. For more information about these options, see "Understanding How SAS Forecast Studio Interprets Events" on page 106.

   NOTE: If you selected multiple events, then specifying the value of the **Usage in System-Generated Models** column for one event changes the value for all the selected events.

Adding an Event to a Custom Model

Add an Event to an ARIMA Model

To add an event only to an ARIMA model, complete the following steps:

1. Open the ARIMA Model dialog box.
2. In the selection pane, click **Events**. All of the events in the project are available from the Events panel. To create a new event, click **New**.
3. In the **Include in Model** column, select the check box of the events that you want to include in the model.
4. For those events that you included in the model, click **NONE** in the **Transfer Function** column.

Figure 12.11 Events in the ARIMA Model Dialog Box

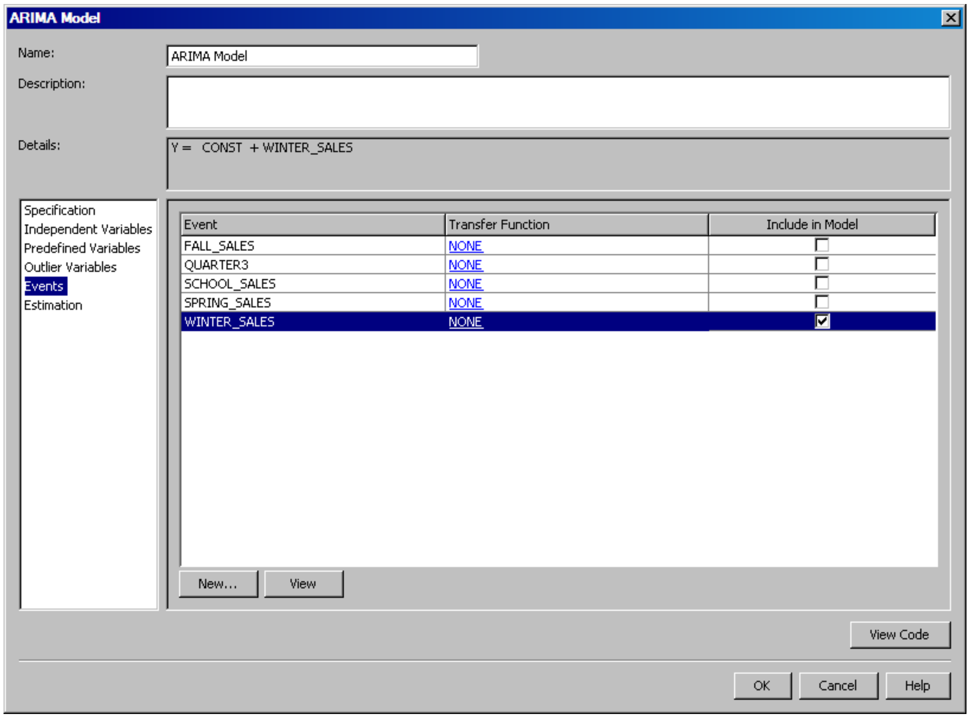

5. In the Transfer Function dialog box, specify the lagging periods, simple differencing, seasonal differencing, numerator factors, and denominator factors.

   NOTE: Predefined events that are applied to custom ARIMA models must be manually differenced to match the specification.

   Click **OK**.

You can now see the impact of the event from the **Details** box. The following example shows a model with the WINTER_SALES event with simple and seasonal differencing.

Figure 12.12 Event in the ARIMA Model Specification

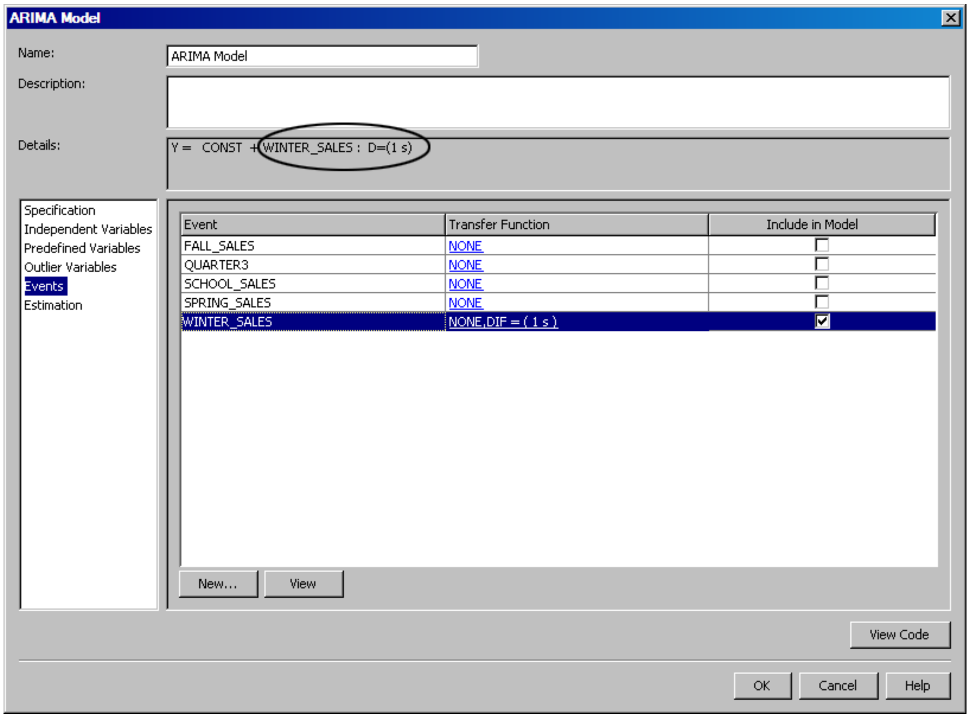

Add an Event to a Subset (Factored) ARIMA Model

To add an event only to a subset ARIMA model, complete the following steps:

1. Open the Subset ARIMA Model dialog box.
2. In the selection pane, click **Events**. All of the events in the project are available from the Events panel. To create a new event, click **New**.
3. In the **Include in Model** column, select the check box of the events that you want to include in the model.
4. For those events that you included in the model, click **NONE** in the **Transfer Function** column.
5. In the Transfer Function dialog box, specify the lagging periods, differencing order, numerator polynomial, and denominator polynomial.

   Click **OK**.

You can now see the impact of the event from the **Details** box. The following example shows a model with the WINTER_SALES event that includes a lag of one period.

Figure 12.13 Event in the Subset ARIMA Model Specification

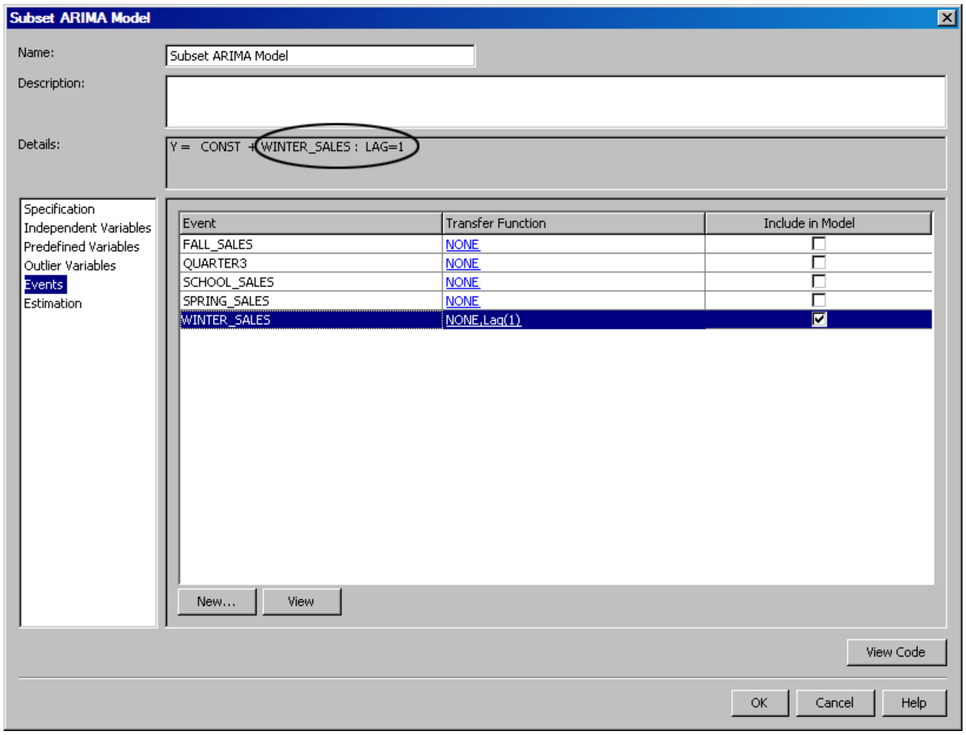

Add an Event to an Unobserved Components Model

To add an event only to an unobserved components model, complete the following steps:

1. Open the Unobserved Components Model dialog box.
2. In the selection pane, click **Events**. All of the events in the project are available from the Events panel. To create a new event, click **New**.
3. In the **Include in Model** column, select the check box of the events that you want to include in the model.
4. For those events that you included in the model, click **NONE** in the **Transfer Function** column.
5. In the Transfer Function dialog box, specify the lagging periods and differencing order.

   Click **OK**.

You can now see the impact of the event from the **Details** box. The following example shows a model with the WINTER_SALES event that includes a lag of two periods and a differencing order of d=(1,2).

Figure 12.14 Event in the Unobserved Components Model Specification

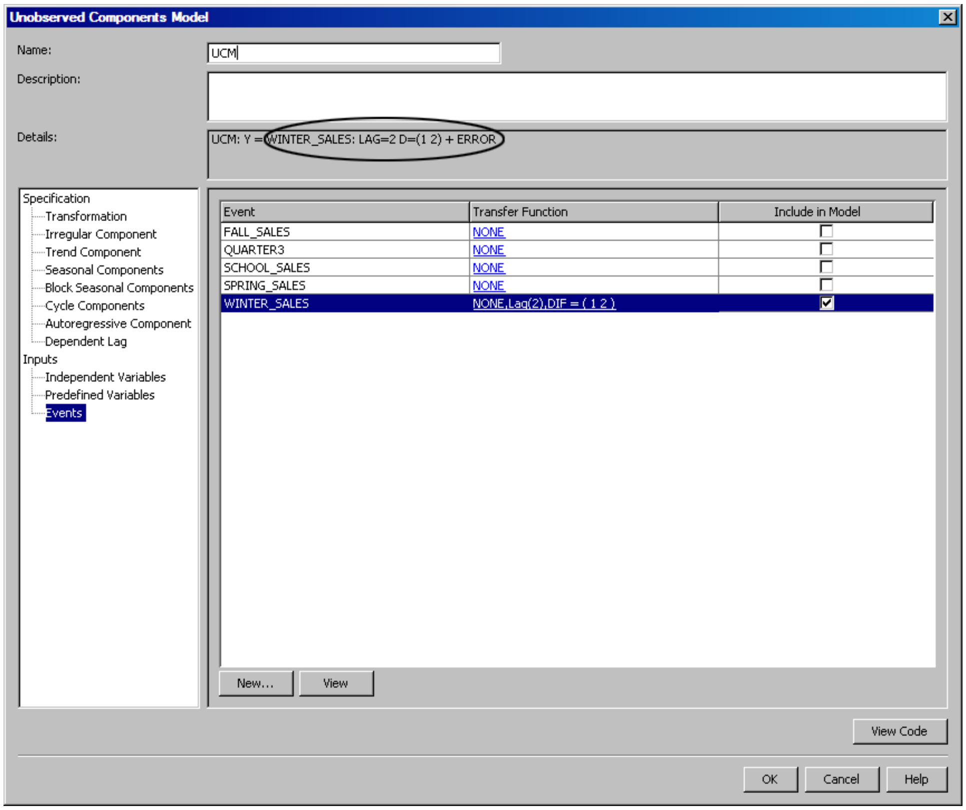

Add an Event to a Multiple Regression Model

To add an event only to a multiple regression model, complete the following steps:

1. Open the Multiple Regression dialog box.
2. In the selection pane, click **Events**. All of the events in the project are available from the Events panel. To create a new event, click **New**.
3. In the **Include in Model** column, select the check box of the events that you want to include in the model.

You can now see the impact of the event from the **Details** box. The following example shows a model with the FALL_SALES and WINTER_SALES events

Figure 12.15 Event in the Multiple Regression Specification

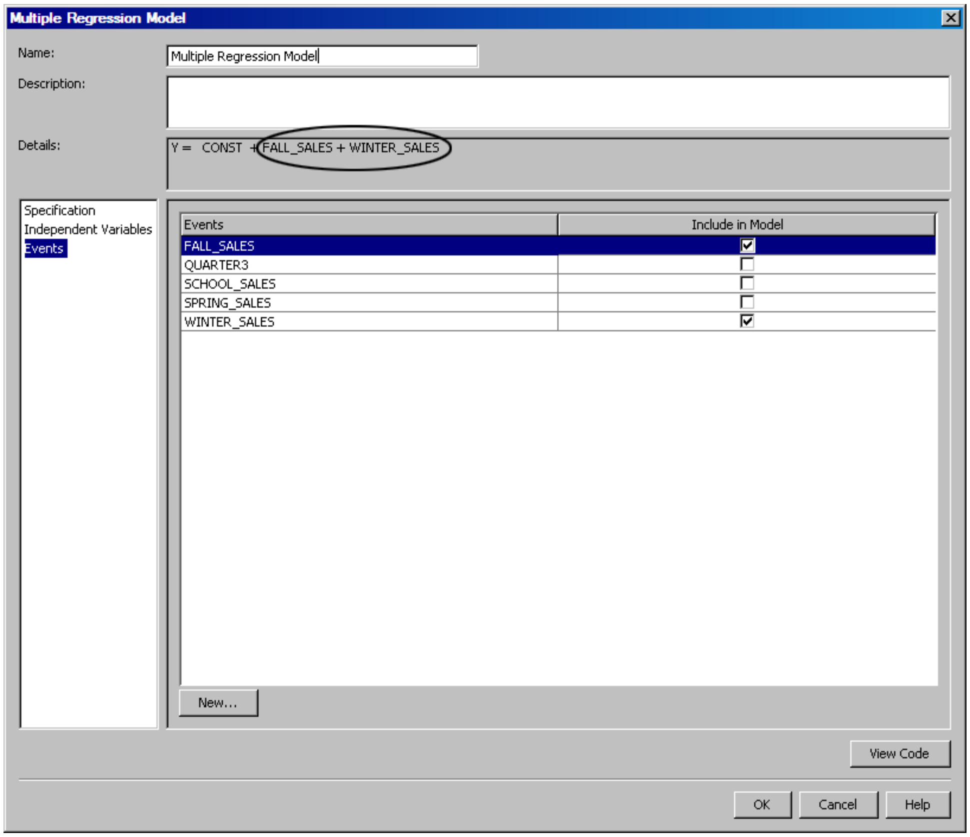

View the Effect of Events

You can view the effect of events from the forecast plot or the data table.

To view the effect of events, complete the following steps:

1. Select **Project** → **Edit Forecasting View Properties**. The Forecasting View Properties dialog box opens.
2. For the forecast plot and data table, select the **Effect of events** check box.

Chapter 13

Working with Overrides

Contents

Understanding Overrides

What Is an Override?

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

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

You can specify overrides only for future values. Series that have an associated override are identified by 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 121.

When Can You Specify an Override?

You can specify overrides only for future values.

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

- You can enter overrides only for the time ID values in the horizon. The horizon is determined by the value that you specify for the **Number of periods to forecast (horizon)** option in the Forecasting Settings dialog box and the end date for the series. The end date for the series is the largest time ID with a nonmissing value for the dependent variable.
- You cannot add overrides when the value for the **Number of periods to forecast (horizon)** option is less than the number of the **Calculate statistics of fit over an out-of-sample range** option. You cannot edit the time ID values in the horizon from the data table, and the Override Calculator is not available if this condition is met.

  If you specified any overrides before this condition was met, then these previous overrides are not removed. The override icons are still visible in the hierarchy. If you make any additional changes to the values of the forecasting options that remove this condition, then these overrides will become visible in the Forecasting View.
- You cannot add a negative override if negative forecasts are not allowed in the project. To allow negative forecasts, you must select the **Allow negative forecasts** check box in the Forecast panel of the Forecasting Settings dialog box.

For more information about these options, see "Specify the Number of Periods to Forecast" on page 73.

How Are Overrides Processed?

How SAS Forecast Studio processes overrides depends on whether any override conflicts are detected and how SAS Forecast Studio resolves conflicts between locked overrides. When setting up a hierarchy, you can choose whether the conflicts are resolved before reconciliation or the conflicts are ignored. For more information about the advanced reconciliation settings, see "Understanding Reconciliation" on page 125.

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

1. When you add an override, a "Reconciliation needed" message appears at the top of the Forecasting View.

2. When you click **Reconcile**, SAS Forecast Studio checks for conflicts in the locked overrides.

   - If no conflicts are detected, SAS Forecast Studio performs an override reconciliation. During the reconciliation, one of the following events occurs:
     - The hierarchy is fully reconciled, and the reconciliation message in the Forecasting View disappears.
     - SAS Forecast Studio cannot reconcile the hierarchy. A "Reconciliation failed" message appears at the top of the Forecasting View. Click **View** to view the unreconciled nodes. For more information, see "Working with Reconciliation Failures" on page 129.

       NOTE: Unreconciled nodes occur only if you restrict the direction of the reconciliation. For more information, see "Specify the Advanced Reconciliation Options" on page 128.
   - If conflicts are detected, the Override Conflicts dialog box opens. You view and resolve any conflicts.

     After all the conflicts are resolved, SAS Forecast Studio reconciles the hierarchy. During the reconciliation, one of the following events occurs:
     - The hierarchy is fully reconciled, and the reconciliation message in the Forecasting View disappears.
     - SAS Forecast Studio cannot reconcile the hierarchy. A "Reconciliation failed" message appears at the top of the Forecasting View. Click **View** to view the unreconciled nodes. For more information, see "Working with Reconciliation Failures" on page 129.

       NOTE: Unreconciled nodes occur only if you restrict the direction of the reconciliation. For more information, see "Specify the Advanced Reconciliation Options" on page 128.

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

1. When you add an override, a "Reconciliation needed" message appears at the top of the Forecasting View.
2. When you click **Reconcile**, 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 "Reconciliation failed" message appears at the top of the Forecasting View. Click **View** to view the unreconciled nodes. For more information, see “Working with Reconciliation Failures” on page 129.

     The Reconciliation Failure Report lists the unreconciled nodes and any conflicts if they were detected but ignored. If there is no restriction on the direction of the reconciliation, then resolving any override conflicts will eliminate the unreconciled nodes. However, if you restricted the direction of the reconciliation, then resolving all the override conflicts might reduce, but not necessarily eliminate, the number of unreconciled nodes. For more information about these reconciliation options, see “Specify the Advanced Reconciliation Options” on page 128.

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 117.
- By using the Override Calculator. This method is good if you want to create overrides for several time periods. For more information, see “Understanding the Override Calculator” on page 117.

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 Number of Periods to Forecast” on page 73.

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 121.

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 121.

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 114.

Understanding the Override Calculator

Add an Override by Using the Override Calculator

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

1. Select **Series** → **Override Calculator**. The Override Calculator opens.

   If the Override Calculator is not available, see "When Can You Specify an Override?" on page 114.

2. Select the time periods that you want to override.

   NOTE: Any time periods that you have selected in the data table in the Forecasting View are automatically selected.

3. Specify how to calculate the overrides for the selected time periods. You can base the override values on an existing value, such as the reconciled forecast, the statistical forecast, current overrides, or the final forecast. You can also choose to set the override values for the selected time periods to a specific value. This value could be assigned proportionally, evenly, or equally (all selected time periods are assigned the same value).

4. Click **OK** to save your changes and close the Override Calculator.

Example 1: Basing the Override on an Existing Value

In this example, you want to create overrides for the Feb2003 and March2003 time periods that are 15% less than the value of the final forecast.

To create these override values, complete the following steps:

1. Select the Region2 > Line2 series in the hierarchy view.
2. In the Forecasting View, select the Feb2003 and March2003 time periods in the data table.
3. Select **Series** → **Override Calculator**. The Override Calculator opens.
4. For how to calculate the override values, select **Base on existing values**. Then specify the following options:
   a) From the drop-down list, select **Final Forecast**.
   b) In the first box, select **-**.
   c) In the text box, type `15`.
   d) From the drop-down list, select **%**.

   The Override Calculator appears similar to the following:

   Figure 13.1 Creating an Override Based on the Final Forecast

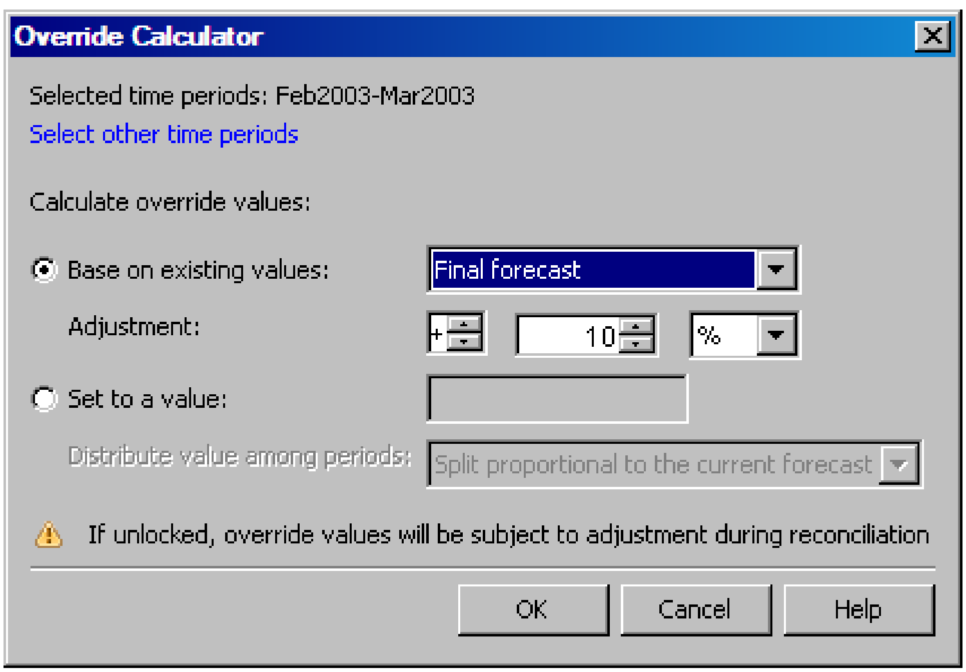

5. Click **OK**.

The new override values appear in the Override row of the data table.

Figure 13.2 Override Values for Feb2003 and Mar2003 Time Periods

☑ Active series

	Nov2002	Dec2002	Jan2003	Feb2003	Mar2003	Apr2003	May2003	Jun2003	Jul2003	Aug2003
Historical Data	1695.000000	1556.000000	.	.	.	.	.	.	.	.
Forecast Model	1758.823104	1614.022406	1705.956111	1746.442398	1675.032112	1780.458664	1835.094515	1843.410438	1722.348345	1796.526790
Reconciled Forecast	1698.890060	1566.899131	1689.332776	1808.974670	1741.591621	1802.317549	1850.576464	1930.003697	1754.415494	1776.628120
Override				1989.872137	1915.750783					
Lock Override				☑	☑					
Final Forecast	.	.	1583.209764	1808.974670	1741.591621	1802.317549	1850.576464	1930.003697	1754.415494	1776.628120

Example 2: Splitting a Value Proportionally Between Time Periods

In this example, you want to create overrides for the Feb2003 and Mar2003 time periods. However, you want these overrides set to a specific value. The value for each time period is determined by the proportion of the current forecasts for the two time periods.

To create these override values, complete the following steps:

1. Select the Region2 > Line2 series in the hierarchy view.
2. In the Forecasting View, select the Feb2003 and Mar2003 time periods in the data table.
3. Select **Series** → **Override Calculator**. The Override Calculator opens.
4. For how to calculate the override values, select **Set to a value**. Then specify the following options:
   a) In the text box, type `2500`.
   b) From the **Distribute value among periods** drop-down list, select **Split proportional to the current forecast**.

The Override Calculator appears similar to the following:

Figure 13.3 Setting an Override to a Specific Value

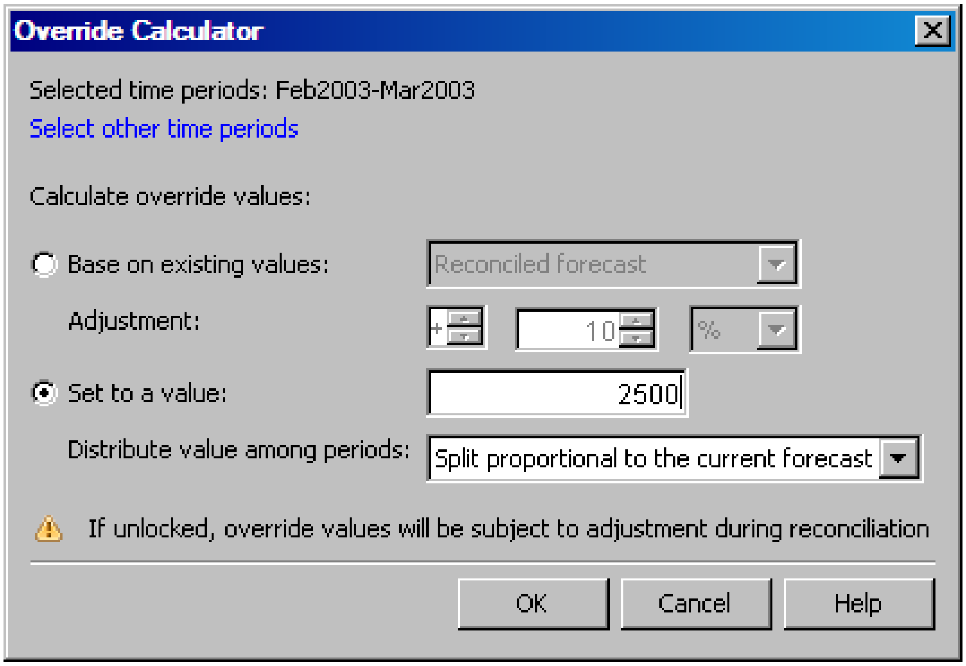

5. Click **OK**.

The new override values appear in the Override row of the data table.

Figure 13.4 Override Values for Feb2003 and Mar2003 Time Periods

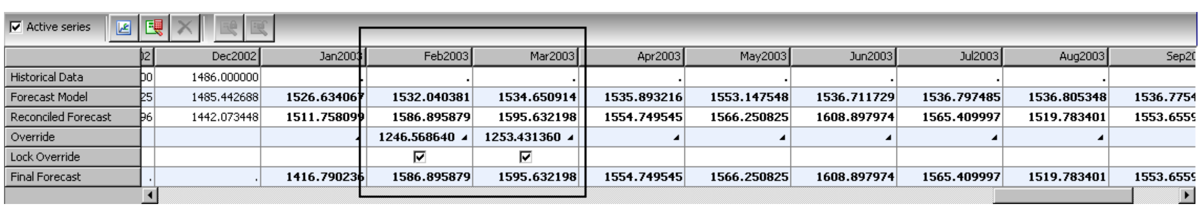

To determine the proportion for each override value, SAS Forecast Studio uses the following calculations:

Time Period	Proportion of Total Forecast	Override Value
Feb2003	1586.895879/3182.527859 = .498627	.498627 * 2500 = 1246.568540
Mar2003	1595.632198/3182.527859 = .501373	.501373 * 2500 = 1253.431360

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 Lock Override 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 several overrides, select the overrides in the data table and select **Series** → **Lock Selected Overrides**. To unlock several overrides, select the overrides in the data table and **Series** → **Unlock Selected Overrides**.

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

Add a Note to a Series

When you specify an override for a forecasted value, you might want to document why you specified this override. In SAS Forecast Studio, you can add notes to series to document any changes that you might have made. The note is automatically saved with the series when you select a different series in the hierarchy or the table view. Series that have an associated note are identified in the hierarchy with a Note icon. The notes persist when you close the project.

To enter a note, complete the following steps:

1. In the hierarchy or table view, select a series.
2. Select **View** → **Notes**. The Notes box appears at the bottom of the **Hierarchy View** tab or the **Table View** tab.
3. In the text box, type your note. This note cannot exceed 2,000 characters.

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 of the override that you want to delete and click .

- To remove all the overrides in the current forecast horizon, select **Project** → **Remove All Overrides**. In the warning that appears, click **Remove** to remove the overrides.
- To delete the overrides in a series, select the overrides that you want to delete and select **Series** → **Delete Selected 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 13.5 Example of an Override Conflict

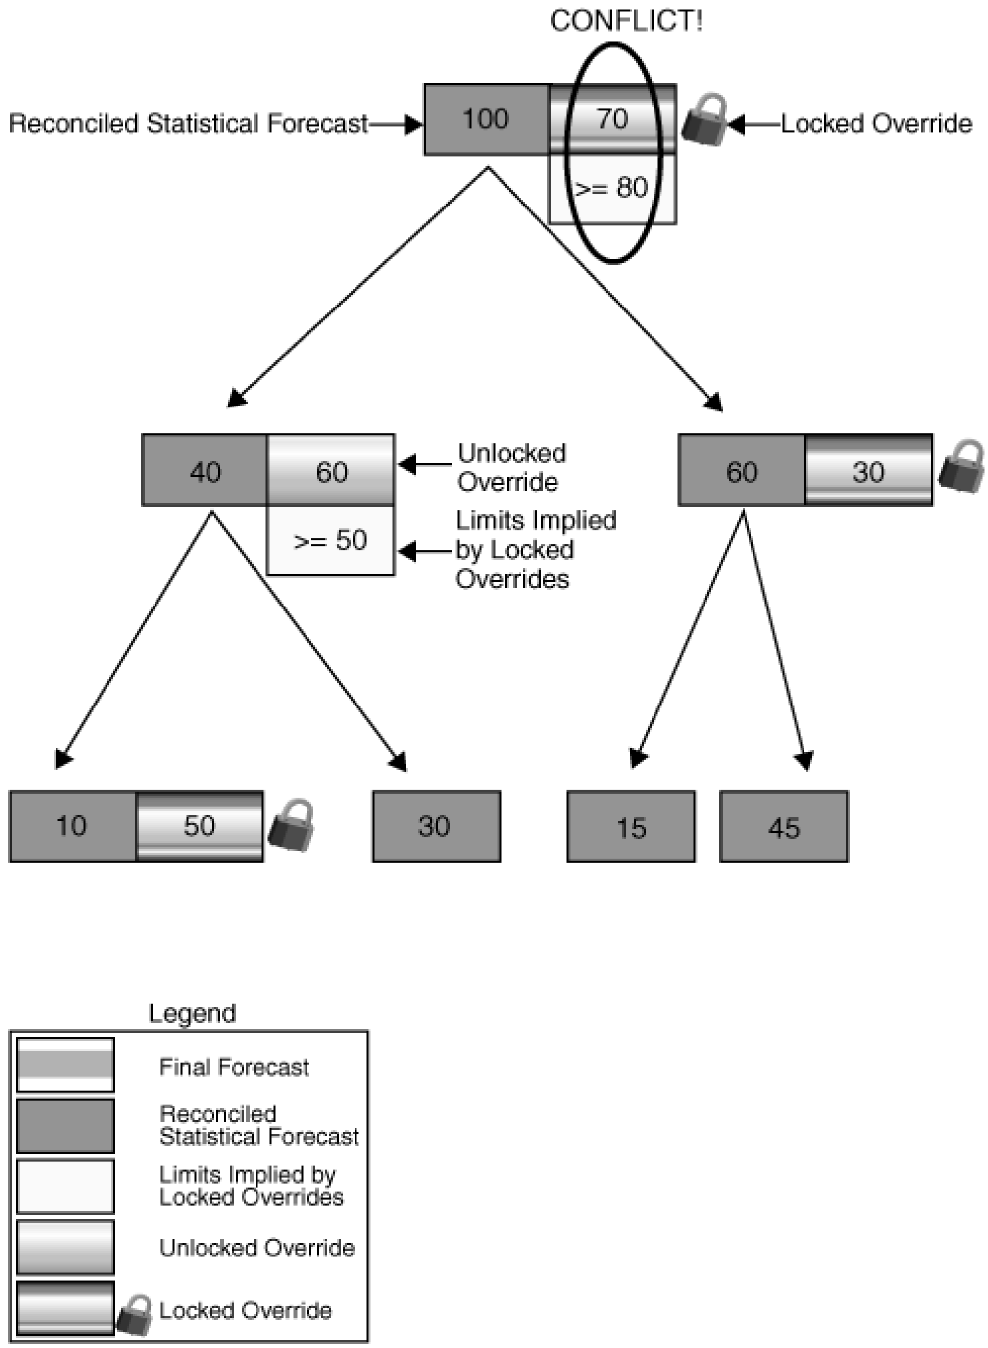

Resolve Override Conflicts

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

To resolve conflicts, complete the following steps:

1. From the **Date** pane, select a time period. The right pane of the dialog box shows the conflict for the time period that you selected.

2. To resolve a conflict, you can do the following in the Parent and Children tables:

   - Clear the check boxes for the overrides that can be unlocked.
   - Change the override value in the table.

   NOTE: For more information about when you can specify an override, see "When Can You Specify an Override?" on page 114.

3. (Optional) Repeat steps 1 and 2 for all of the time periods in the **Date** pane.

4. Click **Reconcile**. SAS Forecast Studio attempts to reconcile the hierarchy. After this reconciliation has completed, a message appears. Click **Yes** to close this message.

Chapter 14

Understanding Hierarchy Reconciliation

Contents

Understanding Reconciliation

If you add an override to data that is forecasted hierarchically, then you must reconcile the hierarchy before SAS Forecast Studio can calculate the final forecast. SAS Forecast Studio might not be able to reconcile the hierarchy for a variety of reasons. These failures appear in the Reconciliation Failure dialog box.

How SAS Forecast Studio reconciles the hierarchy depends on the values that you set for the advanced reconciliation options. Using these options, you can specify the following:

- 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 the historical mean.
- how you want to resolve conflicts between locked overrides
- whether to restrict the direction of reconciliation

Depending on the combination of options that you select, your results could be slightly different. For more information about how to set these options, see "Specify the Advanced Reconciliation Options" on page 128.

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

How to Resolve conflicts	**Reconciliation direction**	**Result**
Require that conflicts be resolved before reconciliation	**No restriction**	When reconciling the hierarchy, SAS Forecast Studio determines whether there are any conflicts. If conflicts exist, they appear in the Override Conflicts dialog box. In this dialog box, you perform either of the following tasks: • You can close this dialog box, and the hierarchy remains unreconciled. • You can resolve any conflicts, and SAS Forecast Studio reconciles the hierarchy.
Ignore conflicts	**No restriction**	You might select these options if you want to reconcile the hierarchy as best as possible, but still allow a few conflicts. When reconciling the hierarchy, SAS Forecast Studio determines whether there are any conflicts. However, because any conflicts should be ignored, the Override Conflicts dialog box does not open if conflicts are detected. Instead the Reconciliation Failure dialog box 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.)

How to Resolve conflicts	Reconciliation direction	Result
Require that conflicts be resolved before reconciliation	**Restrict direction**	Before reconciling the hierarchy, SAS Forecast Studio determines whether there are any conflicts. If conflicts exist, they appear in the Override Conflicts dialog box. In this dialog box, you perform either of the following tasks: • You can close this dialog box, and the hierarchy remains unreconciled. • You can resolve any conflicts, and SAS Forecast Studio reconciles the hierarchy. 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 dialog box opens and shows the reconciliation failures.
Ignore conflicts	**Restrict direction**	You might want to select these options if you want SAS Forecast Studio to reconcile the hierarchy with restrictions but without intervention. Before reconciling the hierarchy, SAS Forecast Studio determines whether there are any conflicts. However, because any conflicts should be ignored, the Override Conflicts dialog box does not open if conflicts are detected. Instead the Reconciliation Failure dialog box 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.

Reconcile the Hierarchy

You can reconcile the hierarchy in the following ways:

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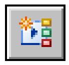 in the toolbar.
- If the "Reconciliation needed" message appears in the Forecasting View, then you can click **Reconcile**.

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. Select **Project** → **Hierarchy and Variable Settings**. The Hierarchy and Variable Settings dialog box opens.
2. Click **Advanced**. The Advanced Reconciliation Settings dialog box opens.
3. Select the disaggregation method.
4. Specify whether to ignore zero forecasts during reconciliation.
5. Select the **Use confidence interval widths in reconciliation magnitude calculations** check box to specify that the loss function for top-down reconciliation be weighted by the inverse of the variance of the input forecasts. For more information about this option, see the HPFRECONCILE procedure in the *SAS High-Performance Forecasting User's Guide*.
6. Specify how to resolve conflicts between locked overrides.
7. Specify whether to restrict the direction of the reconciliation.
8. Click **OK**.

Working with Reconciliation Failures

About the Reconciliation Failure Report

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

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

An unreconciled node could occur for a variety of reasons. For example, unreconciled nodes might result from override conflicts or if you restricted the direction of the reconciliation direction. For more information about how to specify the reconciliation direction, see "Specify the Advanced Reconciliation Options" on page 128.

How to View Unreconciled Nodes

To view and resolve any reconciliation failures, complete the following steps:

1. If you have unreconciled nodes in your hierarchy, then you can view the reconciliation failures in the following ways:
   - Select **Project** → **Reconciliation Failure**.
   - Click **View** next to the "Reconciliation failed" message at the top of the workspace.
2. To resolve reconciliation failures that are the result of override conflicts, select **Resolve conflicts**. The Overrides Conflict dialog box opens.
3. Resolve any conflicts in the Overrides Conflict dialog box and click **Reconcile**. For more information, see "Resolve Override Conflicts" on page 124.

Example: Resolving Unreconciled Nodes

In the following example, the reconciliation failure report lists six unreconciled nodes. This example shows how you can determine which unreconciled nodes are due to override conflicts. It also explains other reasons why you might get an unreconciled node and how you might choose to resolve these failures.

After SAS Forecast Studio generates the forecasts for the project, you notice the "Reconciliation failed" message at the top of the workspace.

Figure 14.1 Reconciliation Failed Message in the Workspace

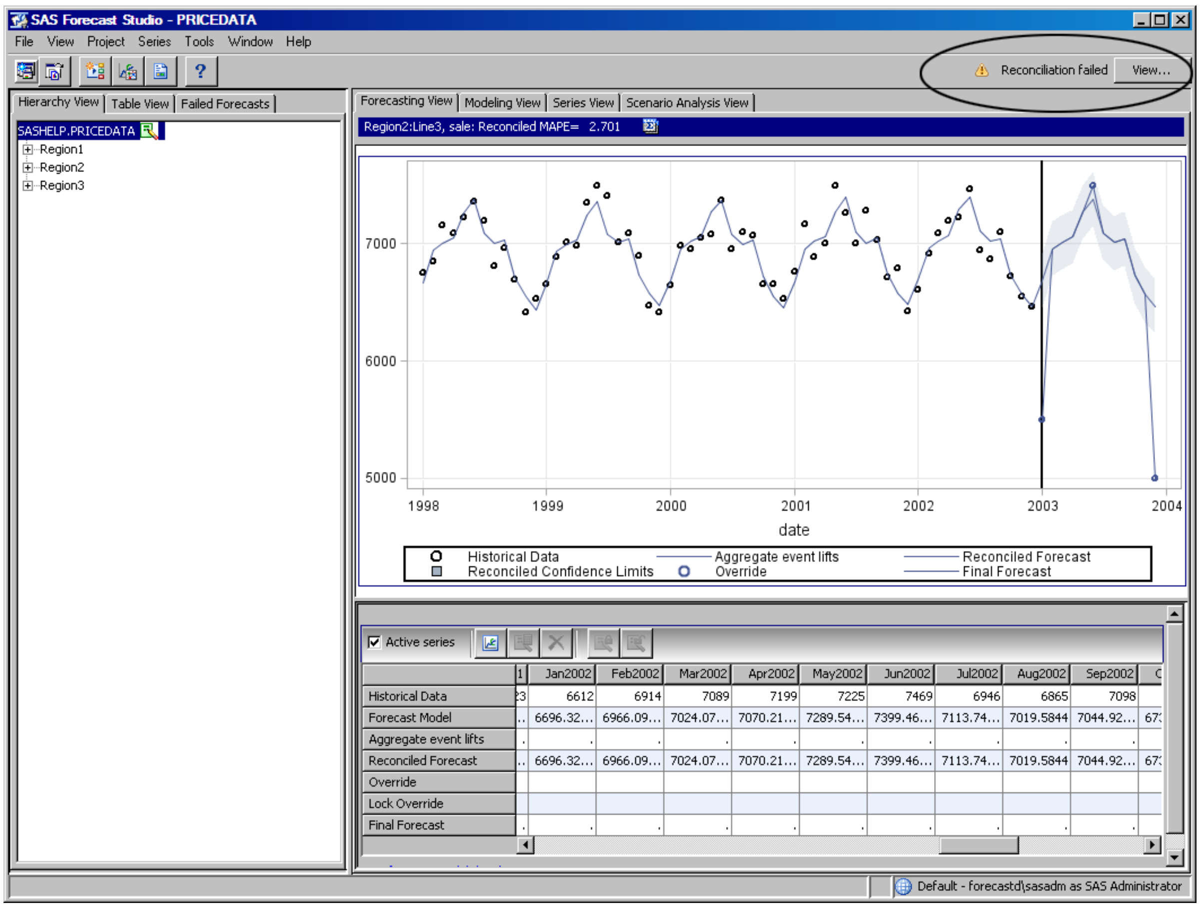

To view and resolve the reconciliation errors, complete the following steps:

1. Click **View** to open the Reconciliation Failures dialog box. In this dialog box, you see six unreconciled nodes.

Figure 14.2 Reconciliation Failures with Override Conflicts

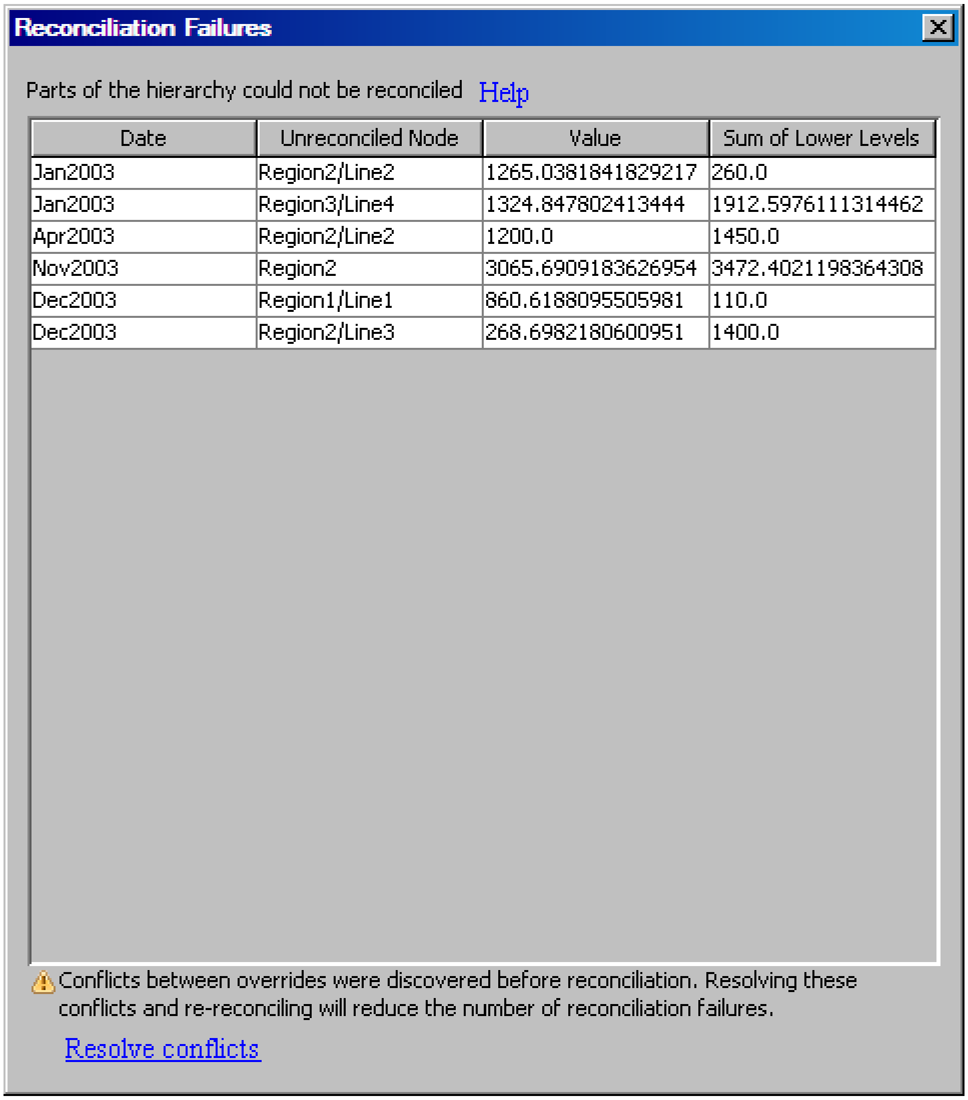

Date	Unreconciled Node	Value	Sum of Lower Levels
Jan2003	Region2/Line2	1265.0381841829217	260.0
Jan2003	Region3/Line4	1324.847802413444	1912.5976111314462
Apr2003	Region2/Line2	1200.0	1450.0
Nov2003	Region2	3065.6909183626954	3472.4021198364308
Dec2003	Region1/Line1	860.6188095505981	110.0
Dec2003	Region2/Line3	268.6982180600951	1400.0

2. Click **Resolve Conflicts** to determine how many of these unreconciled nodes are due to override conflicts. The Override Conflicts dialog box opens. Out of the six reconciliation failures, only the Apr2003 time period contains an override conflict.

3. To resolve the conflict for the Apr2003 time period, clear the **Locked** check box for the parent node. Now, the override value for the Region2/Line2 node is a guideline for the final forecast value. The value of the final forecast reflects the value of the override, but the final forecast and the unlocked override do not have to be the same value.

Figure 14.3 Reconciling the Apr2003 Node by Using the Override Conflicts Dialog Box

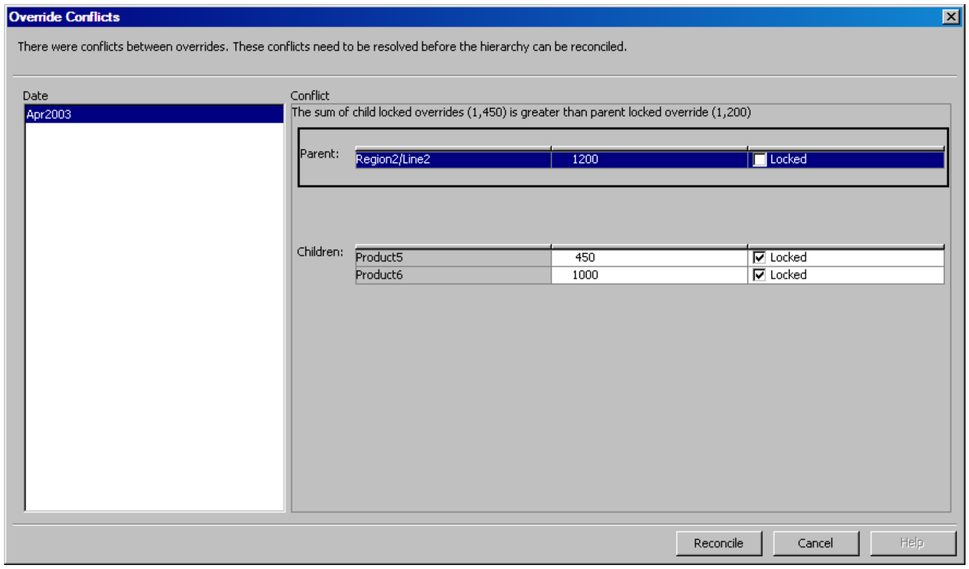

4. Click **Reconcile**. When SAS Forecast Studio reconciles the hierarchy, the following message appears.

Figure 14.4 Message About Hierarchy Reconciliation

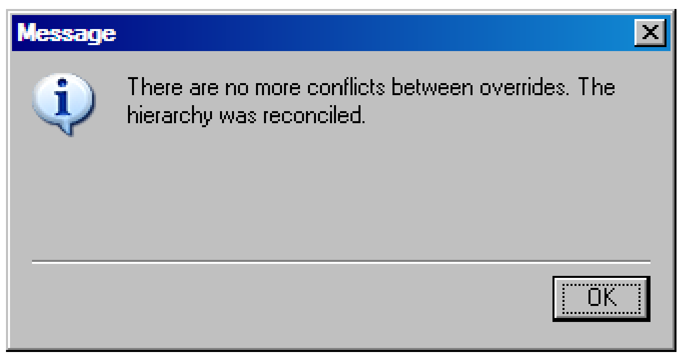

The remaining reconciliation failures are not due to override conflicts. Instead, these failures are the result of the advanced reconciliation settings.

5. To view the advanced reconciliation settings, select Select **Project** → **Hierarchy and Variable Settings**. The Hierarchy and Variable Settings dialog box opens.

6. Click **Advanced**. The Advanced Reconciliation Settings dialog box opens.

Figure 14.5 Advanced Reconciliation Settings Dialog Box

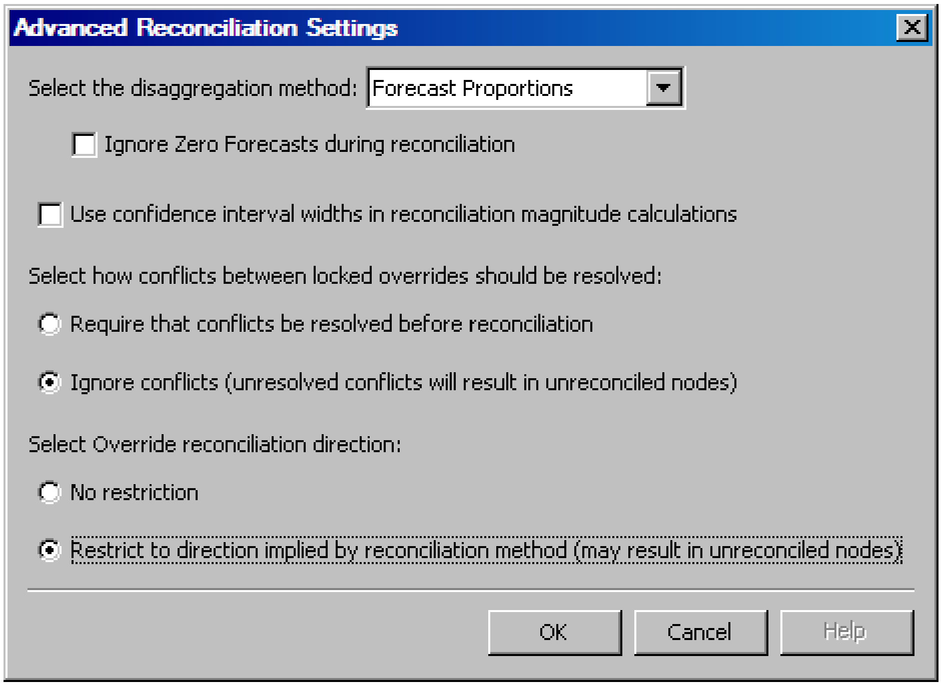

The five remaining reconciliation errors are most likely due to the reconciliation direction. Unreconciled nodes might occur when you restrict the direction of the reconciliation.

7. To unrestrict the direction of the reconciliation, select the **No restriction** option and click **OK**.

8. Click **OK** in the Hierarchy and Variable Settings dialog box. SAS Forecast Studio reconciles the hierarchy, and no additional reconciliation failures are found.

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.
- **Sum of Values** is selected as the aggregation method.
- The **Allow negative forecasts** option is not checked in the **Forecasts** panel of the Forecasting Settings dialog box. For more information, see "Specify the Number of Periods to Forecast" on page 73.
- The **Use confidence interval widths in reconciliation magnitude calculations** option is not checked in the Advanced Reconciliation Settings dialog box. For more information, see "Specify the Advanced Reconciliation Options" on page 128.

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

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 133.

Explanation of the Reconciliation Process

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

Figure 14.6 Example 1: SAS Forecast Studio Honors All Locked Overrides

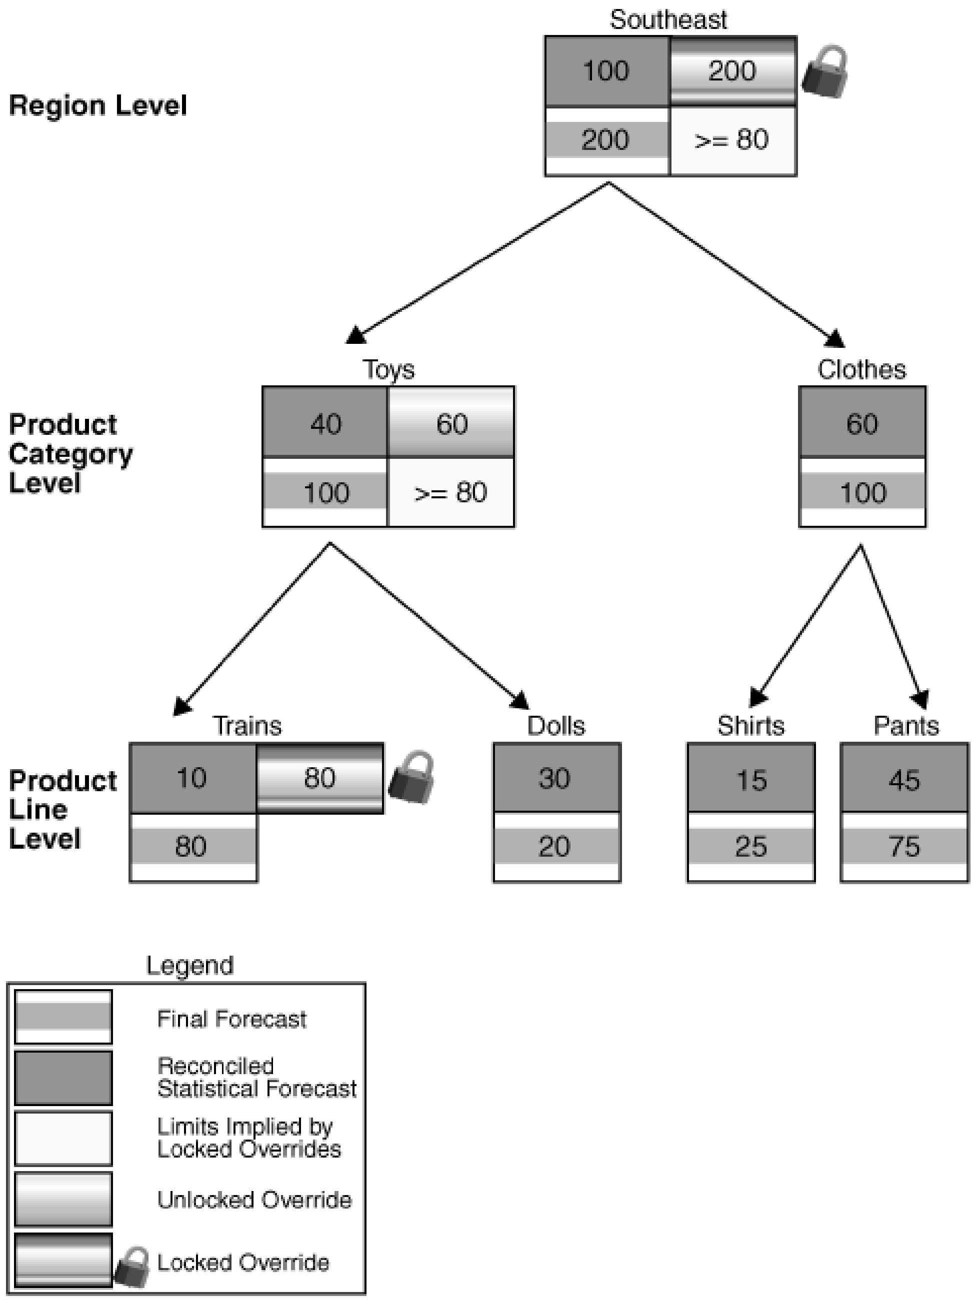

The following table explains the reconciliation process for each level of the hierarchy:

Level in the hierarchy	**Reconciliation process for the nodes in this level**
Region	Because you selected **Top Down** as the reconciliation method, SAS Forecast Studio starts reconciling the hierarchy at the highest level in the hierarchy. In this example, the highest level is the Region level. The reconciled statistical forecast for the Southeast node was 100; however, you specified 200 as a locked override for this level. SAS Forecast Studio honors this locked override, and consequently, the final forecast for the Southeast node is 200.
Product Category	In the Product Category level, there are two 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 two nodes, so the final forecast for both the Toys and Clothes nodes is 100.
Product Line	In the Product Line level, there are four nodes. The two nodes below the Toys node are the Trains and Dolls nodes. The two 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

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 133.

Explanation of the Reconciliation Process

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

Figure 14.7 Example 2: Implicit Bottom-Up Method Used to Reconcile Hierarchy

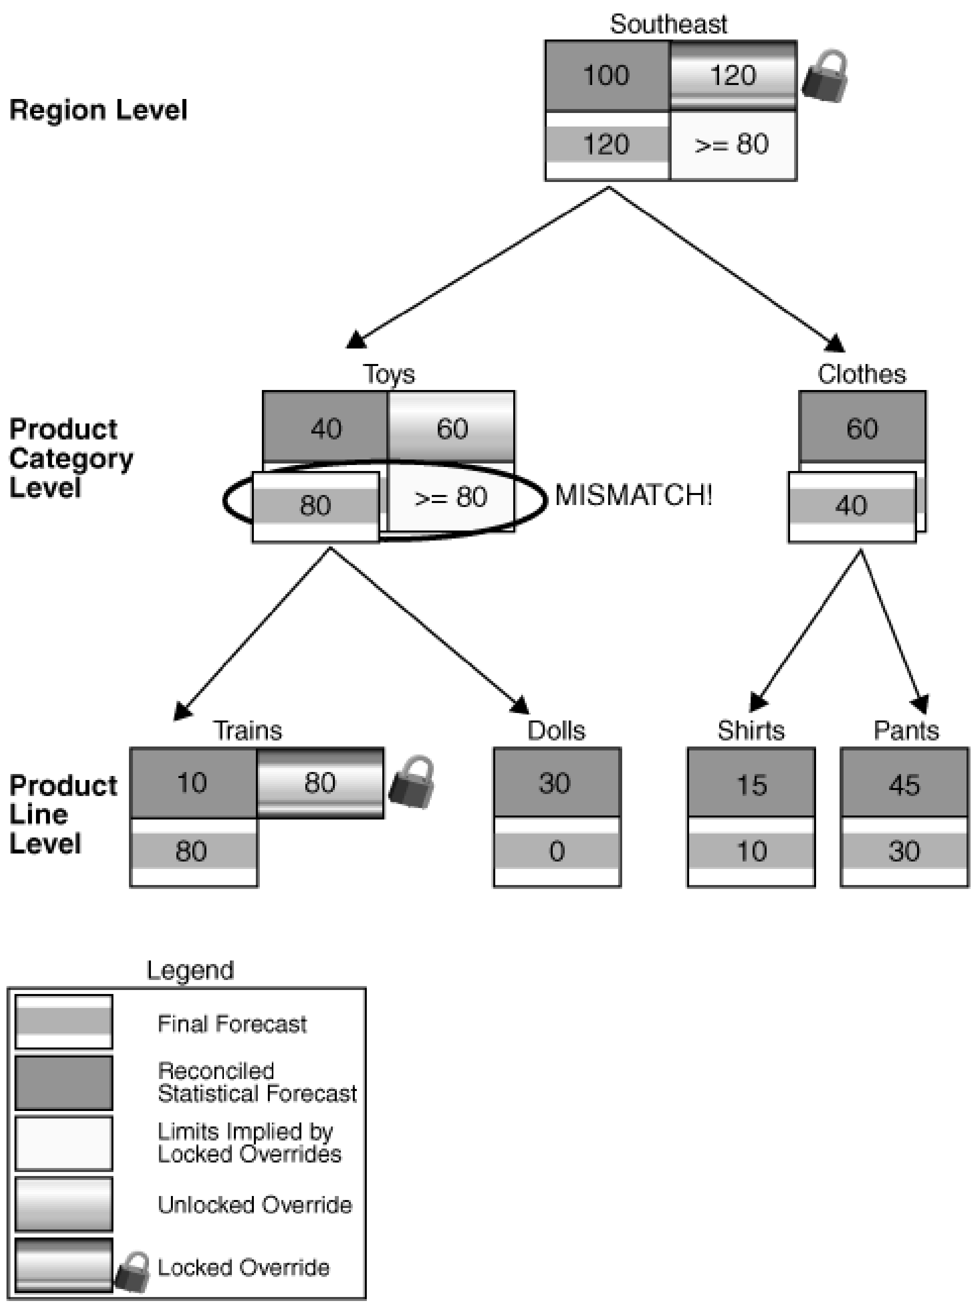

The following table explains the reconciliation process for each level of the hierarchy:

Level in the hierarchy	**Reconciliation process for the nodes in this level**
Region	Because you selected **Top Down** as the reconciliation method, SAS Forecast Studio starts reconciling the hierarchy at the highest level in the hierarchy. In this example, the highest level is the Region level. The reconciled statistical forecast for the Southeast node is 100; however, you specified 120 as a locked override for this level. SAS Forecast Studio honors this override, and consequently, the final forecast for the Southeast node is 120.
Product Category	In the Product Category level, there are two 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 two 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 four nodes. The two nodes below the Toys node are the Trains and Dolls nodes. The two 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

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 133.

Explanation of the Reconciliation Process

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

Figure 14.8 Example 3: Unreconciled Node in Hierarchy

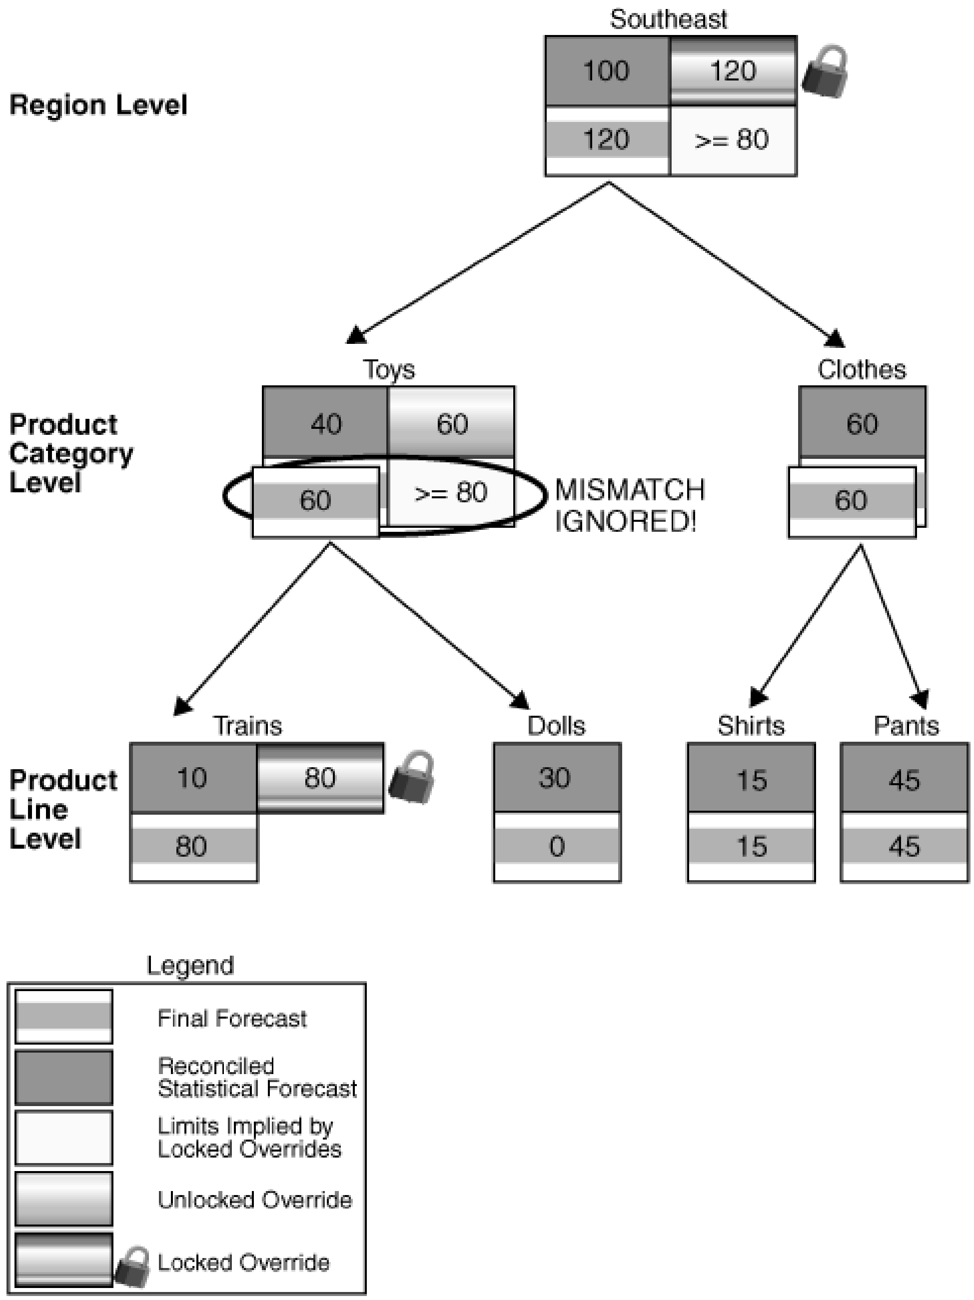

The following table explains the reconciliation process for each level of the hierarchy:

Level in the hierarchy	**Reconciliation process for the nodes in this level**
Region	Because you selected **Top Down** as the reconciliation method, SAS Forecast Studio starts reconciling the hierarchy at the highest level in the hierarchy. In this example, the highest level is the Region level. The reconciled statistical forecast for the Southeast node is 100; however, you specified 120 as a locked override for this level. SAS Forecast Studio honors this override, and consequently, the final forecast for the Southeast node is 120.
Product Category	In the Product Category level, there are two 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 two 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 dialog box.

Level in the hierarchy	**Reconciliation process for the nodes in this level**
Product Line	In the Product Line level, there are four nodes. The two nodes below the Toys node are the Trains and Dolls nodes. The two 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

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 133.

Explanation of the Reconciliation Process

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

Figure 14.9 Example 4: Bottom-Up Method of Reconciliation with No Conflicts

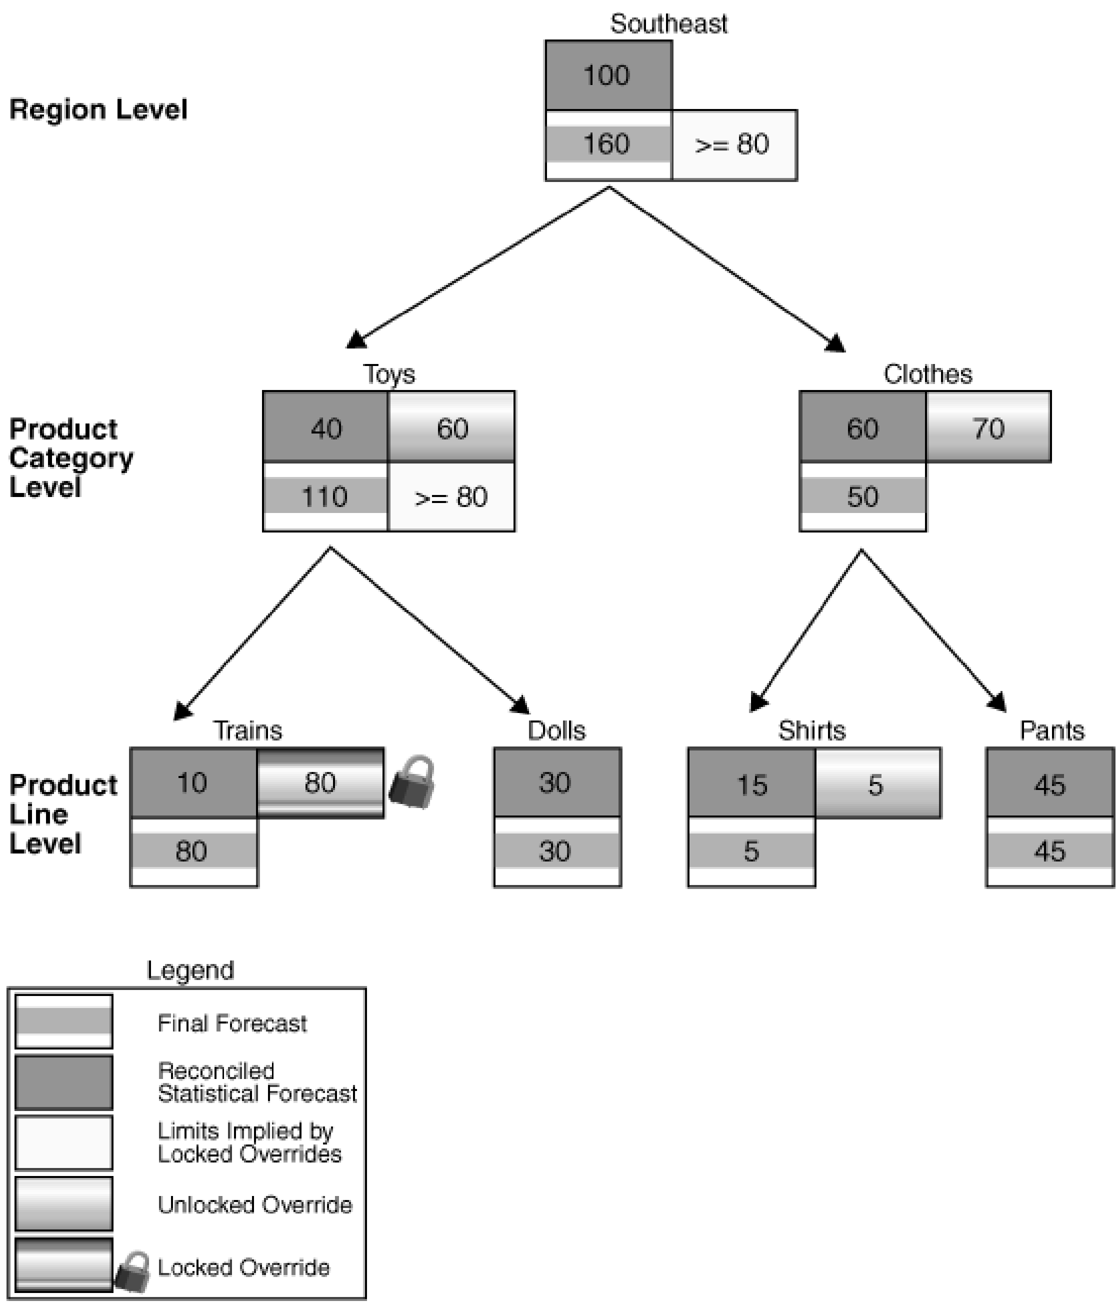

The following table explains the reconciliation process for each level of the hierarchy:

Level in the hierarchy	**Reconciliation process for the nodes in this level**
Product Line	Because you selected **Bottom Up** as the reconciliation method, SAS Forecast Studio starts reconciling the hierarchy at the lowest level in the hierarchy. In this example, the lowest level is the Product Line level. In the Product Line level, there are four nodes. The two nodes below the Toys node are the Trains and Dolls nodes. The two 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 two nodes—the Toys node and the Clothes node. SAS Forecast Studio uses the final forecasts of the nodes at the lower level to determine the final forecast for the nodes at the Product Category level. • For the Toys node, the final forecast for the Trains node (80) is added to the final forecast for the Dolls node (30), and so the final forecast for the Toys node is 110. Note that the unlocked override of 60 for the Toys node has no effect. • For the Clothes node, the final forecast for the Shirts node (5) is added to the final forecast for the Pants node (45), so the final forecast for the Clothes node is 50. Again, SAS Forecast Studio does not honor the unlocked override that you specified for the Clothes node.
Region	The values of the final forecasts for the nodes in the Product Category level determine the final forecast for the Southeast node. The combination of the final forecasts for the Toys node and the Clothes node equals 160, so the final forecast for the Southeast node is 160.

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

Purpose of the Example

This example shows the following:

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

Setup for the Example

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

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

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 133.

Explanation of the Reconciliation Process

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

Figure 14.10 Example 5: Implicit Top-Down Pass in the Bottom-Up Method of Reconciliation

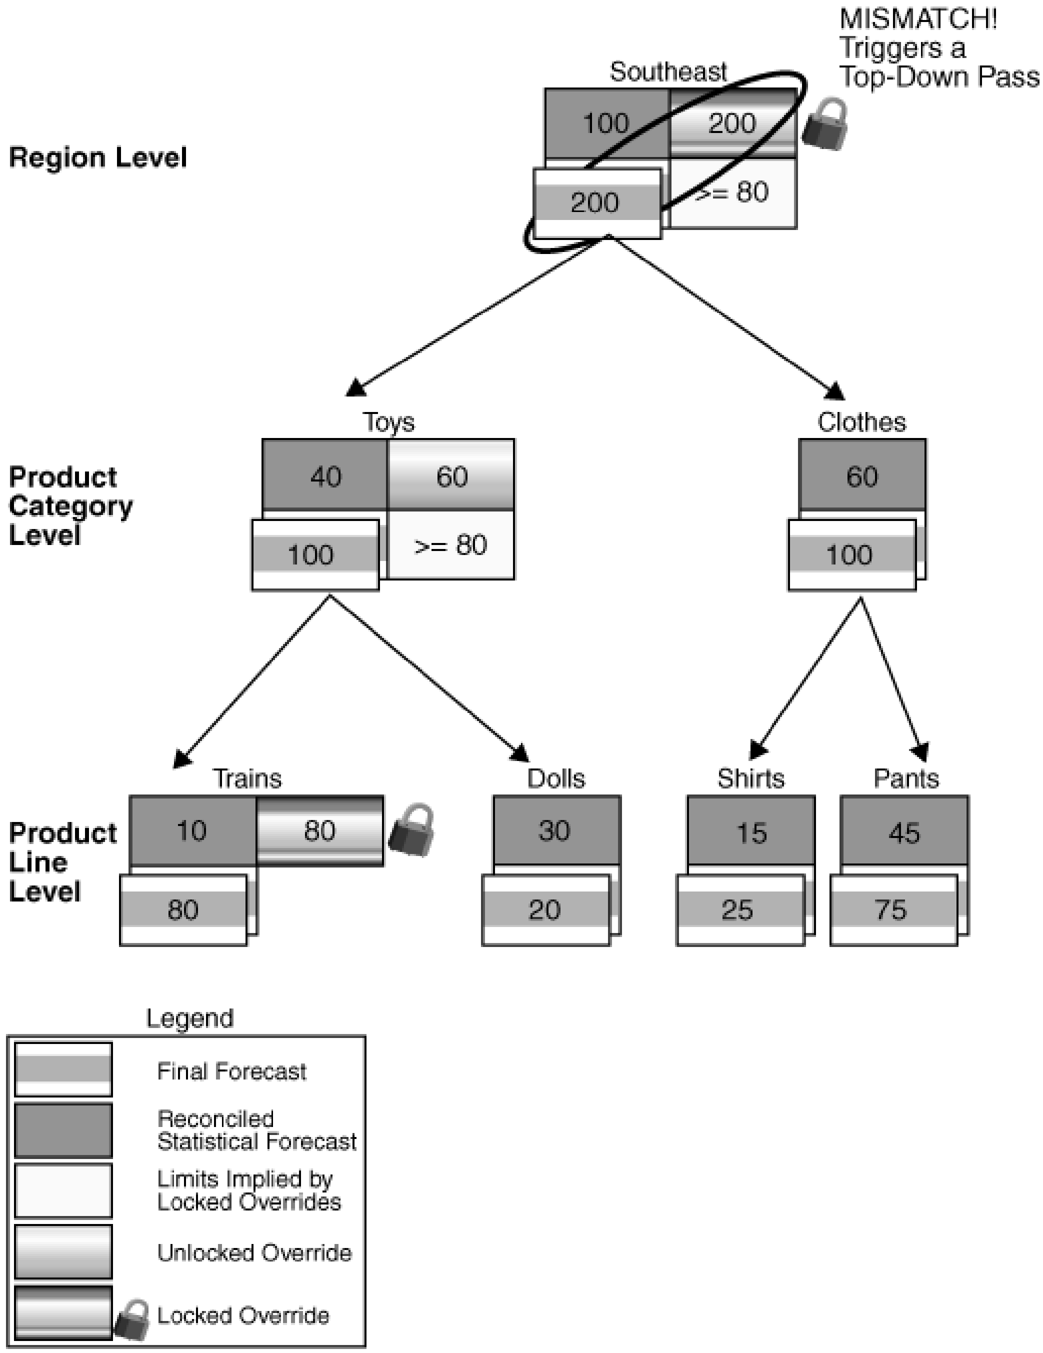

The following table explains the bottom-up reconciliation process for each level of the hierarchy:

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

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

Level in the hierarchy	**Reconciliation process for the nodes in this level**
Region	Because SAS Forecast Studio is now performing an implicit top-down reconciliation, it replaces the final forecast of 170 that was calculated by the bottom-up method of reconciliation with the 200 that you specified as the locked override.
Product Category	You specified an unlocked override of 60 for the Toys node. SAS Forecast Studio uses this unlocked override to determine the proportions during the implicit top-down pass. The unlocked override that you specified for the Toys node is 60, and the final forecast for the Clothes node is 60. Because these nodes are equal, SAS Forecast Studio divides the 200 value from the Southeast node equally between these two 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

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 133.

Explanation of the Reconciliation Process

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

Figure 14.11 Example 6: Implied Limits Have Precedence over an Unlocked Override

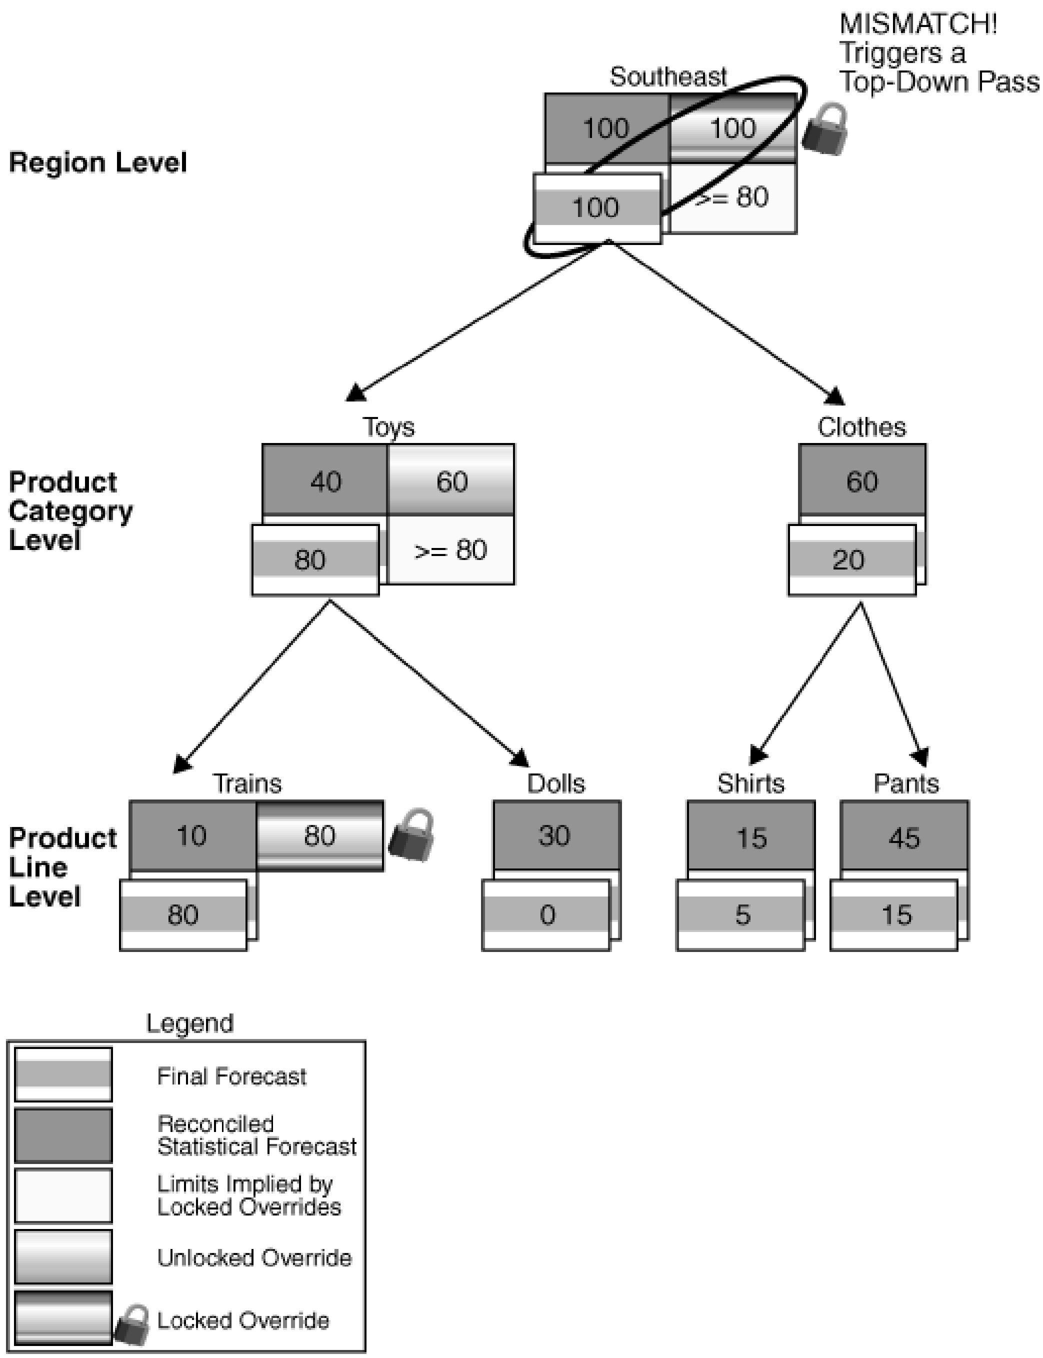

The following table explains the bottom-up reconciliation process for each level of the hierarchy:

Level in the hierarchy	**Reconciliation process for the nodes in this level**
Product Line	Because you selected **Bottom Up** as the reconciliation method, SAS Forecast Studio starts reconciling the hierarchy at the lowest level in the hierarchy. In this example, the lowest level is the Product Line level. In the Product Line level, there are four nodes. The two nodes below the Toys node are the Trains and Dolls nodes. The two 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 two 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

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 133.

Explanation of the Reconciliation Process

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

Figure 14.12 Example 7: Middle-Out Method of Reconciliation with No Conflicts

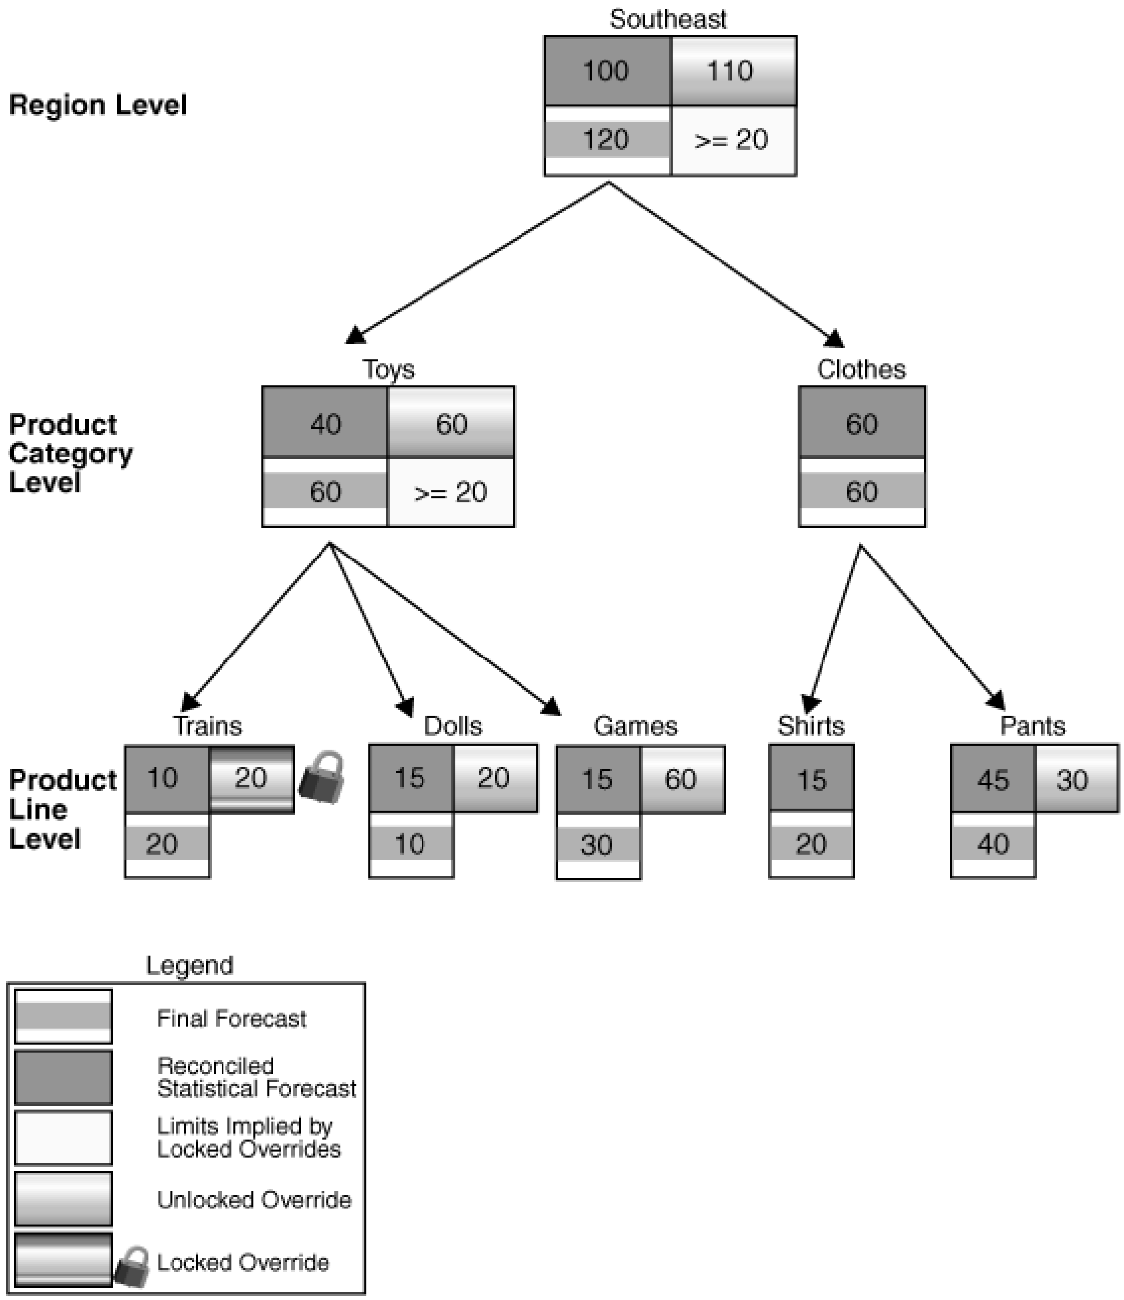

The following table explains the reconciliation process for each level of the hierarchy:

Level in the hierarchy	**Reconciliation process for the nodes in this level**
Product Category	Because you selected **Middle Out** as the reconciliation method, SAS Forecast Studio starts reconciling the hierarchy at the Product Category level. In the Product Category level, there are two 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 five nodes. The three nodes below the Toys node are the Trains, Dolls, and Games nodes. The two 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

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 133.

Explanation of the Reconciliation Process

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

Figure 14.13 Example 8: Middle-Out Method of Reconciliation with an Unreconciled Node

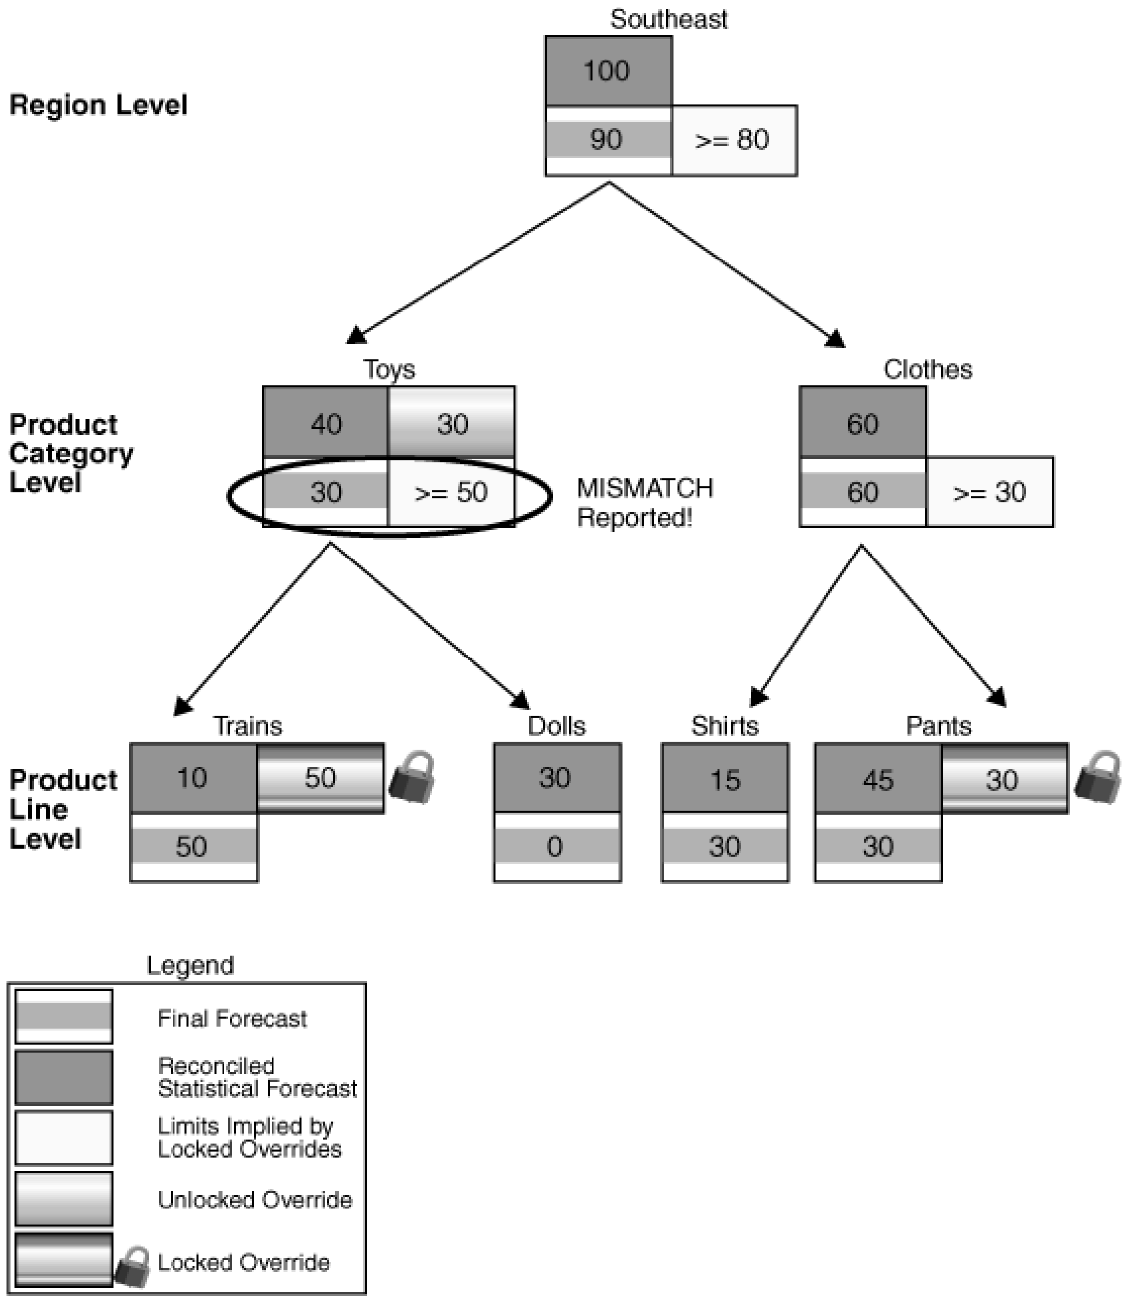

The following table explains the reconciliation process for each level of the hierarchy:

Level in the hierarchy	**Reconciliation process for the nodes in this level**
Product Category	Because you selected **Middle Out** as the reconciliation method, SAS Forecast Studio starts reconciling the hierarchy at the Product Category level. In the Product Category level, there are two 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 dialog box. For the Clothes node, the final forecast is equal to the reconciled statistical forecast of 60. Because 60 is greater than the limits that are implied by a locked override in the lower level of the hierarchy, there are no unreconciled nodes in this branch of the hierarchy.
Region	SAS Forecast Studio calculates the final forecast for the Southeast node by using the bottom-up method of reconciliation. Because you had a conflict in the Toys node in the Product Category level of the hierarchy, SAS Forecast Studio uses the limit that is implied by the locked override in the Trains node to determine the final forecast for the Southeast node. The final forecasts for the Toys node (30) and the Clothes node (60) are combined to create the final forecast for the Southeast node, so the final forecast for the Southeast node is 90.

Level in the hierarchy	**Reconciliation process for the nodes in this level**
Product Line	SAS Forecast Studio calculates the final forecast for the nodes in the lowest level in the hierarchy by using the top-down method of reconciliation. In the Product Line level, there are four nodes. The two nodes below the Toys node are the Trains and Dolls nodes. The two 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 15

Analyzing the Time Series

Contents

Modifying the Properties of the Time ID Variable

Modify the Interval and Seasonal Cycle Length

You assign a variable to the time ID role when you create a project. You cannot change the time ID variable after the project is correct. For more information about the time ID role, see "Time ID Variables" on page 40.

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

To modify the interval and seasonal cycle time ID variable, complete the following steps:

1. Select **Project** → **Forecasting Settings**. The Forecasting Settings dialog box opens.

2. In the selection pane, click **Time ID**.

3. Select the interval of the values for the time ID variable. If you select **Weekday** as the interval, then you can select which days are the weekend or inactive days in the week by clicking **Weekend**.

4. Specify 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.

5. Specify the shift or offset for the interval. The units of the shift depend on the interval.

6. Specify the length of the 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.

7. (Optional) Edit the format of the time ID variable.

8. Click **OK**.

Examples of Changing the Time Interval

Example 1: Change the First Day of the Week

In SAS Forecast Studio, the first day of the week is Sunday. However, in your company, Wednesday is the first day of the week. To accurately forecast your data, SAS Forecast Studio needs to use an interval of WEEK.2.

To change the first day of the week to Wednesday, complete the following steps:

1. Select **Project** → **Forecasting Settings**. The Forecasting Settings dialog box opens.

2. In the selection pane, click **Time ID**.

3. From the **Interval** drop-down list, select **Week**.

4. In the **Shift** box, type `4`.

5. Click **OK**.

For more information about the numbers that SAS Forecast Studio assigns to each day of the week, see "Time Intervals in SAS Forecast Studio" on page 40.

Example 2: Specify a Yearly Interval

In SAS Forecast Studio, a year is from January 1-December 31. However, for a couple of years, your company's fiscal year began in June and ended in May. To accurately forecast your data, SAS Forecast Studio needs to use an interval of YEAR2.6.

To specify an interval of YEAR2.6, complete the following steps:

1. Select **Project** → **Forecasting Settings**. The Forecasting Settings dialog box opens.
2. In the selection pane, click **Time ID**.
3. From the **Interval** drop-down list, select **Year**.
4. In the **Multiplier** box, type `2`. The multiplier specifies that there are two years in the interval.
5. In the **Shift** box, type `6`. The shift specifies that the year should begin at the sixth month (or June).

Figure 15.1 Specifying June as the Beginning of the Year for Two Years

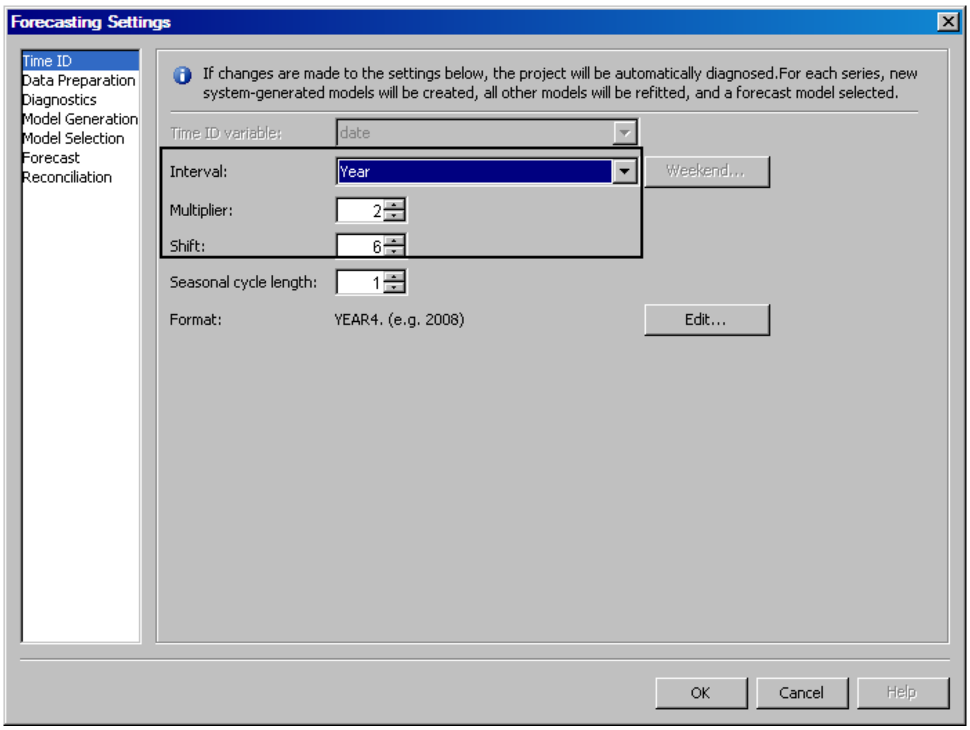

6. Click **OK**.

Example 3: Specify an Interval with a Weekend

If your data is weekly, SAS Forecast Studio assumes that you have data for every day in the week. If your company is closed on a weekend (such as Friday, Saturday, and Sunday), your data does not contain any observations for these three days. You can specify a four day week (Monday-Thursday) by specifying the days in a weekend in SAS Forecast Studio.

To specify the days in the weekend, complete the following steps:

1. Select **Project** → **Forecasting Settings**. The Forecasting Settings dialog box opens.
2. In the selection pane, click **Time ID**.
3. Select the time ID variable in the input data source.
4. From the **Interval** drop-down list, select **Weekday**.
5. Click **Weekend**. The Weekend dialog box opens.
6. Select the check boxes for **Sunday**, **Friday**, and **Saturday**. Click OK.

Figure 15.2 Weekend Days Are Friday, Saturday, and Sunday

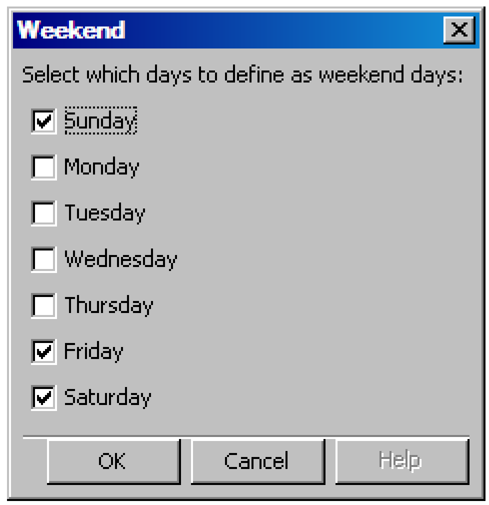

7. In the Forecasting Settings dialog box, click **OK**.

After SAS Forecast Studio reforecasts your project, select **Project** → **SAS Code** to view the SAS code for the project. In the code, you see that the interval equals WEEKDAY167W where 1=Sunday, 6=Friday, and 7=Saturday.

Edit the Format of the Time ID Variable

SAS Forecast Studio automatically determines the date format from the time ID variable. However, you can change this default format. For more information about the time ID formats, see *SAS Language Reference: Dictionary*.

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

1. Select **Project** → **Forecasting Settings**. The Forecasting Settings dialog box opens.

2. In the selection pane, click **Time ID**.
3. Click **Edit**. The Edit Time ID Format dialog box opens.
4. Select the format that you want to use.

   NOTE: The WEEKU*w*., WEEKV*w*., and WEEKW*w*. formats are not supported.
5. Click **OK** in the Edit Time ID Format dialog box.
6. Click **OK** in the Forecasting Settings dialog box.

View the Time Series Plots

From the Series 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. Open the Series View and click **Plots**.
2. From the drop-down list, select the plot that you want to view.

View the Time Series Tables

From the Series 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 View, complete the following steps:

1. Open the Series View and click **Tables**.
2. From the drop-down list, select the table that you want to view.

Transform a Time Series

To improve forecasting accuracy, you can transform a time series. Parameter estimation is performed using the transformed series. The predictions and confidence limits for the transformed model are calculated from the transformed time series and these parameter estimates.

To transform a time series, complete the following steps:

1. Open the Series View.
2. Select the functional transformation for the dependent and independent variables. If you select **Box-Cox** as the functional transformation, then you must specify a parameter value between -5 and 5.
3. Specify the simple difference.
4. Specify the seasonal difference.

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

Viewing the Results of a Time Series Decomposition

About Time Series Decomposition

Using several plots and tables in the Series View, you can perform seasonal decomposition on the working series. Classic decomposition breaks the time series down into trend, cycle, seasonal, and irregular components. The trend and cycle components are often combined to form the trend-cycle component.

Seasonal decomposition includes a seasonal component and additional decomposition methods. The decomposition method determines the mode of the seasonal adjustment decomposition to be performed. In SAS Forecast Studio, there are four decomposition methods: additive, log-additive, pseudo-additive, and multiplicative.

Table 15.1 Seasonal Adjustment Formulas

Component	Decomposition Method	Formula
original series	multiplicative	$O_t = TC_t S_t I_t$
	additive	$O_t = TC_t + S_t + I_t$
	log-additive	$log(O_t) = TC_t + S_t + I_t$
	pseudo-additive	$O_t = TC_t(S_t + I_t - 1)$
trend-cycle component	multiplicative	centered moving average of O_t
	additive	centered moving average of O_t
	log-additive	centered moving average of $log(O_t)$
	pseudo-additive	centered moving average of O_t
seasonal-irregular component	multiplicative	$SI_t = S_t I_t = O_t / TC_t$
	additive	$SI_t = S_t + I_t = O_t - TC_t$
	log-additive	$SI_t = S_t + I_t = log(O_t) - TC_t$
	pseudo-additive	$SI_t = S_t + I_t - 1 = O_t / TC_t$
seasonal component	multiplicative	seasonal Averages of SI_t
	additive	seasonal Averages of SI_t
	log-additive	seasonal Averages of SI_t
	pseudo-additive	seasonal Averages of SI_t
irregular component	multiplicative	$I_t = SI_t / S_t$
	additive	$I_t = SI_t - S_t$
	log-additive	$I_t = SI_t - S_t$
	pseudo-additive	$I_t = SI_t - S_t + 1$
trend-cycle-seasonal component	multiplicative	$TCS_t = TC_t S_t = O_t / I_t$
	additive	$TCS_t = TC_t + S_t = O_t - I_t$
	log-additive	$TCS_t = TC_t + S_t = O_t - I_t$
	pseudo-additive	$TCS_t = TC_t S_t$
trend component	multiplicative	$T_t = TC_t - C_t$
	additive	$T_t = TC_t - C_t$
	log-additive	$T_t = TC_t - C_t$
	pseudo-additive	$T_t = TC_t - C_t$
cycle component	multiplicative	$C_t = TC_t - T_t$
	additive	$C_t = TC_t - T_t$
	log-additive	$C_t = TC_t - T_t$
	pseudo-additive	$C_t = TC_t - T_t$
seasonally adjusted series	multiplicative	$SA_t = O_t / S_t = TC_t I_t$
	additive	$SA_t = O_t - S_t = TC_t + I_t$
	log-additive	$SA_t = O_t / exp(S_t) = exp(TC_t + I_t)$
	pseudo-additive	$SA_t = TC_t I_t$

The trend-cycle component is computed from the centered moving average for the s-period as follows:

$$TC_t = \sum_{k=-\lfloor s/2 \rfloor}^{\lfloor s/2 \rfloor} y_{t+k} / s$$

The seasonal component is obtained by averaging the seasonal-irregular component for each season where $0 \le j \le T/s$ and $1 \le k \le s$.

$$S_{k+js} = \sum_{t=k \bmod s} \frac{SI_t}{T/s}$$

The seasonal components are normalized to sum to one (multiplicative) or zero (additive).

How to View the Time Series Decomposition

1. Open the Series View.
2. Select one of the following plots or tables:
   - decomposition panel
   - a components plot, such as cycle component, irregular component, seasonal-irregular component, seasonal component, trend-cycle-seasonal component, trend-cycle component, or trend component
   - seasonal adjusted series
   - seasonally adjusted series (percent change)

Here is an example of how the decomposition plots and tables appear in the Series View:

Figure 15.3 Decomposition Plots and Tables in Series View

Chapter 16

Working with Models

Contents

Overview of Model Repositories

What Is a Model Repository?

All of the models that are available to a SAS Forecast Studio project are stored in a model repository. 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 Appendix B, "Reserved Names in SAS Forecast Studio."

Types of Models in the Model Repository

To open the model repository, select **Project** → **Model Repository**. SAS Forecast Studio has default models, custom models, and models from an external list. You can determine the model type from the value in the **Type** column.

Figure 16.1 Model Repository

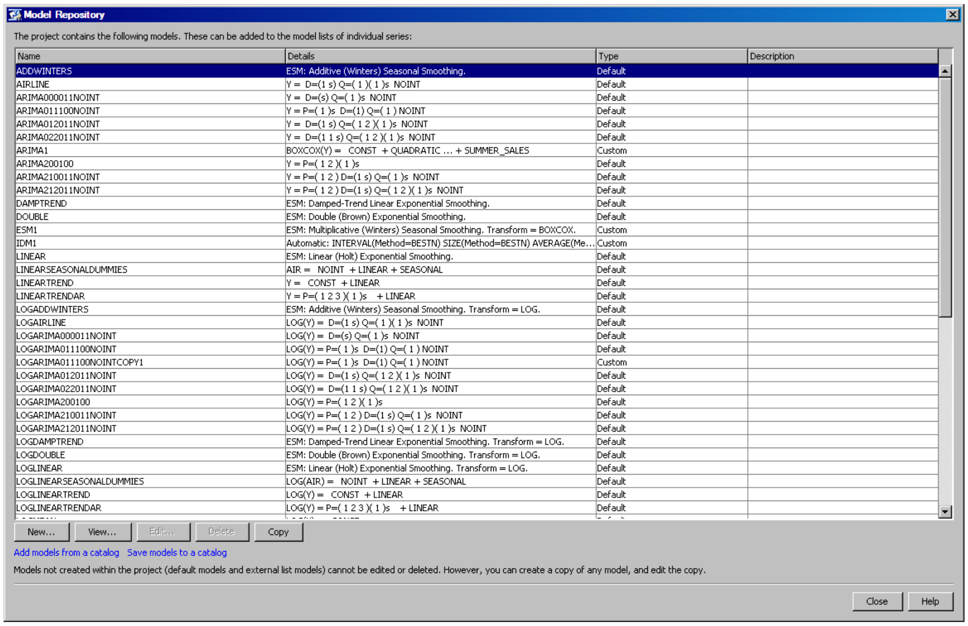

SAS Forecast Studio supports the following types of models:

default models
: 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.

custom models
: models that you create in SAS Forecast Studio. When you create a model for a particular series, that model is automatically added to this model repository. After the model is added, you can apply it to other series in the project. Because these are user-defined models, you can add, edit, copy, or delete custom models.

models from an external list
: models from the external list that was specified in the **Model Generation** panel of the Forecasting Setting dialog box. You cannot edit or delete these models. However, you can copy a model from an external list and customize it to meet your needs.

How SAS Forecast Studio Fits Models

For a project, you can specify the classes of models that SAS Forecast Studio should try to fit to the data. These models make up the model selection list for a series. When SAS Forecast Studio diagnoses a project, SAS Forecast Studio attempts to fit all the models in the model selection list. SAS Forecast Studio runs a series of diagnostics to determine the characteristics of the data (such as seasonality or intermittency), and it avoids models that are inappropriate for the data.

- If the diagnostics determine that a series is intermittent, then you cannot use continuous time series models, such as ARIMA, exponential smoothing, or unobserved components models.
- If the diagnostics determine that a series is continuous, then you cannot use an intermittent time series model, such as the intermittent demand model.

SAS Forecast Studio uses the value of the selection criterion to determine the best-performing model, or also called the forecast model. The results of the forecast model are displayed in the forecast plot and data table in the Forecasting View. In the model selection list, the best-performing model appears in bold.

If none of the models in the model selection list fit, then SAS Forecast Studio attempts to use the model selection list that ships with SAS Forecast Studio. This default model selection list consists of exponential smoothing models.

NOTE: : The status of a model can change depending on the other models in the model selection list. If you change a model selection list (for example, you add, edit, or delete a model from the list or you reforecast a series) and you have selected the **Perform seasonality test** option in the Forecasting Settings dialog box, then the seasonality test that is used to re-create the model selection list is more rigorous than the test that is used when SAS Forecast Studio diagnoses and forecasts the entire project. The more rigorous test could result in models that previously fit appearing as failed in the model selection list. For more information about the **Perform seasonality test** option, see "Set the Diagnose Properties for Each Series" on page 72.

Select the Models to Include in the Model Selection List

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.

To select the models to fit to each series, complete the following steps:

1. Select **Project** → **Forecasting Settings**. The Forecasting Settings dialog box opens.
2. In the selection pane, click **Model Generation**.
3. Select the system-generated models to create. By default, SAS Forecast Studio creates an ARIMA model and an exponential smoothing model. You can also choose to create unobserved components models.
4. To use an external model selection list that contains the models that you want to use, select the **Models from an external list** check box and click **Browse**. The Select Model Selection List dialog box opens, and you can select the model selection list that you want to use.
5. Specify whether to fit the exponential smoothing models that SAS Forecast Studio generates only at the lowest level of the hierarchy.
6. Click **OK**.

Specifying the Forecast Model

What Is the Forecast Model?

The forecast model is the model that SAS Forecast Studio uses to generate a numerical prediction of the future values for the time series. By default, SAS Forecast Studio chooses the best-performing model in the model selection list as the forecast model. The best-performing model is determined by the holdout sample and the selection criterion that you specified. The best-performing model can change over time as the performance changes for the other models in the list.

You can override the automatic selection by SAS Forecast Studio and specify the forecast model. When you reforecast the project, SAS Forecast Studio selects the forecast model based on the current forecast model and how you specified to reforecast the data. In SAS Forecast Studio, you can specify where to start the process for generating a forecast. For more information about each of the steps in this process, see “Process for Generating a Forecast” on page 71.

- You select a system-generated model as the forecast model. Then reforecast the project and have SAS Forecast Studio start at the Diagnose step. You also specify that SAS Forecast Studio should ignore your selection for the forecast model. In this case, SAS Forecast Studio selects the best-performing model.
- You change the forecast model for one or more series. Then you reforecast the project and have SAS Forecast Studio start at the Select step. SAS Forecast Studio keeps the forecast model that you selected rather than selecting the best-performing model from all the fitted models. If you specify that SAS Forecast Studio should ignore your selection for the forecast model, then SAS Forecast Studio selects the best-performing model.

For more information about the reforecasting options, see “Reforecasting a Series or Project” on page 75.

Select the Forecast Model

By default, SAS Forecast Studio determines the best-performing model. However, if you want to select a forecast model, complete the following steps:

1. Select **View** → **Details** → **Modeling View**.
2. In the model selection list, select the model that you want to use as the forecast model and select **Set this model as forecast model**.

When you change the forecast model for a series that is part of a hierarchy, a "Reconciliation needed" warning appears. Click **Reconcile** to reconcile the hierarchy.

Assessing the Accuracy of the Models

The Selection Criterion

SAS Forecast Studio uses the selection criterion to determine the most accurate model. One of the most common selection criteria in business forecasting is the mean absolute percent error (MAPE). The MAPE is the average of all the individual absolute percent errors. Therefore, this criterion shows the size of the forecast error relative to the magnitude of the actual value. When MAPE is your selection criterion, the model with the smallest MAPE value is the best-performing model. In SAS Forecast Studio, MAPE is the default selection criterion.

You should change the selection criterion after forecasts are generated only to investigate the robustness of the model selection. If the same model is the best performing when the selection criterion is MAPE or when the selection criterion is Root Mean Squared Error (RMSE), then you can say that the model selection is robust with respect to a selection criterion of MAPE or RMSE. Choosing a selection criterion after the forecasts are generated can result in forecast bias.

Using a Holdout Sample and an Out-of-Sample Region

The accuracy of the model can be calculated for a holdout sample of data at the end of each time series that was not used to construct models. Using a holdout sample to judge accuracy is often referred to as an honest assessment because it simulates fitting and deploying a model and then judging the accuracy in a live environment. Although using a holdout sample is preferred, it might not be feasible. For example, a time series might be too short to enable the use of a holdout sample.

It is recommended that you use an out-of-sample region along with the holdout sample to determine how the model performs with future forecasts.

Specify the Holdout Sample and Selection Criterion for Model Selection

To specify a holdout sample and the selection criterion for model selection, complete the following steps:

1. Select **Project** → **Forecasting Settings**. The Forecasting Settings dialog box opens.
2. In the selection pane, click **Model Selection**.
3. To specify a holdout sample to be used for model selection, select the **Use holdout sample for model selection** check box. The holdout sample is a subset of the dependent time series that

ends at the last nonmissing observation. The statistics for the selection criterion are computed using only the holdout sample. The default value is 2.

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

If you specify both the holdout sample to use for model selection and the maximum percentage of the series for the holdout sample, then the size of the holdout sample is *min*(*holdout-sample*,*maximum percentage*T), where T is the length of the dependent time series with beginning and ending missing values removed.

4. From the **Selection criterion** drop-down list, select 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 D, "Statistics of Fit."

Specify an Out-of-Sample Region

To specify an out-of-sample region, complete the following steps:

1. Select **Project** → **Forecasting Settings**. The Forecasting Settings dialog box opens.
2. In the selection pane, click **Forecast**.
3. In the **Calculate statistics of fit over an out-of-sample range** box, specify 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.

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** → **Details** → **Modeling View** to open the Modeling View.
2. From the model selection list, select a model.
3. Click **Plots** and select the plot that you want to view.

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.

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** → **Details** → **Modeling View** to open the Modeling View.
2. From the model selection list, select a model.
3. Click **Tables** and select the table that you want to view.

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.

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.

To compare models, complete the following steps:

1. Select **View** → **Details** → **Modeling View**.
2. To compare all of the models for the series, select **Series** → **Compare Models**. The Compare Models dialog box opens.
3. Compare the plots and statistics of fit for each model in the series.

   Each model is plotted in a different color.

Figure 16.2 Plots Tab in Compare Models Dialog Box

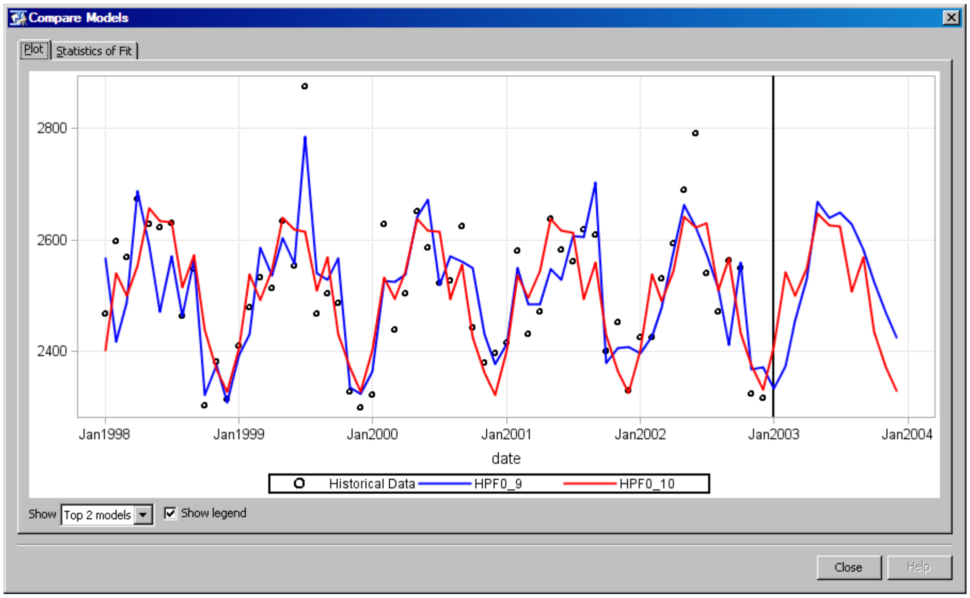

For each statistic of fit, the cell for the best-performing model is shaded in green. For more information about the available statistics, see Appendix D, "Statistics of Fit."

Figure 16.3 Statistics of Fit Tab in Compare Models Dialog Box

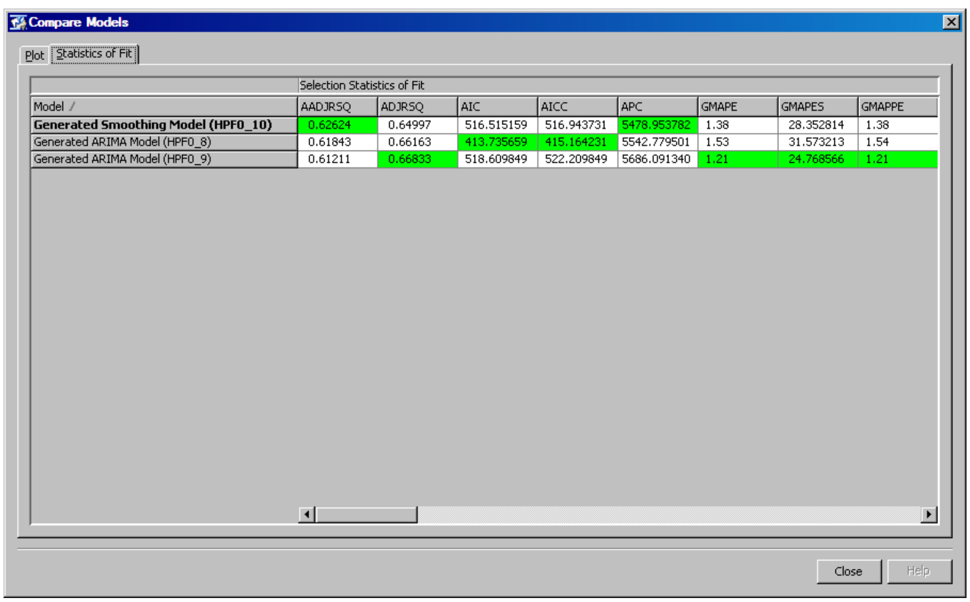

	Selection Statistics of Fit							
Model	AADJRSQ	ADJRSQ	AIC	AICC	APC	GMAPE	GMAPES	GMAPPE
Generated Smoothing Model (HPF0_10)	0.62624	0.64997	516.515159	516.943731	5478.953782	1.38	28.352814	1.38
Generated ARIMA Model (HPF0_8)	0.61843	0.66163	413.735659	415.164231	5542.779501	1.53	31.573213	1.54
Generated ARIMA Model (HPF0_9)	0.61211	0.66833	518.609849	522.209849	5686.091340	1.21	24.768566	1.21

4. Click **Close** after you have finished your comparison.

Add a Model to a Series

If you are working in the model selection list, then SAS Forecast Studio adds the model to the current series.

To add a model to the series that is currently being displayed in the Modeling View, complete the following steps:

1. Select **Series** → **Add Model** or click 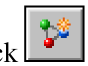 in the model selection list in the Modeling View. The Add Model dialog box opens.
2. Specify how to add a new model to the series. You can create a new model or fit an existing model.
3. Click **OK**.

Importing and Exporting Models

Add Models from a Catalog

To add models to the model repository from a catalog, complete the following steps:

1. Select **Project** → **Model Repository**. The Model Repository dialog box opens.
2. Click **Add models from a catalog**. The Import Model Repository dialog box opens.
3. Select the catalog that contains the models that you want to add and click **OK**. The new models are added to the model repository.

Save Models to a Catalog

You might want to share a custom model with other users at your site. When you save these models to a catalog, the models still remain in the model repository for your project. If you edit these models, these changes are not saved to the models in the catalog. The changes appear only in your project.

To export a model, complete the following steps:

1. Select **Project** → **Model Repository**. The Model Repository dialog box opens.

2. Select the models that you want to export and click **Save models to a catalog**. The Export Model Repository dialog box opens.
3. Select the catalog where you want to save the models and click **OK**.

Create a New Model Based on a Default Model

You can use the default models that ship with SAS Forecast Studio as a basis for a user-defined model.

To create a new custom model based on a default model, complete the following steps:

1. Select **Project** → **Model Repository**. The Model Repository dialog box opens.
2. Select the default model that you want to use as the basis for the custom model and click **Copy**. 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.
3. Edit the model options and click **OK**. For more information about the options that are available for each model, see Chapter 17, "Creating User-Defined Models."

SAS Forecast Studio saves the new model and displays it in the Model Repository dialog box.

Edit a Model

About Editing Models

In SAS Forecast Studio, you can edit a model in the following ways:

- from the model selection list in the Modeling View
- from the model repository

When you edit a model that is used by multiple series, then SAS Forecast Studio automatically updates the model in all of the series. You can see if a model is used in multiple series by using the search feature in the model specification dialog box.

Figure 16.4 Searching for All Series That Use the Current Model

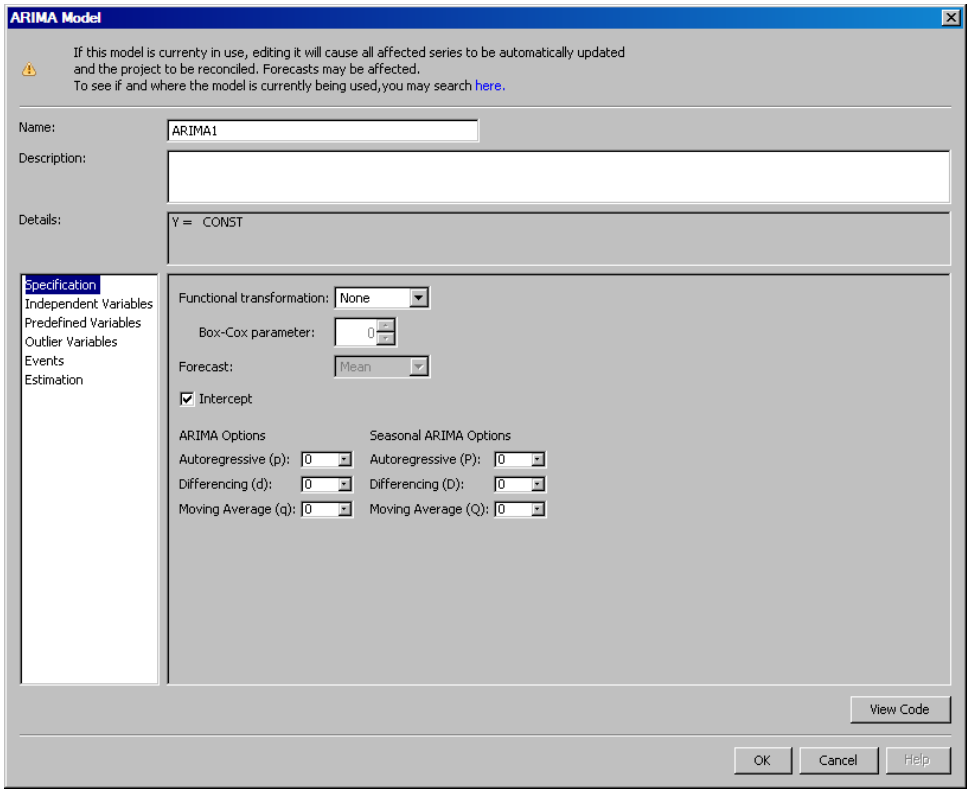

After the search is complete, you can specify whether to update the model for the current series or for all series that use the model.

Edit a Model from the Model Selection List

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

1. Select **View** → **Details** → **Modeling View** to open the Modeling View. The model selection list appears at the top of this view.

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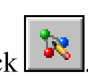. The model specification dialog box for that model opens.

3. In the model specification dialog box, specify the options that you want to edit and click **OK**. For more information about the options that are available for each model, see Chapter 17, “Creating User-Defined Models.”

NOTE: You cannot edit the following types of models:

- default models that shipped with SAS Forecast Studio.

- models that SAS Forecast Studio created using the HPFDIAGNOSE procedure. You select the models that SAS Forecast Studio generates from the Models pane in the Forecasting Settings dialog box. For more information, see "Select the Models to Include in the Model Selection List" on page 180.
- imported models. For more information, see "Importing and Exporting Models" on page 186.

Edit a Model from the Model Repository

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

1. Select the series in the hierarchy or table view and select **Project** → **Model Repository**. The Model Repository dialog box opens.
2. Select the custom model that you want to edit from the list and click **Edit**. The model specification dialog box for that model opens.
3. In the model specification dialog box, specify the options that you want to edit and click **OK**. For more information about the options that are available for each model, see Chapter 17, "Creating User-Defined Models."

Copy Models

To copy a model, complete the following steps:

1. Select **View** → **Details** → **Modeling View** to open the Modeling View.
2. In the model selection list, select the model that you want to copy and click 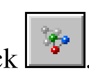.
3. Edit the model specification. By default, the name of the copied model is *model-name*Copy*n* where *model-name* is the original model name and *n* is the lowest integer value available. This name appears in the model selection list unless you edit it. For more information, see "Edit a Model" on page 187.

   NOTE: The copied model appears in the model selection list even if you cancel out of the model specification dialog box.

Remove Models from the Model Selection List

When you delete a model from the model selection list, the model is not deleted from the model repository, but it no longer appears in the model selection list.

To delete a model from the model selection list, select the model name and click 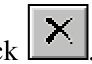. You are prompted to confirm this removal.

NOTE: If you delete all the models in the model selection list, then SAS Forecast Studio automatically fits an exponential smoothing model so that a forecast can be produced for the series. This exponential smoothing model is automatically removed when you add another model to the model selection list.

Chapter 17

Creating User-Defined Models

Contents

ARIMA Models

Overview of ARIMA Models

When you create an Autoregressive Integrated Moving Average (ARIMA) model, you can specify the autoregressive and moving average polynomials of an ARIMA model. You can also include events and independent variables in the model.

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

Create an ARIMA Model

To create an ARIMA model, complete the following steps:

1. Select **Project** → **Model Repository**. The Model Repository dialog box opens.

   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 186.

2. Click **New**. The New Model dialog box opens.

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

Figure 17.1 ARIMA Model Dialog Box

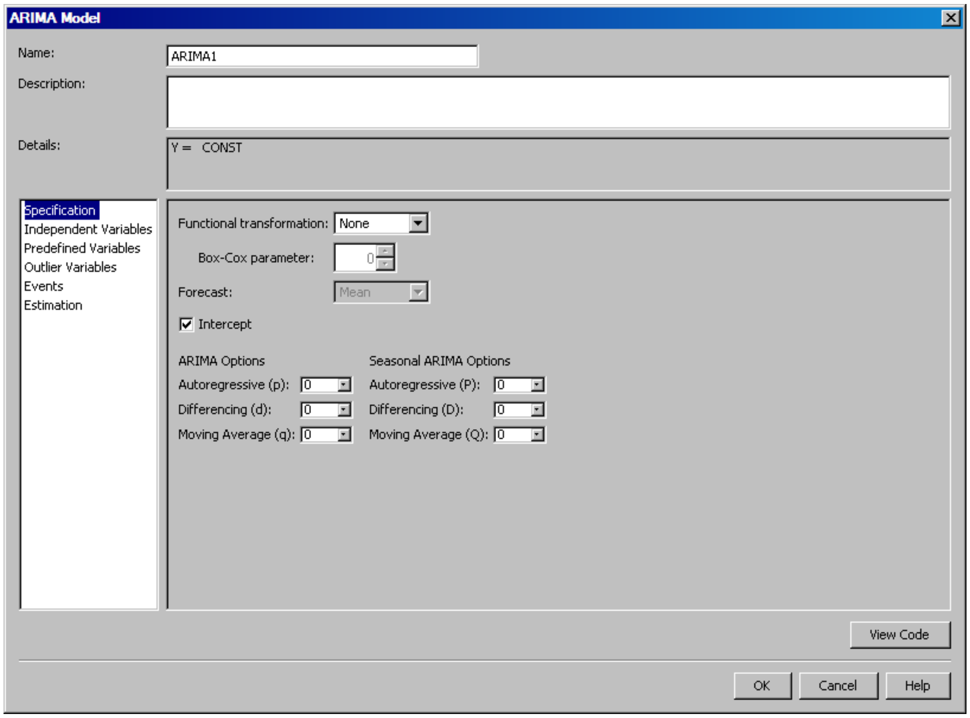

4. Specify a unique name for the model. By default, SAS Forecast Studio automatically generates a unique name for the model.

   NOTE: SAS Forecast Studio has a list of reserved model names. For more information, see Appendix B, "Reserved Names in SAS Forecast Studio."

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

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

6. To specify a basic nonseasonal (ARIMA(p, d, q)) or seasonal (ARIMA(p, d, q)(P, D, Q))s model, select **Specification** in the selection pane. You can specify the following options:

   - which functional transformation to apply to the dependent series.

     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. Specify an integer value for the following orders:

- the nonseasonal and seasonal autoregressive orders (p and P).
  If you set p = 4, then you are implying autoregressive orders of (1 2 3 4) in the nonseasonal component of the model. If you set p = 4 and P = 3, then you are implying autoregressive orders of (1 2 3 4) in the nonseasonal component of the model and autoregressive orders of (1 2 3) in the seasonal component of the model. In the Details pane, this order appears as follows: P = (1 2 3 4)(1 2 3)s.
  NOTE: The seasonal length of the order is identified by the placeholder "s".
- the nonseasonal and seasonal differencing orders (d and D).
  If you set d = 4, then you are implying differencing orders of (1 1 1 1) in the nonseasonal component of the model. If you set d = 4 and D = 3, then you are implying (1 1 1 1) in the nonseasonal component of the model and (s s s) in the seasonal component of the model. In the Details pane, this order appears as follows: D = (1 1 1 1 s s s).
- the nonseasonal and seasonal moving average orders (q and Q).
  If you set q = 4, then you are implying moving average orders of (1 2 3 4) in the nonseasonal component of the model. If you set q = 4 and Q = 3, then you are implying (1 2 3 4) in the nonseasonal component of the model and (1 2 3) in the seasonal component of the model. In the Details pane, this order appears as follows: Q = (1 2 3 4)(1 2 3)s.

NOTE: For a simple nonseasonal ARIMA model, you do not have to specify the orders for the seasonal component (P D Q). The season length (such as monthly, daily, and so on) is implied by the time ID variable in the project.

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

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

7. To include an independent variable in the model, select **Independent Variables** in the selection pane. For each independent variable that you want to include, select the check box in the **Include in Model** column.

   If you include only one independent variable in the model, then click **Transfer Function** to specify whether the independent variable enters the model as a simple regressor or if it enters the model after some modifications, such as lagging or differencing. You can specify these modifications by using the Transfer Function dialog box.

   a) Select the functional transformation to apply to the time series.
   b) Specify the delay (or lag) for the input series.
   c) Specify the differencing orders for the input series. How you specify the differencing order depends on the options that are available.
      - If the **Simple differencing** and **Seasonal differencing** options are available, then you can specify the integer value to use for these differencing orders.
      - If the **Differencing orders** option is available, then you must specify the syntax for the differencing order. For more information about the syntax, see Appendix F, "Autoregressive, Differencing, and Moving Average Orders."

NOTE: For ARIMA and Subset (factored) ARIMA models, the differencing order for new events is the same as the differencing order that is specified for the model. You can view the differencing order for the model in the Details pane of the ARIMA Model dialog box or the Subset ARIMA Model dialog box.

d) Specify whether you want an ordinary or dynamic regression for the numerator and denominator polynomials.

How you specify these polynomials depends on the options that are available.

- If the **Numerator factors** and **Denominator factors** options are available, then you can specify the integer value to use for the simple and seasonal factors.
- If the **Numerator polynomial** and **Denominator polynomial** options are available, then you must use the following syntax to specify the numerator (NUM) and denominator (DEN) polynomial of the transfer function.

NUM= ***order***

NUM= (*lag, . . . , lag*) . . . (*lag, . . . , lag*)

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

DEN= ***order***

DEN= (*lag, . . . , lag*) . . . (*lag, . . . , lag*)

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

For information about how to specify the polynomial order, see the autoregressive (p) option in Appendix F, "Autoregressive, Differencing, and Moving Average Orders."

e) Click **OK**.

8. To include a predefined variable in the model, select **Predefined Variables** in the selection pane. For each predefined variable that you want to include, select the check box in the **Include in Model** column.

9. To include an event in the model, select **Events** in the selection pane. For each event that you want to include in the model, select the check box in the **Include in the Model** column.

10. To specify the estimation options, select **Estimation** in the selection pane.

11. Click **OK**.

Detecting Outliers in ARIMA Models

About Detecting Outliers

By default, SAS Forecast Studio automatically detects outliers in the ARIMA models. These outliers show you changes in the level of the response series that are not accounted for by the model. In SAS Forecast Studio, these changes can be additive outliers or level shifts. For each outlier, SAS Forecast Studio creates a dummy regressor or indicator variable. Then the model and the dummy regressors are fit to the data.

SAS Forecast Studio uses the following iterative process to detect outliers:

1. Identify a significant outlier.
2. Determine whether this outlier falls within the specified significance level.
3. Add the outlier to the model.
4. Continue identifying and adding outliers until no more significant outliers are found or until the number of iterations in this process exceeds the maximum number of outliers that are allowed in the model. You can specify this maximum number as an integer or a percentage value.

Specify How to Detect Outliers

To specify whether to detect outliers, complete the following steps:

1. Select **Project** → **Forecasting Settings**. The Forecasting Settings dialog box opens.
2. In the selection pane, click **Diagnostics**.
3. Select the **Detect outliers** check box to specify whether SAS Forecast Studio should detect outliers in the data when fitting an ARIMA model. In the box, specify the maximum number of outliers that can be included in the model.
4. Specify the significance level for outlier detection. The significance level is the cutoff value for outlier detection. This value can range between 0 and 1.
5. Specify the maximum percentage of values in a series that can consist of outliers.

   If you specify both a maximum number and a maximum percentage for outliers, then the number of outliers is *min*(*maximum-number*,*percentage*T), where *T* is the length of the time series with beginning and ending missing values removed. For example, if you specified the maximum number of outliers as 5 and the maximum percentage of outliers as 10, then the number of outliers is *min*(5,0.1T).

Remove an Outlier from the Model

To remove an outlier, complete the following steps:

- Select the ARIMA model that you want to edit from the model selection list or the model repository. When you click **Edit**, the ARIMA model specification dialog box opens.
- In the selection pane, select **Outlier Variables**. By default, any outliers that were automatically detected during the creation of the model are inputs in the model.
- Select the outlier that you want to delete and click **Delete**.

Subset (Factored) ARIMA Models

Overview of Subset (Factored) ARIMA Models

When creating a Subset (factored) ARIMA model, you can specify a general ARIMA model. You can specify the autoregressive and moving average polynomials of arbitrary complexity. You can also specify a general differencing order. The Subset ARIMA Model dialog box is similar to the ARIMA Model dialog box, except that it uses a more general specification of the ARIMA options.

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

Create a Subset (Factored) ARIMA Model

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

1. Select **Project** → **Model Repository**. The Model Repository dialog box opens.

   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 186.

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

Figure 17.2 Subset ARIMA Model Dialog Box

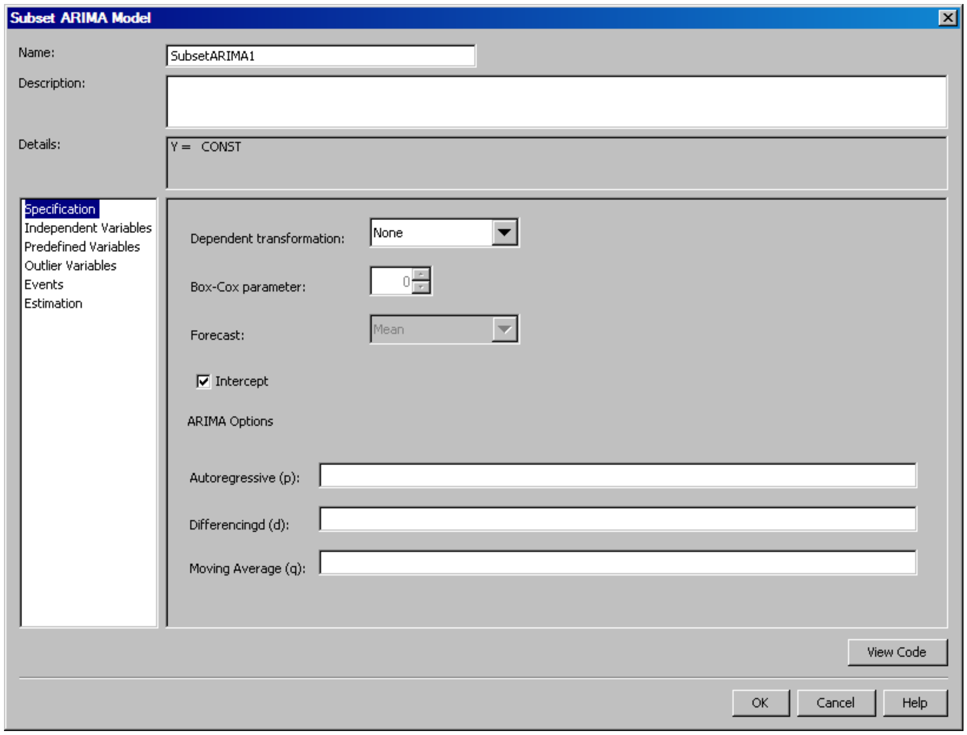

4. Specify a unique name for the model. By default, SAS Forecast Studio automatically generates a unique name for the model.

   NOTE: SAS Forecast Studio has a list of reserved model names. For more information, see Appendix B, "Reserved Names in SAS Forecast Studio."

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

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

6. Select **Specification** in the selection pane to access the following options:

   - which dependent transformation to apply to the dependent series.
   - 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 Appendix F, "Autoregressive, Differencing, and Moving Average Orders."

7. To include an independent variable in the model, select **Independent Variables** in the selection pane. For each independent variable that you want to include, select the check box in the **Include in Model** column.

If you include only one independent variable in the model, then click **Transfer Function** to specify whether the independent variable enters the model as a simple regressor or if it enters the model after some modifications, such as lagging or differencing. You can specify these modifications by using the Transfer Function dialog box.

a) Select the functional transformation to apply to the time series.

b) Specify the delay (or lag) for the input series.

c) Specify the differencing orders for the input series. How you specify the differencing order depends on the options that are available.

   - If the **Simple differencing** and **Seasonal differencing** options are available, then you can specify the integer value to use for these differencing orders.
   - If the **Differencing orders** option is available, then you must specify the syntax for the differencing order. For more information about the syntax, see Appendix F, "Autoregressive, Differencing, and Moving Average Orders."

   NOTE: For ARIMA and Subset (factored) ARIMA models, the differencing order for new events is the same as the differencing order that is specified for the model. You can view the differencing order for the model in the Details pane of the ARIMA Model dialog box or the Subset ARIMA Model dialog box.

d) Specify whether you want an ordinary or dynamic regression for the numerator and denominator polynomials.

   How you specify these polynomials depends on the options that are available.

   - If the **Numerator factors** and **Denominator factors** options are available, then you can specify the integer value to use for the simple and seasonal factors.
   - If the **Numerator polynomial** and **Denominator polynomial** options are available, then you must use the following syntax to specify the numerator (NUM) and denominator (DEN) polynomial of the transfer function.

   NUM= ***order***

   NUM= (*lag, . . . , lag*) . . . (*lag, . . . , lag*)

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

   DEN= ***order***

   DEN= (*lag, . . . , lag*) . . . (*lag, . . . , lag*)

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

For information about how to specify the polynomial order, see the autoregressive (p) option in Appendix F, "Autoregressive, Differencing, and Moving Average Orders."

e) Click **OK**.

8. To include a predefined variable in the model, select **Predefined Variables** in the selection pane. For each predefined variable that you want to include, select the check box in the **Include in Model** column.

9. To delete an outlier from the model, select **Outlier Variables** in the selection pane. By default, any outliers that were automatically detected during the creation of the model are inputs in the model. 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.

10. To include an event in the model, select **Events** in the selection pane. For each event that you want to include in the model, select the check box in the **Include in the Model** column.

11. To specify the estimation options, select **Estimation** in the selection pane.

12. Click **OK**.

Unobserved Components Models

Overview of Unobserved Components Model

Unobserved components models can include events and independent variables.

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

Create an Unobserved Components Model

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

1. Select **Project** → **Model Repository**. The Model Repository dialog box opens.

   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 186.

2. Click **New**. The New Model dialog box opens.

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

Figure 17.3 Unobserved Components Model Dialog Box

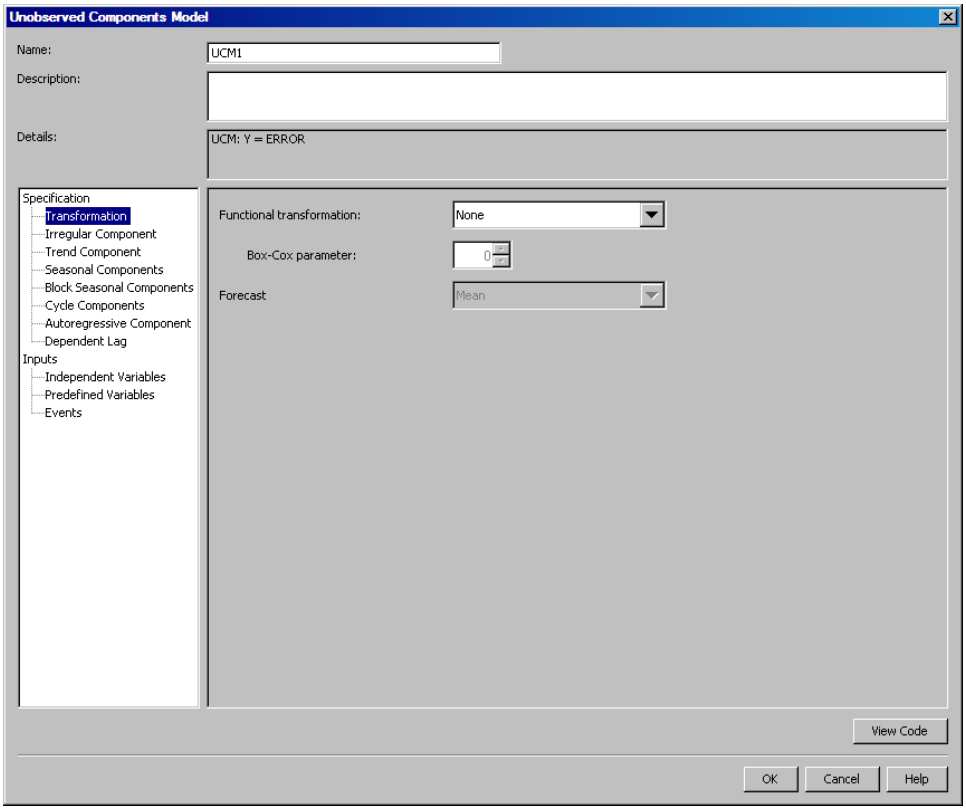

4. Specify a unique name for the model. By default, SAS Forecast Studio automatically generates a unique name for the model.

   NOTE: SAS Forecast Studio has a list of reserved model names. For more information, see Appendix B, "Reserved Names in SAS Forecast Studio."

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

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

6. To specify a transformation, select **Transformation** in the selection pane. You can then specify the following options:

   - the functional transformation for the model.
   - how to calculate the forecasts. Forecasts can be based on the mean or median. By default, the mean value is used.

7. To create an irregular component that corresponds to the overall random error in the model, select **Irregular Component** in the selection pane.

   a) Select the **Irregular component** check box.
   b) (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.

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

8. To define the trend component for the model, select **Trend Component** in the selection pane. 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.

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

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

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

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

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

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

9. To specify a seasonal component, select **Seasonal Components** in the selection pane. You can specify a maximum of three seasonal components.

   a) In the Seasonal Components pane, click **New**. The New Seasonal Component dialog box opens.

   b) Specify the type of seasonal component. A seasonal component can be one of two types: dummy or trigonometric (which is the default).

   c) 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.

   d) Specify the initial value for the disturbance variance.

   e) 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.

   f) Click **OK**.

10. To create a block-seasonal component, select **Block Seasonal Components** in the selection pane. You can specify a maximum of three block-seasonal components.

    a) In the Block Seasonal Components pane, click **New**. The New Blockseasonal Component dialog box opens.

    b) Specify the type of block-seasonal component. A block-seasonal component can be one of two types: dummy (which is the default) or trigonometric.

c) Specify the block size. The block size can be any integer larger than or equal to 2.

d) 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.

e) 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.

f) To supply an initial value for the variance, select the **Define initial variance** check box. Specify the initial value for the disturbance variance.

g) 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.

11. To specify a cycle component, select **Cycle Components** in the selection pane. You can specify up to 50 cycles in a model. By default, the cycle components are estimated from the data. You can also create additional cycle components. A default name is assigned to each cycle component. You cannot change this name.

    a) In the Cycle Components pane, click **New**. The New/Edit Cycle Component dialog box opens.

    b) 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.

    c) 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.

    d) 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.

    e) To fix the values of the component parameters, select the **Do not perform estimation (noest)** check box. This option is available for the initial period, initial damping factor, and initial disturbance variance.

12. To specify an autoregressive component, select **Autoregressive Component** in the selection pane.

    a) Select the **Autoregressive component** check box.

    b) (Optional) Specify an initial damping value to use during the parameter estimation process, select the **Define initial damping** check box, and then specify a value in the box. This value must be between -1 and 1, including -1 but excluding 1.

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

    d) To fix the values of the component parameters to those specified for the damping and variance, select the **Do not perform estimation (noest)** check box.

13. To specify the forecast variable lags to be included a predictors in the model, select **Dependent Lag** in the selection pane.

    a) Select the **Lag dependent variable** check box.

    b) 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)

    c) 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.

    d) 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.

14. To include an independent variable in the model, select **Independent Variables** in the selection pane. For each independent variable that you want to include, select the check box in the **Include in Model** column.

    If you include only one independent variable in the model, then click **Transfer Function** to specify whether the independent variable enters the model as a simple regressor or if it enters the model after some modifications, such as lagging or differencing. You can specify these modifications by using the Transfer Function dialog box.

a) Select the functional transformation to apply to the time series.
b) Specify the delay (or lag) for the input series.
c) Specify the differencing orders for the input series. How you specify the differencing order depends on the options that are available.
   - If the **Simple differencing** and **Seasonal differencing** options are available, then you can specify the integer value to use for these differencing orders.
   - If the **Differencing orders** option is available, then you must specify the syntax for the differencing order. For more information about the syntax, see Appendix F, "Autoregressive, Differencing, and Moving Average Orders."
d) Click **OK**.

15. To include a predefined variable in the model, select **Predefined Variables** in the selection pane. For each predefined variable that you want to include, select the check box in the **Include in Model** column.

16. To include an event in the model, select **Events** in the selection pane. For each event that you want to include in the model, select the check box in the **Include in the Model** column.

17. Click **OK**.

Exponential Smoothing Models

Overview of Exponential Smoothing Models

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

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

Create an Exponential Smoothing Model

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

1. Select **Project** → **Model Repository**. The Model Repository dialog box opens.

   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 186.

2. Click **New**. The New Model dialog box opens.

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

Figure 17.4 Exponential Smoothing Model Dialog Box

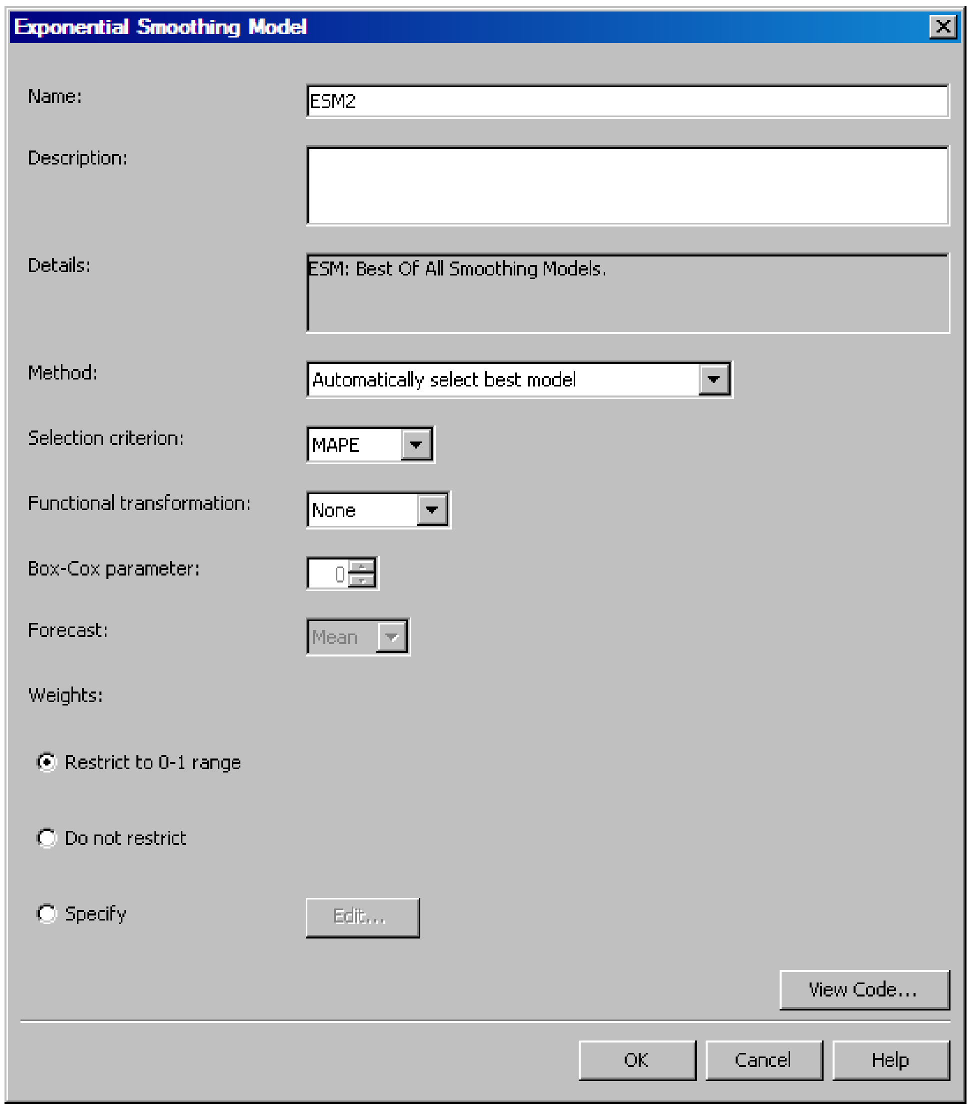

NOTE: The Exponential Smoothing Model dialog box also opens when you edit a model. The options that are available when you are creating or editing a model are the same.

4. 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 Appendix B, "Reserved Names in SAS Forecast Studio."

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

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

6. From the **Method** drop-down list, select the forecasting method to be used to forecast the time series.

7. Select the selection criterion to use. The default selection criterion is the project default.

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

   - You selected one of the best models as the forecasting method.
   - In the **Functional transformation** drop-down list, you select **Auto**, which means that the transformation for the dependent series is automatically chosen based on the model selection criteria.

8. From the **Functional transformation** drop-down list, select the dependent series transformation to be applied to the time series.

   NOTE: You can specify a transformation only when the values of the series are positive.

9. Specify how to calculate the forecasts. Forecasts can be based on the mean or median. By default, the mean value is used.

10. Specify whether to restrict all weights to the 0 to 1 range.

    To specify a smoothing weight, complete the following steps:

    a) Select **Specify** as the weight and click **Edit**. The Edit Weights dialog box opens.
    b) Restrict the values and set the initial value for the parameters. Only the parameters that are relevant for a model are available.
    c) To specify that the model parameters are fixed values, select the **Do not perform estimation (noest)** check box. 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.
    d) To restrict the weights to stable values, select the **Restrict weights to stable values** check box.

Intermittent Demand Models

Overview of Intermittent Demand Models

Intermittent demand models are 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, their models will not work for intermittent time series data.

Create an Intermittent Demand Model

To create an intermittent demand model, complete the following steps:

1. Select **Project** → **Model Repository**. The Model Repository dialog box opens.

   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 186.

2. Click **New**. The New Model dialog box opens.

3. From the drop-down list, select **Intermittent demand** as the model type. The Intermittent Demand Model dialog box opens.

Figure 17.5 Intermittent Demand Model Dialog Box

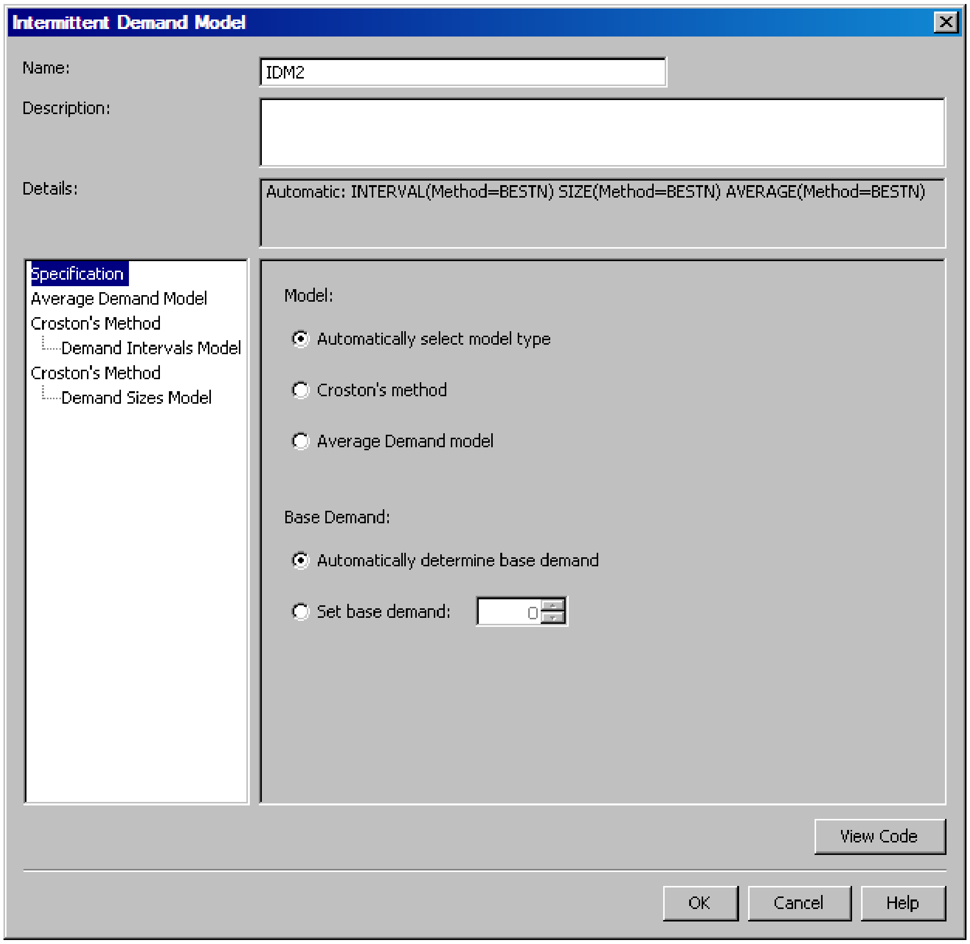

4. Specify a unique name for the model. By default, SAS Forecast Studio automatically generates a unique name for the model.

   NOTE: SAS Forecast Studio has a list of reserved model names. For more information, see Appendix B, "Reserved Names in SAS Forecast Studio."

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

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

6. To specify which intermittent demand method to use, select **Specification** in the selection pane. You can specify the following options:

   - the model type. You can use Croston's method for the model, or you can select an average demand model. You can also have SAS Forecast Studio automatically determine the model type.
   - 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.

7. To create an average demand model, a Croston's method demand intervals model, or a Croston's method demand sizes model, select **Average Demand**, **Demand Intervals Model**, or **Demand Sizes Model** in the selection pane and specify the following options:

   - the forecasting method to use to forecast the time series.
   - the selection criterion to use. The default selection criterion is the project default.

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

     – You selected **Automatically select best non-seasonal model** as the forecasting method.

     – In the **Functional transformation** drop-down list, you select **Auto**, which means that the transformation for the dependent series is automatically chosen based on the model selection criteria.
   - the functional transformation to apply to the time series.
   - 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.

     To specify a smoothing weight, complete the following steps:

     a) Select **Specify** as the weight and click **Edit**. The Edit Weights dialog box opens.

     b) Restrict the values and set the initial value for the parameters. Only the parameters that are relevant for a model are available.

     c) To specify that the model parameters are fixed values, select the **Do not perform estimation (noest)** check box. 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.

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

8. Click **OK**.

Multiple Regression Models

Overview of the Multiple Regression Model

You can create a multiple regression model with autocorrelated errors. You can specify which independent variables and events to include in the model.

Create a Multiple Regression Model

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

1. Select **Project** → **Model Repository**. The Model Repository dialog box opens.

   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 186.

2. Click **New**. The New Model dialog box opens.

3. From the drop-down list, select **Multiple Regression** as the model type. The Multiple Regression Model dialog box opens.

   Figure 17.6 Multiple Regression Model Dialog Box

   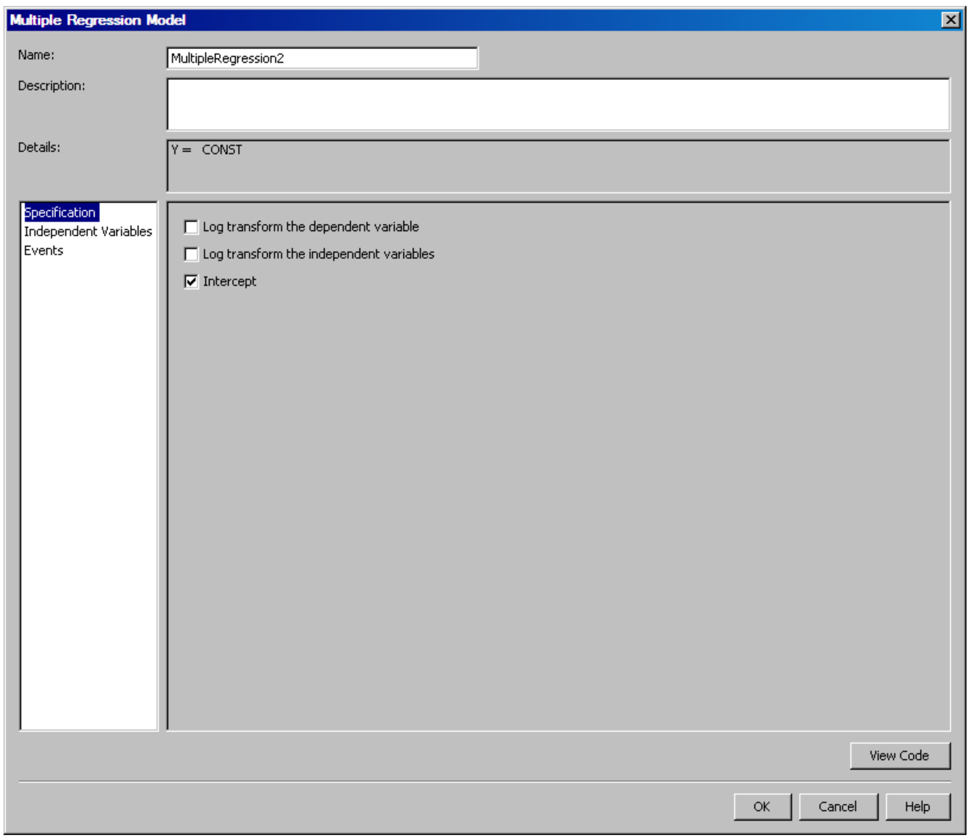

   NOTE: The Multiple Regression Model dialog box also opens when you edit a model. The options that are available when you are creating or editing a model are the same.

4. 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 Appendix B, "Reserved Names in SAS Forecast Studio."

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

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

6. Select **Specification** in the selection pane to specify the following options:

   - whether to log transform the dependent variable.
   - whether to log transform the independent variables.
   - whether to include the intercept in the model. By default, the intercept is included.

7. To include an independent variable in the model, select **Independent Variables** in the selection pane. For each independent variable that you want to include, select the check box in the **Include in Model** column. If you include only one independent variable in the model, then click **Transfer Function** to specify the transfer function options.

8. To include an event in the model, select **Events** in the selection pane. For each event that you want to include in the model, select the check box in the **Include in the Model** column.

9. Click **OK**.

Moving Average Models

Overview of Moving Average Models

Using SAS Forecast Studio, you can create a moving average model. The formula for the moving average model with width k is $y_t = [y_{(t-1)} + \ldots + y_{(t-k)}]/k + error$.

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

Create a Moving Average Model

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

1. Select **Project** → **Model Repository**. The Model Repository dialog box opens.

   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 186.

2. Click **New**. The New Model dialog box opens.

3. From the drop-down list, select **Moving Average** as the model type. The Moving Average Model dialog box opens.

   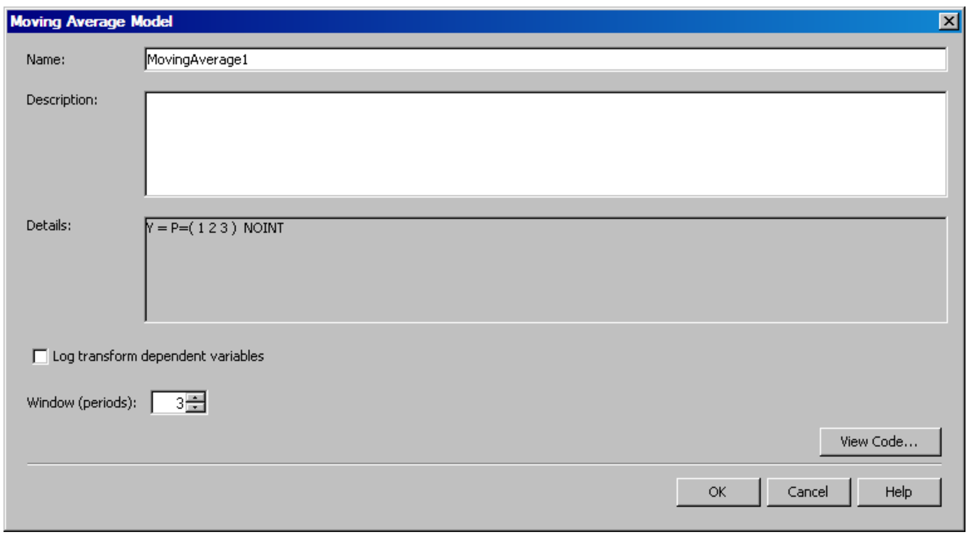

   Figure 17.7 Moving Average Model Dialog Box

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

4. 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 Appendix B, "Reserved Names in SAS Forecast Studio."

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

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

6. Select the **Log transform dependent variables** check box to log transform the dependent variable.

7. In the **Window (periods)** box, specify the number of periods for the moving average. The default value is 3.

8. Click **OK**.

Curve Fitting Models

Overview of Curve Fitting Models

Curve fitting models enable you to identify trends and relationships in your time series data. Using SAS Forecast Studio, you can create a curve fitting model with a linear or quadratic trend.

Create a Curve Fitting Model

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

1. Select **Project** → **Model Repository**. The Model Repository dialog box opens.

   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 186.

2. Click **New**. The New Model dialog box opens.

3. From the drop-down list, select **Curve Fitting** as the model type. The Curve Fitting Model dialog box opens.

Figure 17.8 Curve Fitting Model Dialog Box

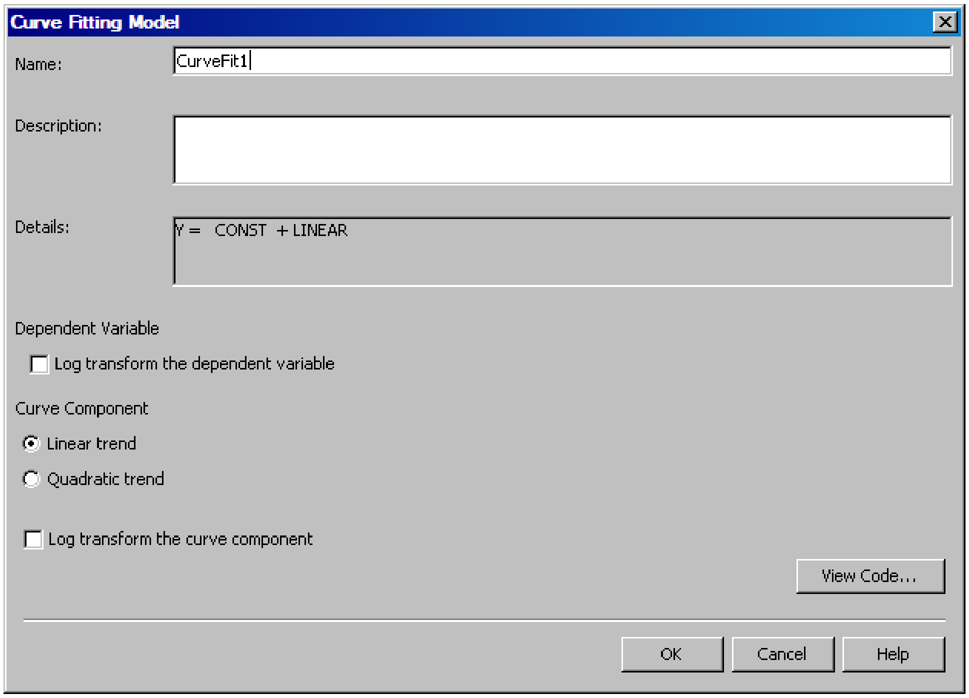

NOTE: The Curve Fitting Model dialog box also opens when you edit a model. The options that are available when you are creating or editing a model are the same.

4. 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 Appendix B, "Reserved Names in SAS Forecast Studio."

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

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

6. Select the **Log transform the dependent variable** check box if you want to log transform the dependent variable.

7. Select the curve component for the model. You can choose from a linear or quadratic trend.

8. Specify whether to log transform the curve component.

9. Click **OK**.

Random Walk Models

Overview of Random Walk Models

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

You can also create the following random walk models:

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

Seasonal Random Walk without Drift
: ARIMA(0, 1, 0)(0, 1, 0)s with no intercept

Seasonal Random Walk with Drift
: ARIMA(0, 1, 0)(0, 1, 0)s

Create a Random Walk Model

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

1. Select **Project** → **Model Repository**. The Model Repository dialog box opens.

   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 186.

2. Click **New**. The New Model dialog box opens.

3. From the drop-down list, select **Random walk** as the model type. The Random Walk Model dialog box opens.

Figure 17.9 Random Walk Model Dialog Box

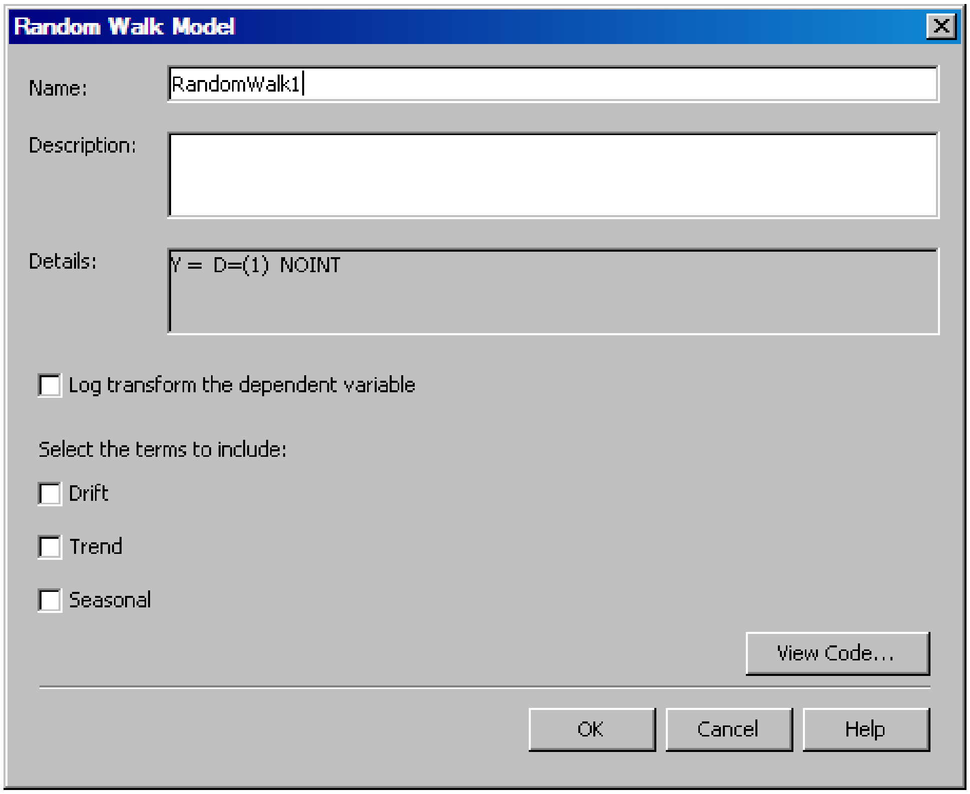

NOTE: The Random Walk Model dialog box also opens when you edit a model. The options that are available when you are creating or editing a model are the same.

4. 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 Appendix B, "Reserved Names in SAS Forecast Studio."

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

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

6. Select whether to log transform the dependent variable. When you select this option, the model uses the data for the log transformed series.

7. Select whether to include a drift term, trend term, and seasonal term in the model.

8. Click **OK**.

Chapter 18

Performing Scenario Analysis

Contents

What Is Scenario Analysis?

SAS Forecast Studio enables you to quickly generate a large number of forecasts. However, you might want to see how the generated forecasts change when you manipulate the future values of the independent variables. For example, what happens to the forecasts when the price increases? This functionality is called scenario analysis and is available from the Scenario Analysis View in SAS Forecast Studio. In this view, you can create different scenarios without having to create multiple SAS Forecast Studio projects.

The Scenario Analysis View enables you to quickly create and run scenarios, compare different scenarios, save scenarios for future use, and apply the forecasted values from a scenario in your project. Although the Scenario Analysis View is integrated in SAS Forecast Studio and uses the data in the current project, the values that you change in this view do not change the generated forecasts in your project. However, integrating this functionality in SAS Forecast Studio ensures that the analysis is completed properly.

The basis of each scenario is a model that contains the input variables that you want to change. In SAS Forecast Studio, you can specify how these input variables are used in model generation. Each model in the project can use all, some, or none of the input variables.

To perform a scenario analysis, SAS Forecast Studio uses the HPFENGINE procedure to generate score files for the forecasts. For more information about the procedures and functions that SAS Forecast Studio uses, see "Using Forecasting Model Score Files and DATA Step Functions" in *SAS High-Performance Forecasting User's Guide*.

Create a Scenario

To create a scenario, complete the following steps:

1. In the hierarchy or table view, select the series that you want to use in the analysis.
2. Select **View** → **Details** → **Scenario Analysis View**.
3. In the Scenario Analysis View, click **New**. The Create New Scenario dialog box opens.
4. Specify a valid SAS name for the scenario. You can also provide a description.
5. Select the model that uses the inputs that you want to change in the analysis and click **OK**. The model now appears in the Scenario Analysis View.
6. In the input table, edit the future values for the independent variable. You can edit these values individually, or you can select multiple values and edit them using the Input Calculator.

Editing the Values of the Input Variables

About Editing Input Variables

You can edit the future values of the input variables in the following ways:

- by typing the new value in the input table. You can use this method when you want to specify a single value.
- by using the input calculator. This method is good if you want to create new values for several time periods.

Edit a Value from the Input Table

To edit the value of an input variable from the input table, complete the following steps:

1. In the scenario table, select a model.
2. In the input table, double-click in the cell for the time period that you want to override.
3. Type the new value and press ENTER.

Edit Values by Using the Input Calculator

To edit the values for multiple time periods, complete the following steps:

1. Click  to open the Input Calculator.
2. Select the time periods that you want to override.

   NOTE: Any time periods that you have selected in the input table in the Scenario Analysis View are automatically selected.
3. Specify how to calculate the input values for the selected time periods. You can base the input values on existing values or you can set the input values to a specified value.
4. Click **OK** to save your changes and close the input calculator.

Run a Scenario

To run a scenario, complete the following steps:

1. In the Scenario Analysis View, select a scenario from the table.
2. Click **Run Scenario**.

The values in the Scenario Forecast row of the input table are updated to reflect the changes to the input values.

Compare Scenarios

To compare scenarios, complete the following steps:

1. In the Scenario Analysis View, click **Compare**. The Compare Scenarios dialog box opens.
2. Select the check boxes of the scenarios that you want to compare. The selected plots appear in the **Scenario plots** area.
3. (Optional) Select the **Show legend** check box.
4. Click **Close** when you are finished.

Set Scenario Forecast Values as Overrides

If you want to apply the results from your scenario analysis to the selected series, you can add the scenario forecasts as overrides.

To set the scenario forecasts as overrides, complete the following steps:

1. In the hierarchy or table view, select the series that you want to use in the analysis.
2. Select **View** → **Details** → **Scenario Analysis View**.
3. In the Scenario Analysis View, click **Set scenario forecast values as overrides**.

In the Scenario Analysis View, the values of the model forecast and scenario forecast are identical. In the Forecasting View, the values of the scenario forecast appear in the Override row of the data table. By default, these overrides are locked.

Copy a Scenario

You might want to edit values in the scenario without losing your original scenario. You can do this by saving the scenario to a different name and then modifying the input values in this new scenario.

To copy a scenario, complete the following steps:

1. Select the scenario that you want to copy.
2. Click **Save As**. The Edit Scenario dialog box opens.

3. Specify a new name for the scenario. You can also modify the description and change the model selection.

4. Click **OK**.

The new scenario appears in the Scenario Analysis View. You can now modify the input values for this copied scenario and compare the results to the original scenario.

Reset Input Values in a Scenario

After modifying the input values for several time periods, you might decide that you want to use the last saved values. To reset the input values to the last saved values, click **Reset**. Any changes that you have made in the input table are lost.

Example: Performing a Scenario Analysis

Creating the Price_Increase Scenario

During the next several months, you expect the price of gasoline to increase in Region2. This increase will result in higher shipping costs for your company, and these higher costs need to be passed on to your customers. Currently, the value for the PRICE variable is 67.54489, and you want to raise the price to 70.12. You can create a scenario that shows how the increase in price changes the generated forecasts.

To perform this scenario analysis, complete the following steps:

1. In the hierarchy, select the Region2 node.

Figure 18.1 Region2 in the Hierarchy

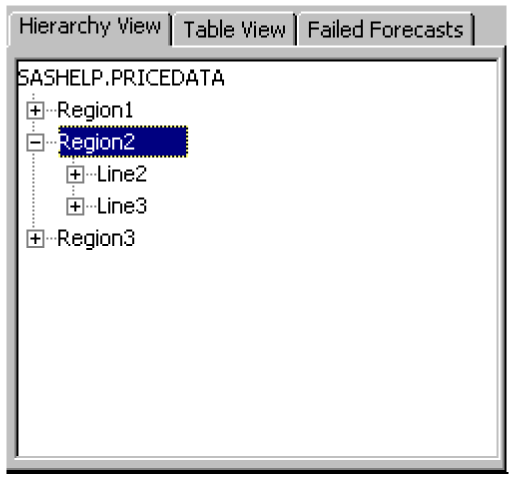

2. Select **View** → **Details** → **Scenario Analysis View** to open the Scenario Analysis View.
3. In the Scenario Analysis View, click **New**. The Create New Scenario dialog box opens.
4. In the **Name** box, type `Price_Increase`.
5. In the **Description** box, type `Reflects increase in gas prices`.
6. From the table, select the **Generated ARIMA Model (HPF0_7)**.

The Create New Scenario dialog appears similar to the following:

Figure 18.2 Select Model in Create New Scenario Dialog Box

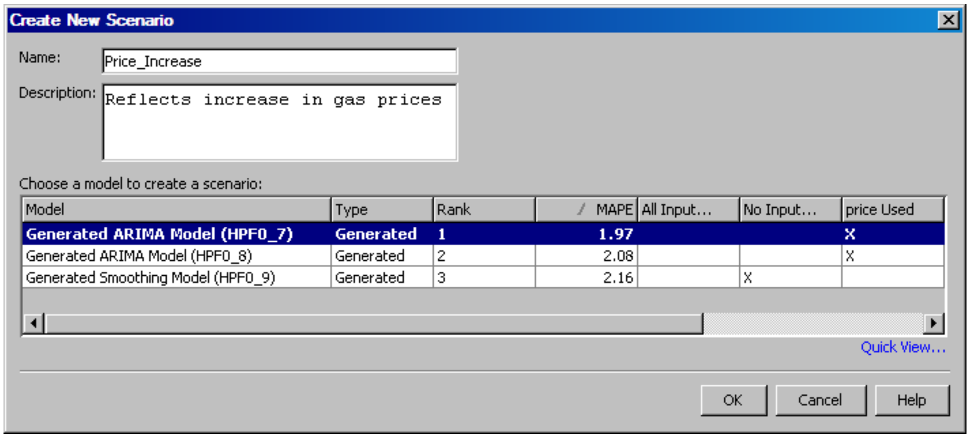

Click **OK**. The model now appears in the Scenario Analysis View.

Figure 18.3 New Price_Increase Scenario

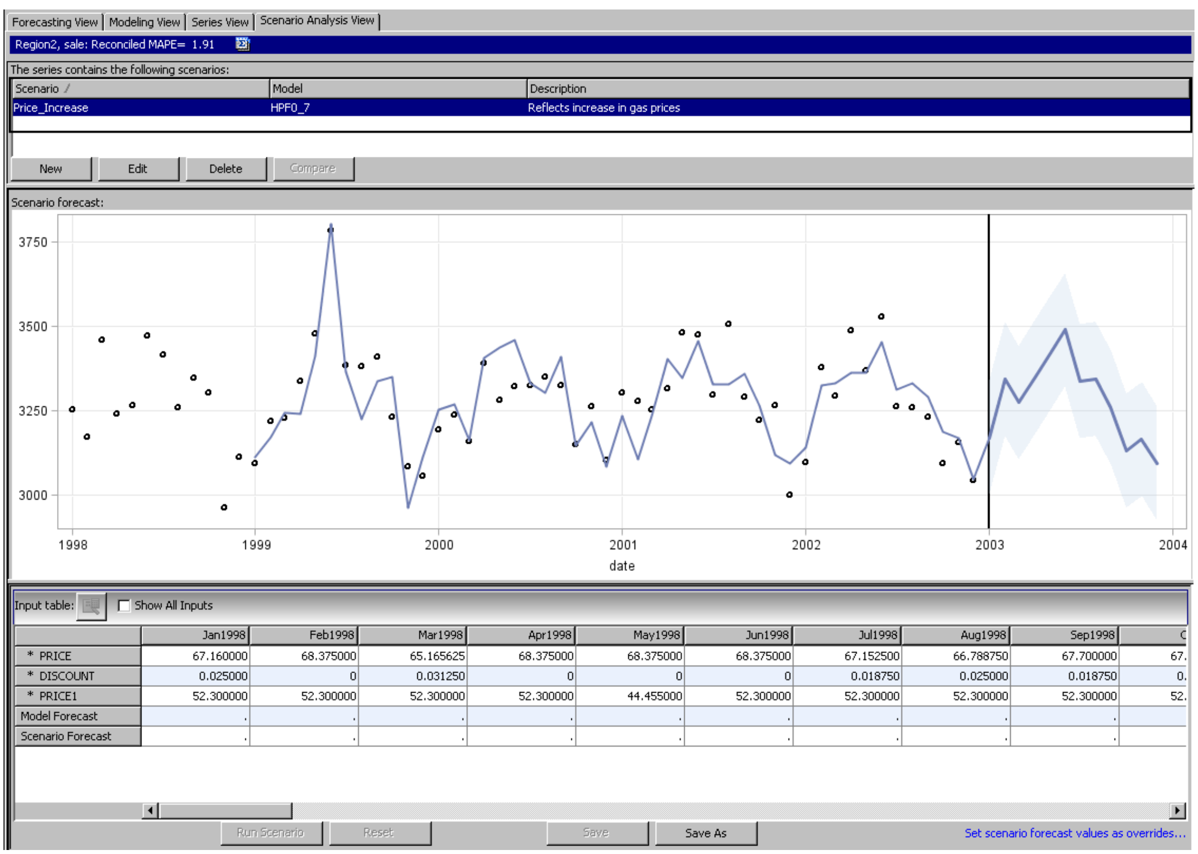

7. In the input table, edit the value of the PRICE variable by completing the following steps:

   a) Select all time periods in the PRICE variable row and click [icon]. The Input Calculator opens.

   b) In the Input Calculator, select **Set to a value** and type `70.12` in the text box.

   c) From the **Distribute value among periods** drop-down list, select **Assign specified value to each period**.

   The Input Calculator appears similar to the following:

   Figure 18.4 Assigning the Value of PRICE by Using the Input Calculator

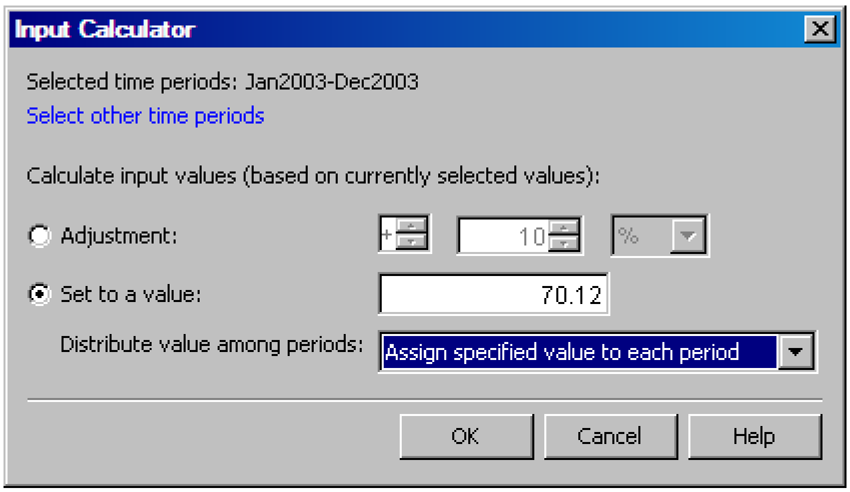

   Click **OK**. In the input table, 70.12000 now appears as the value of PRICE for each time period.

Figure 18.5 New PRICE Values in the Input Table

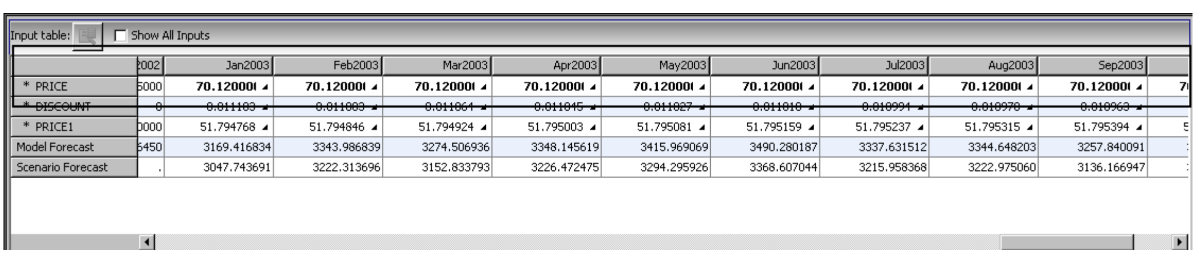

Input table: Show All Inputs

	2002	Jan2003	Feb2003	Mar2003	Apr2003	May2003	Jun2003	Jul2003	Aug2003	Sep2003
* PRICE	5000	**70.12000(**	**70.12000(**	**70.12000(**	**70.12000(**	**70.12000(**	**70.12000(**	**70.12000(**	**70.12000(**	**70.12000(**
* DISCOUNT	0	0.011103	0.011083	0.011064	0.011045	0.011027	0.011010	0.010994	0.010978	0.010963
* PRICE1	0000	51.794768	51.794846	51.794924	51.795003	51.795081	51.795159	51.795237	51.795315	51.795394
Model Forecast	6450	3169.416834	3343.986839	3274.506936	3348.145619	3415.969069	3490.280187	3337.631512	3344.648203	3257.840091
Scenario Forecast	.	3047.743691	3222.313696	3152.833793	3226.472475	3294.295926	3368.607044	3215.958368	3222.975060	3136.166947

8. Click **Run Scenario**.

When you compare the scenario forecast to the model forecast for each time period, you see that an increase in price slightly decreases the generated forecast. For example, in the Jan2003 time period, the model forecast is 3169.416834. However, the forecast that is generated by the scenario is 3047.743691. This is a decrease of 121.673143 units (or a decrease in sales of approximately 3.8%).

Figure 18.6 Difference in Model and Scenario Forecasts for Jan2003

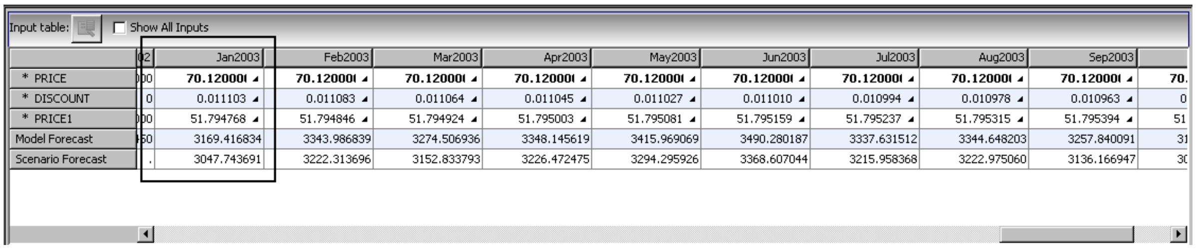

Input table: Show All Inputs

	02	Jan2003	Feb2003	Mar2003	Apr2003	May2003	Jun2003	Jul2003	Aug2003	Sep2003	
* PRICE	00	**70.12000(**	**70.12000(**	**70.12000(**	**70.12000(**	**70.12000(**	**70.12000(**	**70.12000(**	**70.12000(**	**70.12000(**	**70.**
* DISCOUNT	0	0.011103	0.011083	0.011064	0.011045	0.011027	0.011010	0.010994	0.010978	0.010963	0
* PRICE1	00	51.794768	51.794846	51.794924	51.795003	51.795081	51.795159	51.795237	51.795315	51.795394	51
Model Forecast	50	3169.416834	3343.986839	3274.506936	3348.145619	3415.969069	3490.280187	3337.631512	3344.648203	3257.840091	31
Scenario Forecast	.	3047.743691	3222.313696	3152.833793	3226.472475	3294.295926	3368.607044	3215.958368	3222.975060	3136.166947	30

Applying a Scenario to the Project

In addition to the Price_Increase scenario, you create a scenario called Discount_Decrease for the Region2 series. In this scenario, you see how the forecasts would be impacted if you eliminated the discount. You want to see how the scenarios compare. After the comparison, you decide that the Price_Increase scenario is the most likely outcome, and you apply the values for the scenario forecast so that they become the actual forecast values in the project.

To compare and select a scenario, complete the following steps:

1. To compare Price_Increase and Discount_Decrease scenarios, click **Compare**. The Compare Scenarios dialog box opens.

   Figure 18.7 Comparison Plots of the Scenarios

   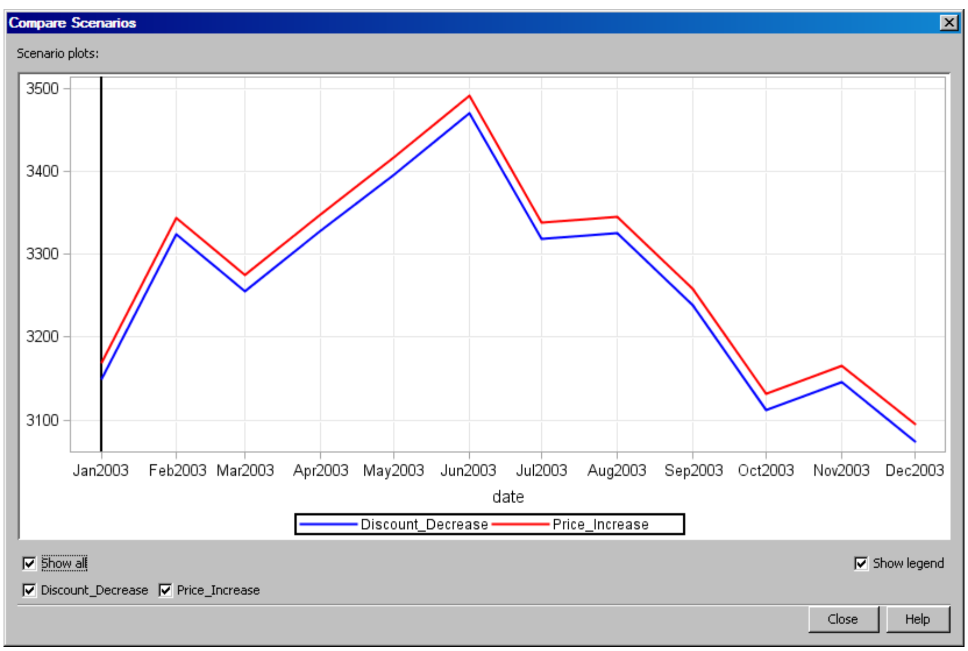

   These plots show that the forecasts for the Price_Increase scenario are slightly higher than the forecasts for the Discount_Decrease scenario. Although the Price_Increase scenario results in a slight decrease in sales, it is the better performing scenario, and you want to apply the values from this scenario to the Region2 series.

Click **Close** to close the Compare Scenarios dialog box.

2. In the Scenario Analysis View, select the **Price_Increase** scenario.
3. Below the input table, click **Set scenario forecast values as overrides**.

Now in the Forecasting View, the values for the scenario forecasts appear in the Override row of the data table.

Figure 18.8 Scenario Forecasts as Overrides in Forecasting View

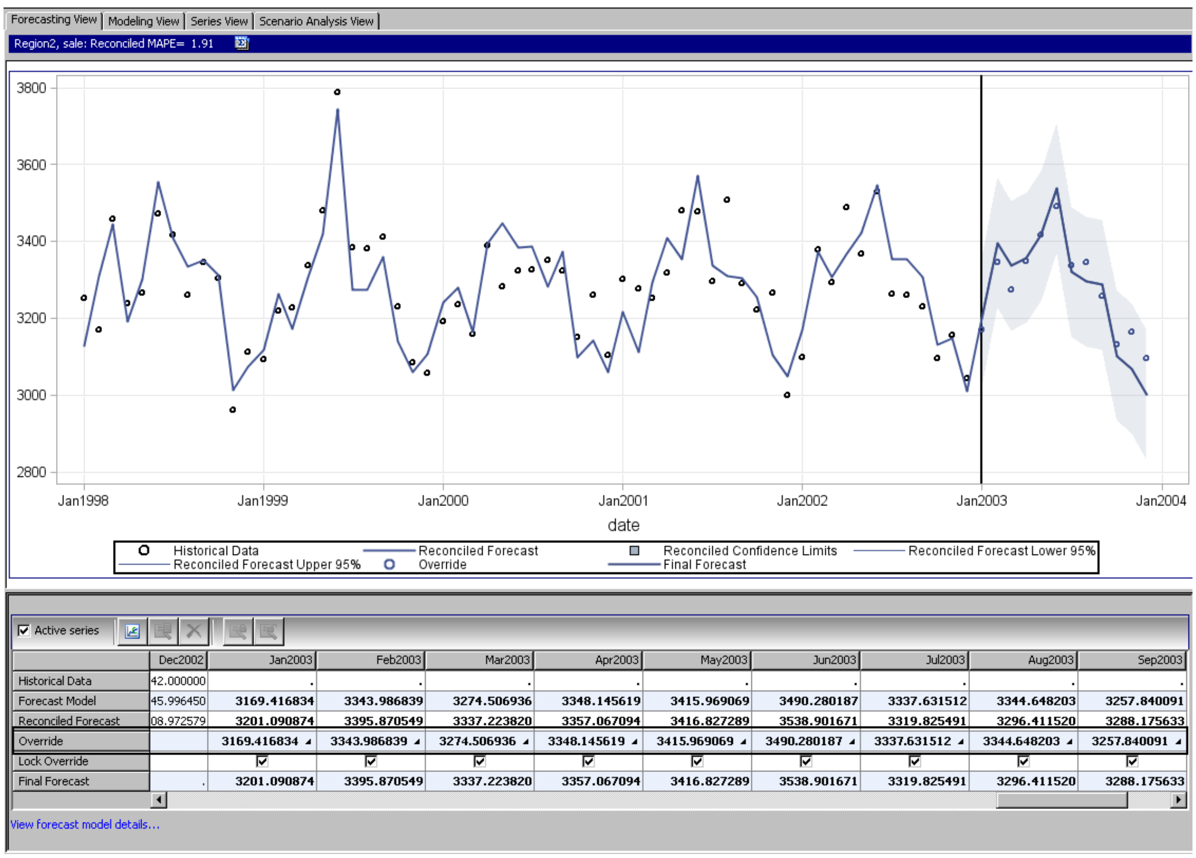

When you reconcile your project, SAS Forecast Studio uses the scenario forecasts as overrides for the Region2 series.

Part V

Advanced Topics

Chapter 19

Working in Batch Mode

Contents

Working with the SAS Code

About the SAS Code

In SAS Forecast Studio, you can view the SAS code for tasks, such as creating a model or event, or for forecasting the entire project. For more information about how to view the code for a project, see "View the Project Code" on page 232.

From the SAS Code dialog box, you view a read-only version of the SAS code that runs for that object or project. You can expand and collapse the code to view all or only portions of the code. You can also save the code to an external file. For more information, see "Save the Code to a File" on page 232.

View the Project Code

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

1. Select **Project** → **SAS Code**. The SAS Code dialog box opens.
2. From the **Task** drop-down list, select the version of the project code that you want to view. You can view the code for a specific task, such as the code for diagnosing the project or the code used in model selection.
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.

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 **Project** → **SAS Code**. The SAS Code dialog box opens.
2. From the **Task** drop-down list, select the version of the project code that you want to view. You can view the code for a specific task, such as the code for diagnosing the project or the code used in model selection.
3. In the SAS Code dialog box, click **Save**. The Save dialog box opens.
4. Select the location where you want to save this file and click **Save**.

Running Projects in Batch Mode

Saving the Project Code to a Batch File

Using SAS Forecast Studio, you can save the project code as a series of SAS files. You can use these files to create batch jobs for forecasting.

NOTE: In order for this code to execute properly in batch mode, you must declare the library references that the code uses to run. You can declare these library references in your `autoexec.sas` file. For help declaring library references, contact the SAS administrator at your site.

There are two versions of each SAS file. The files differ in whether they include the code for importing data.

- If the filename contains DO_NOT_IMPORT_DATA, then the code for importing new data is not included in this batch file. When this code is run, only the data that was available during the project creation is used in the calculations.
- If the filename contains IMPORT_DATA, then the code for importing new data is included in this batch file.

  This data specification code does the following:

  1. creates a copy of the original data that is sorted by the BY variables
  2. creates XML files that describe the hierarchy
  3. accumulates the data for each level

  When this code is run, any new data that is available is used in the calculation.

Descriptions of the Project Code in Each Batch File

The following list describes the project code that is contained in each batch file:

CREATE_PROJECT code files

The CREATE_PROJECT_DO_NOT_IMPORT_DATA and CREATE_PROJECT_IMPORT_DATA files contain the SAS code to re-create the project. This code includes all of the data preparation and forecasting steps.

NOTE: When you create a project in SAS Forecast Studio, you can choose whether to run the project immediately or whether to save the code to a file. If you choose to save the code, then you can run only the CREATE_PROJECT_IMPORT_DATA code files in batch mode.

These files include the following project code:

- Library declarations that create librefs, which reference existing directories inside a project. Library declarations include definitions of macro variables that describe the project location, the data specification location, and the event repository location.
- Catalog declarations for model repositories that are used by the project.
- Macro variable declarations that describe procedure settings.
- PROC HPFDIAGNOSE statements for each level in the project. These are "destructive," which means that SAS Forecast Studio rediagnoses all series in the project and re-creates the model repositories.

  PROC HPFENGINE statements for each level in the project are set to TASK=SELECT. The HPFENGINE procedure selects and fits the models for each series and creates the forecasts.
- If reconciliation is enabled, PROC HPFRECONCILE and PROC HPFENGINE statements for each pair of levels in a sequence, which is determined by the reconciliation level.

- DATA steps for adjusting statistic-of-fit values in the RECSTAT data set.
- DATA steps to generate data sets that contain the final forecasts at all levels.
- Calls to macros that update the final forecasts for any series that contains overrides. If reconciliation is enabled, then the HPFOVREC macro runs and reconciles the overrides. If reconciliation is not enabled, then the MERGOVRD macro runs.
- A statement that clears catalog, libref, and filename declarations.

DIAGNOSE_DESTRUCTIVE code files

The DIAGNOSE_DESTRUCTIVE_DO_NOT_IMPORT_DATA and DIAGNOSE_DESTRUCTIVE_IMPORT_DATA files contain the SAS code to rediagnose the project. This code does not include the data preparation steps.

These files contain the following project code:

- Library declarations that create librefs, which reference existing directories inside a project. Library declarations include definitions of macro variables that describe the project location, the data specification location, and the event repository location.
- Catalog declarations for model repositories that are used by the project.
- Macro variable declarations that describe procedure settings.
- PROC HPFDIAGNOSE statements for each level in the project. These are "destructive," which means that SAS Forecast Studio rediagnoses all series in the project and re-creates the model repositories.

  PROC HPFENGINE statements for each level in the project are set to TASK=SELECT. The HPFENGINE procedure selects and fits the models for each series and creates the forecasts.
- If reconciliation is enabled, PROC HPFRECONCILE and PROC HPFENGINE statements for each pair of levels in a sequence, which is determined by the reconciliation level.
- DATA steps for adjusting statistic-of-fit values in the RECSTAT data set.
- DATA steps to generate data sets that contain the final forecasts at all levels.
- Calls to macros that update the final forecasts for any series that contains overrides. If reconciliation is enabled, then the HPFOVREC macro runs and reconciles the overrides. If reconciliation is not enabled, then the MERGOVRD macro runs.
- A statement that clears catalog, libref, and filename declarations.

DIAGNOSE_NON_DESTRUCTIVE_PROJECT code files

The DIAGNOSE_NON_DESTRUCTIVE_PROJECT_DO_NOT_IMPORT_DATA and DIAGNOSE_NON_DESTRUCTIVE_PROJECT_IMPORT_DATA files contain the SAS code to rediagnose the project. This code does not include the data preparation steps.

These files include the following project code:

- Library declarations that create librefs, which reference existing directories inside a project. Library declarations include definitions of macro variables that describe the project location, the data specification location, and the event repository location.
- Catalog declarations for model repositories that are used by the project.
- Macro variable declarations that describe procedure settings.

- PROC HPFDIAGNOSE statements for each level in the project are not destructive. These statements use the INEST option of the HPFDIAGNOSE procedure to preserve all customizations that are created by user-defined models.

  PROC HPFENGINE statements for each level in the project are set to TASK=SELECT. The HPFENGINE procedure selects and fits the models for each series and creates the forecasts.
- If reconciliation is enabled, PROC HPFRECONCILE and PROC HPFENGINE statements for each pair of levels in a sequence, which is determined by the reconciliation level.
- DATA steps for adjusting statistic-of-fit values in the RECSTAT data set.
- DATA steps to generate data sets that contain the final forecasts at all levels.
- Calls to macros that update the final forecasts for any series that contains overrides. If reconciliation is enabled, then the HPFOVREC macro runs and reconciles the overrides. If reconciliation is not enabled, then the MERGOVRD macro runs.
- A statement that clears catalog, libref, and filename declarations.

SELECT_MODELS code files

The SELECT_MODELS_DO_NOT_IMPORT_DATA and
SELECT_MODELS_IMPORT_DATA code files contain the SAS code to perform model selection, to estimate the parameters of the selected model, and to produce forecasts. This code does not include the data preparation, diagnose, fit, and forecasting steps.

These files include the following project code:

- Library declarations that create librefs, which reference existing directories inside a project. Library declarations include definitions of macro variables that describe the project location, the data specification location, and the event repository location.
- Catalog declarations for model repositories that are used by the project.
- Macro variable declarations that describe procedure settings.
- PROC HPFENGINE statements (TASK=SELECT) for each level in the project. The HPFENGINE procedure selects and fits the models for each series and creates the forecasts.
- A statement that replaces missing values for stochastic independent variables.
- If reconciliation is enabled, PROC HPFRECONCILE and PROC HPFENGINE statements for each pair of levels in a sequence, which is determined by the reconciliation level.
- DATA steps for adjusting statistic-of-fit values in the RECSTAT data set.
- DATA steps to generate data sets that contain the final forecasts at all levels.
- Calls to macros that update the final forecasts for any series that contains overrides. If reconciliation is enabled, then the HPFOVREC macro runs and reconciles the overrides. If reconciliation is not enabled, then the MERGOVRD macro runs.
- A statement that clears catalog, libref, and filename declarations.

FIT_MODELS code files

The FIT_MODELS_DO_NOT_IMPORT_DATA and FIT_MODELS_IMPORT_DATA code files contain the SAS code to estimate parameters using the model that you specified and then to create a forecast. No model selection is performed. This code does not include the data preparation and diagnose steps.

These files include the following project code:

- Library declarations that create librefs, which reference existing directories inside a project. Library declarations include definitions of macro variables that describe the project location, the data specification location, and the event repository location.
- Catalog declarations for model repositories that are used by the project.
- Macro variable declarations that describe procedure settings.
- PROC HPFENGINE statements (TASK= FIT) for each level in the project. The HPFENGINE procedure refits selected models for each series and creates forecasts.
- A statement that replaces missing values for stochastic independent variables.
- If reconciliation is enabled, PROC HPFRECONCILE and PROC HPFENGINE statements for each pair of levels in a sequence, which is determined by the reconciliation level.
- DATA steps for adjusting statistic-of-fit values in the RECSTAT data set.
- DATA steps to generate data sets that contain the final forecasts at all levels.
- Calls to macros that update the final forecasts for any series that contains overrides. If reconciliation is enabled, then the HPFOVREC macro runs and reconciles the overrides. If reconciliation is not enabled, then the MERGOVRD macro runs.
- A statement that clears catalog, libref, and filename declarations.

FORECAST_MODELS code files

The FORECAST_MODELS_DO_NOT_IMPORT_DATA and FORECAST_MODELS_IMPORT_DATA contains the SAS code to reforecast the model and parameters. When refitting the model parameters, SAS Forecast Studio uses the estimate of the previous parameter as a starting point for reestimation. This code does not include the data preparation, diagnose, and fit model steps.

These files include the following project code:

- Library declarations that create librefs, which reference existing directories inside a project. Library declarations include definitions of macro variables that describe the project location, the data specification location, and the event repository location.
- Catalog declarations for model repositories that are used by the project.
- Macro variable declarations that describe procedure settings.
- PROC HPFENGINE statements (TASK= FORECAST) for each level in the project. The HPFENGINE procedure re-creates forecasts for each series by using selected and fitted models.
- A statement that replaces missing values for stochastic independent variables.
- If reconciliation is enabled, PROC HPFRECONCILE and PROC HPFENGINE statements for each pair of levels in a sequence, which is determined by the reconciliation level.

- DATA steps for adjusting statistic-of-fit values in the RECSTAT data set.
- DATA steps to generate data sets that contain the final forecasts at all levels.
- Calls to macros that update the final forecasts for any series that contains overrides. If reconciliation is enabled, then the HPFOVREC macro runs and reconciles the overrides. If reconciliation is not enabled, then the MERGOVRD macro runs.
- A statement that clears catalog, libref, and filename declarations.

RECONCILE_FORECASTS_AND_OVERRIDES_DO_NOT_IMPORT_DATA code file

The RECONCILE_FORECASTS_AND_OVERRIDES_DO_NOT_IMPORT_DATA file contains the SAS code to reconcile model forecasts and overrides. This code does not include the data preparation, diagnose, select, fit, and forecast steps. This code can be run only if those steps have been executed at least once.

These files include the following project code:

- Library declarations that create librefs, which reference existing directories inside a project. Library declarations include definitions of macro variables that describe the project location, the data specification location, and the event repository location.
- Catalog declarations for model repositories that are used by the project.
- Macro variable declarations that describe procedure settings.
- If reconciliation is enabled, PROC HPFRECONCILE and PROC HPFENGINE statements for each pair of levels in a sequence, which is determined by the reconciliation level.
- DATA steps for adjusting statistic-of-fit values in the RECSTAT data set.
- DATA steps to generate data sets that contain the final forecasts at all levels.
- Calls to macros that update the final forecasts for any series that contains overrides. If reconciliation is enabled, then the HPFOVREC macro runs and reconciles the overrides. If reconciliation is not enabled, then the MERGOVRD macro runs.
- A statement that clears catalog, libref, and filename declarations.

RECONCILE_FORECASTS_DO_NOT_IMPORT_DATA code file

The RECONCILE_FORECASTS_DO_NOT_IMPORT_DATA file contains the SAS code to reconcile model forecasts. This code does not include the data preparation, diagnose, select, fit, forecast, and override reconciliation steps. This code can be run only if those steps have been executed at least once.

These files include the following project code:

- Library declarations that create librefs, which reference existing directories inside a project. Library declarations include definitions of macro variables that describe the project location, the data specification location, and the event repository location.
- Catalog declarations for model repositories that are used by the project.
- Macro variable declarations that describe procedure settings.
- If reconciliation is enabled, PROC HPFRECONCILE and PROC HPFENGINE statements for each pair of levels in a sequence, which is determined by the reconciliation level.

- DATA steps for adjusting statistic-of-fit values in the RECSTAT data set.
- DATA steps to generate data sets that contain the final forecasts at all levels.
- A statement that clears catalog, libref, and filename declarations.

RECONCILE_OVERRIDES_DO_NOT_IMPORT_DATA code file

The RECONCILE_OVERRIDES_DO_NOT_IMPORT_DATA file contains the SAS code to reconcile overrides. This code does not include the data preparation, diagnose, select, fit, forecast, and reconciliation of the forecast steps. This code can be run only if those steps have been executed at least once.

These files include the following project code:

- Library declarations that create librefs, which reference existing directories inside a project. Library declarations include definitions of macro variables that describe the project location, the data specification location, and the event repository location.
- Catalog declarations for model repositories that are used by the project.
- Macro variable declarations that describe procedure settings.
- Calls to macros that update the final forecasts for any series that contains overrides. If reconciliation is enabled, then the HPFOVREC macro runs and reconciles the overrides. If reconciliation is not enabled, then the MERGOVRD macro runs.
- A statement that clears catalog, libref, and filename declarations.

About the SAS Forecast Server Macros

Summary of SAS Forecast Server Macros

The following SAS macros are provided with SAS Forecast Server software. A SAS macro is a program that generates SAS statements. Macros make it easy to produce and execute complex SAS programs that would be time-consuming to write yourself.

Table 19.1 shows the SAS macros that are included in the SAS Forecast Server software.

Table 19.1 SAS Forecast Server Macros

Macro Name	Description
%FSCLEAR	Clears the project information that is currently stored in the global macro variables.
%FSCOPY	Copies a SAS Forecast Studio project to a new destination. **NOTE:** Both servers must run the same version or higher of SAS Forecast Studio.
%FSCREATE	Creates a new SAS Forecast Studio project in batch mode.
%FSDELARC	Deletes an archived SAS Forecast Studio project.
%FSDELENV	Deletes the specified environment.

Table 19.1 SAS Forecast Server Macros (continued)

Macro Name	Description
%FSDELPRJ	Deletes an existing SAS Forecast Studio project.
%FSEXPALL	Exports all SAS Forecast Studio projects to archive files.
%FSEXPORT	Exports a single SAS Forecast Studio project to an archive file.
%FSGETENV	Creates an output data set that contains the following information: • the project name • the date and time that the project was created • the date and time that the project was modified • the user name of the person who created the project • whether the project is currently open by another user • whether the user can open this project • whether the user can delete this project • whether the project is public
%FSGETPRJ %FSGETURP	Retrieves the metadata about the SAS Forecast Studio projects. Prints to the log the names of the unregistered projects in the environment. You must use the %FSREGPRJ macro to register these projects.
%FSIMPALL	Imports all SAS Forecast Studio projects from archived files.
%FSIMPORT	Imports a SAS Forecast Studio project from an archived file.
%FSLOAD	Opens an existing SAS Forecast Studio project, and loads global macro variables that describe the project.
%FSMIGALL	Migrates all existing SAS Forecast Studio projects to the current version of SAS Forecast Server.
%FSMIGPRJ	Migrates an existing SAS Forecast Studio project to the current version of SAS Forecast Server.
%FSMOVE	Moves a SAS Forecast Studio project to a new destination. **NOTE:** Both servers must run the same version or higher of SAS Forecast Server.
%FSNEWENV	Creates a new environment.
%FSREGENV %FSREGPRJ	Creates a new environment and registers the projects. Registers a project in the metadata.
%FSREN	Renames a single SAS Forecast Studio project.
%FSRUNPRJ	Opens an existing SAS Forecast Studio project, and runs the project at a given stage.
%FSSETOWN	Assigns the owner of a SAS Forecast Studio project.
%FSSETPUB	Determines whether public access to a SAS Forecast Studio project should be enabled.

Table 19.1 SAS Forecast Server Macros (continued)

Macro Name	Description
%FSUNREG	Unregisters an existing SAS Forecast Studio project from the metadata server. **NOTE:** Only an administrative user can deploy this macro.
%FSURGENV	Unregisters the specified environment.

Additional Information

These macros are available automatically for you to use in your SAS programs.

For more information, see the following resources:

- *SAS Forecast Server Administrator's Guide* describes the syntax for each of the macros and provides examples of how to use the macros.
- *SAS Macro Language: Reference* describes the SAS macro facility.

Chapter 20

Working with SAS Forecast Studio Tasks

Contents

Overview of the SAS Forecast Studio Tasks

What Are the SAS Forecast Studio Tasks?

The following forecasting tasks are available from the SAS Add-In for Microsoft Office and SAS Enterprise Guide:

- the Forecast Studio Create Project task to create a SAS Forecast Studio project from an Excel or SAS data source
- the Forecast Studio Open Project task to open the results from a selected series in an existing SAS Forecast Studio project
- the Forecast Studio Submit Overrides task to submit overrides for the forecast data in an existing SAS Forecast Studio project

What Is the SAS Add-In for Microsoft Office?

The SAS Add-In for Microsoft Office extends the functionality of Microsoft Excel, Microsoft Word, and Microsoft PowerPoint by enabling you to access SAS analytics and SAS reporting functionality without any SAS programming experience. The SAS add-in is designed for users who are familiar with these Microsoft Office programs but who might be new to SAS.

When the SAS add-in is installed on your computer, you see the following items automatically integrated into Excel, Word, and PowerPoint:

- a SAS menu and the SAS Analysis Tools toolbar if you are using Microsoft Office 2003
- a **SAS** tab if you are using Microsoft Office 2007.

Figure 20.1 Location of SAS Menu in Microsoft Excel 2003

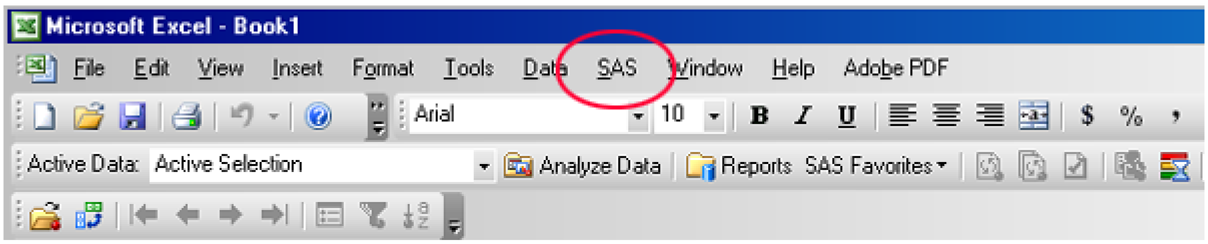

The SAS add-in includes approximately 80 SAS tasks that enable you to perform a variety of analyses. The Forecast Studio Create Project task, the Forecast Studio Open Project task, and the Forecast Studio Submit Overrides task are available in the Analyze Data window. In the SAS Tasks folder, click **Time Series** to view these tasks.

Figure 20.2 Forecasting Tasks in Analyze Data Window

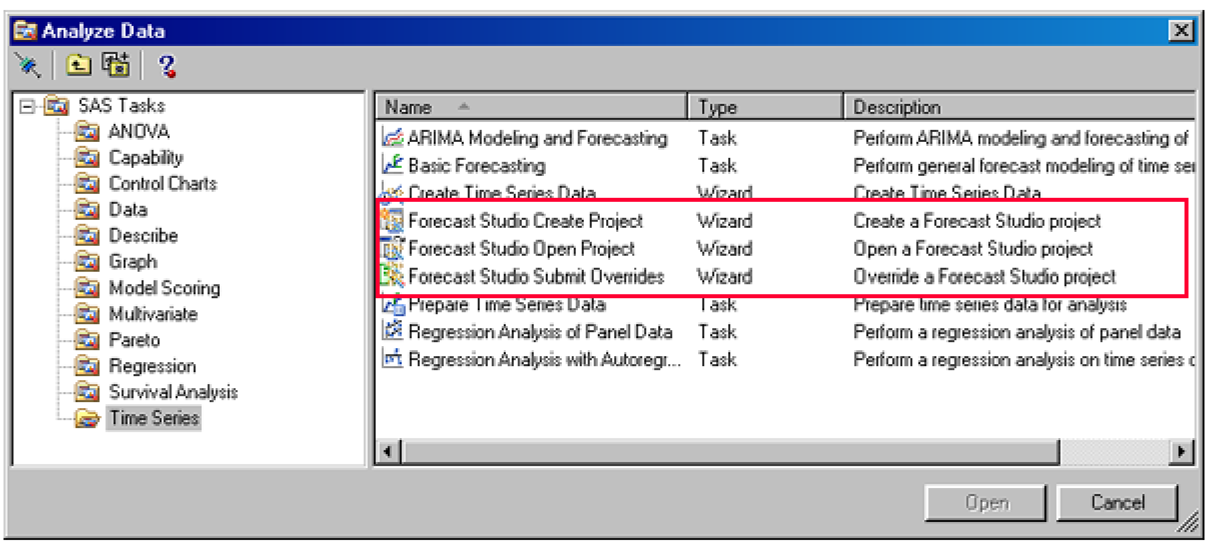

What Is SAS Enterprise Guide?

SAS Enterprise Guide provides a SAS graphical interface that helps you exploit the power of SAS and publish dynamic results in a Microsoft Windows client application. The solution is the preferred

SAS interface for business analysts, programmers and statisticians and a key application in SAS Business Intelligence offerings.

Figure 20.3 SAS Enterprise Guide

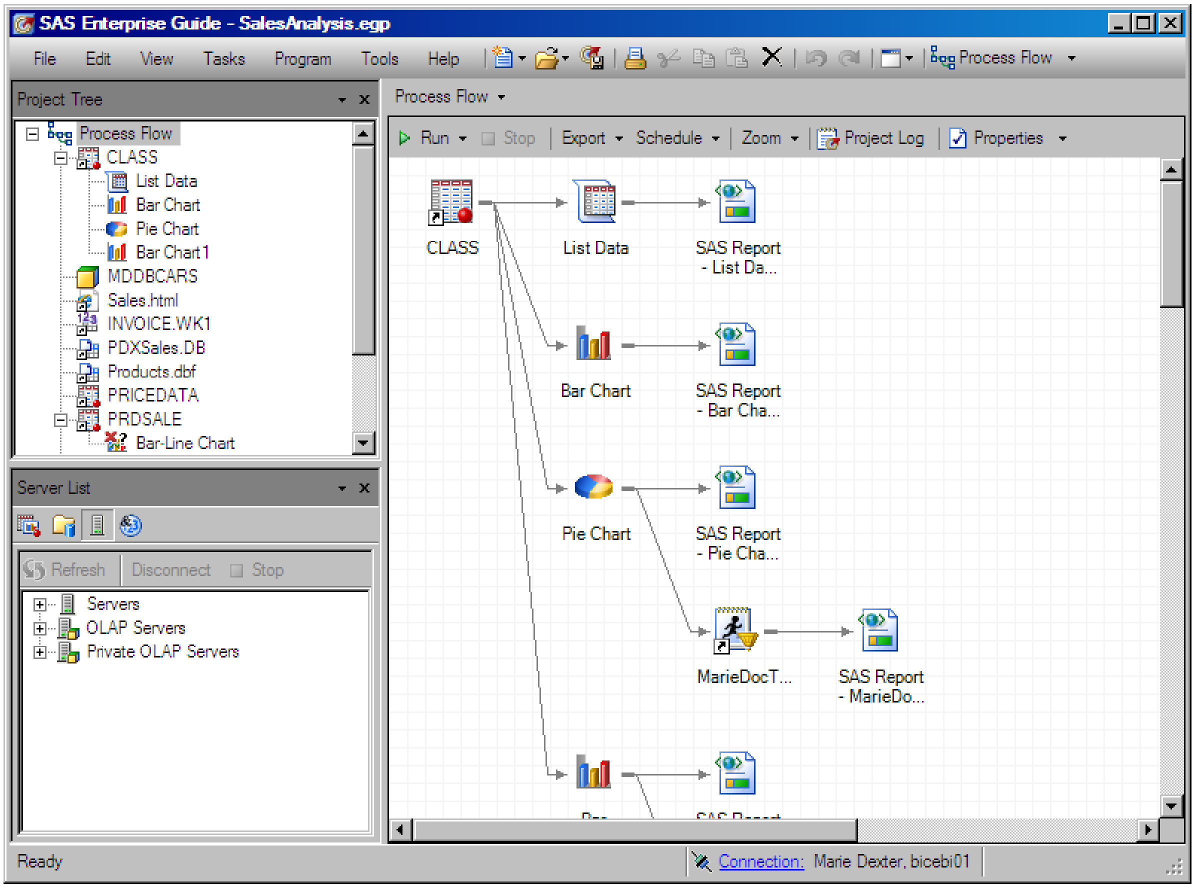

SAS Enterprise Guide includes approximately 80 SAS tasks that enable you to perform a variety of analyses. To access the forecasting tasks, select **Tasks** → **Time Series**.

Prerequisites for Using the SAS Forecast Studio Tasks

In order to use the SAS Forecast Studio tasks in SAS Enterprise Guide or the SAS Add-In for Microsoft Office, your site administrator must have completed the following steps:

- installed SAS Forecast Server 3.1.
- installed SAS Enterprise Guide 4.2 or the SAS Add-In 4.2 for Microsoft Office, which is a product in the SAS Enterprise Business Intelligence Server bundle.

  NOTE: The SAS Forecast Studio tasks were not available in the initial release of SAS Enterprise Guide 4.2 or the SAS Add-In 4.2 for Microsoft Office. You can download these from the

Software Downloads page at http://www.sas.com/download. These tasks will also be available in a maintenance release.

- configured SAS Forecast Studio Server to use the SAS Add-In for Microsoft Office. For more information, see the *SAS Forecast Server Administrator's Guide*.

The Forecast Studio Create Project Task

The Forecast Studio Create Project task enables you to create a SAS Forecast Studio project by using an Excel or SAS data source. Using this task, you can specify the forecasting variables, choose whether to forecast the data hierarchically, and specify the forecast horizon for the new SAS Forecast Studio project.

The Forecast Studio Open Project Task

The Forecast Studio Open Project task enables you to select a series from an existing SAS Forecast Studio project and specify the result types to open in Microsoft Excel, Word, or PowerPoint.

Here are some examples of the results that you can open in Microsoft Excel, Word, or PowerPoint:

- the forecasting plot
- the data table
- all of the models that have been defined for the current project
- the estimates, standard errors, and significance tests for each of the model parameters
- tables for each of the components that are used by the model
- all the available statistics of fit for the forecasting model
- the forecast values of the dependent variable for the time range of the forecast horizon

After opening these results, you can perform additional analyses using the other SAS tasks that ship with the SAS Add-In for Microsoft Office or SAS Enterprise Guide.

The Forecast Studio Submit Overrides Task

The Forecast Studio Submit Overrides task enables you to submit overrides for the forecast data in an existing SAS Forecast Studio project. After you submit an override, SAS Forecast Server reconciles the project. How the override is applied depends on whether the data in the project is forecasted hierarchically.

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

If SAS Forecast Server detects an override conflict when reconciling the project, then the SAS Enterprise Guide or the SAS add-in user sees an error message stating that the project cannot be reconciled and that the project must be opened in SAS Forecast Studio to reconcile the conflicts.

When the project is opened in SAS Forecast Studio, the "Reconciliation needed" warning appears at the top of the workspace. Click **Reconcile** to open the Override Conflicts dialog box to see all of the overrides that were submitted. To resolve the override conflicts, the SAS Forecast Studio user can modify the overrides by using the Override Conflicts dialog box, or the SAS Forecast Studio user can close this dialog box and update overrides by using the data table.

Additional Information

For more information about how to use each of these tasks, see the Help for the SAS Add-In for Microsoft Office or SAS Enterprise Guide.

Part VI

Appendixes

Appendix A

Troubleshooting Tips

Contents

Viewing 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.

To view the SAS log that SAS Forecast Studio creates when generating the forecast, select **Tools** → **SAS Log**. If an error occurs when SAS Forecast Studio is generating the forecasts, then you can click  in the status bar to open the SAS log.

Troubleshooting the Time ID Variable

In order to create the time series data, SAS Forecast Studio requires a specific format for the input data. For more information about how SAS Forecast Studio creates the time series data, see Chapter 5, "Preparing an Input Data Set for SAS Forecast Studio."

By default, SAS Forecast Studio reviews the first 10,000 observations in the time series data. If SAS Forecast Studio identifies three unique values for the time ID variable, then SAS Forecast Studio can automatically set the time interval. For more information about how to specify the number of observations to determine the time interval, see *SAS Forecast Server Administrator's Guide*.

Common errors include the following:

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

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

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

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

    NOTE: If the specified time interval is too small, then a large number of observations could have missing values. For example, you select WEEK as the interval, but there is no data for many of the weeks. A larger interval, such as MONTH, should be selected.

  – The values of the time ID variable are evenly spaced values at a specific time interval, but SAS Forecast Studio cannot detect the time interval. For example, SAS Forecast Studio cannot detect the time interval if the data represents a multi-weekday interval, such as the WEEKDAY3. interval.

    In the Specify Time ID Interval dialog box, select **Manually set the time interval and other time ID properties** option and specify the time interval to use.

  – The wrong variable has been selected as the time ID variable.

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

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

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

Troubleshooting Other Types of Variables

When you assign a variable (other than the time ID variable) to a role, you might get an error if the data values do not meet a specific requirement in SAS Forecast Studio. Some common errors include the following:

- 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.

- SAS Forecast Studio cannot load or find the format that is associated with a variable. When you define a custom format, you save these formats in a SAS catalog. In order for SAS Forecast Server to find the format catalog and recognize your custom format, you must specify the FMTSEARCH system option in your startup files.

  To use a custom format in SAS Forecast Studio, complete the following steps:

  1. Create the custom format and add it to a SAS catalog. For more information about custom formats and how to create them, see *SAS Language Reference: Dictionary*.
  2. In a text editor, open the SAS configuration file or the SAS autoexec file. These files can be found in the installation directory for SAS Forecast Studio.
  3. Add the FMTSEARCH system option to this startup file, so that SAS Forecast Studio knows where to find the format catalog.

     Use the following syntax to add the FMTSEARCH system option:

     FMTSEARCH=(*catalog-specification-1*...*catalog-specification-n*)

     The value for *catalog-specification* can be either a libref or libref.catalog. If only the libref is provided, then SAS assumes that FORMATS is the name of the catalog.
  4. Save the startup file.

- The following errors are specific to the BY variable. These errors are detected after you finish the New Project wizard. You will not see this error when you assign a variable to the BY variable role.

  – If your BY variable contains a format that maps multiple data values to the same formatted value, then SAS Forecast Server does not create the project. Mapping several data values to the same formatted value is not allowed.

  – Numeric variables that you assign to the BY variables role must have discrete values. If the values for the BY variable are continuous, then SAS Forecast Studio does not create the project.

Appendix B

Reserved Names in SAS Forecast Studio

Contents

Reserved Variable Names

Overview of Reserved Variable Names

For each project, SAS Forecast Studio creates several output data sets. The variable names in your input data set cannot match any of the variable names in these output data sets. The variable names in your input data set also cannot start with an underscore. If you try to assign a variable to a role and the variable name matches either of these conditions, then an error message appears.

For more information about the output data sets that SAS Forecast Studio creates, see Appendix C, "Output Data Sets in the Project Directory."

Reserved Variable Names

The following table lists alphabetically the variables that are used by SAS Forecast Studio.

Variable Name	**Description**
_*VariableName*	Any variable name that begins with an underscore
AADJRSQ	Amemiya's adjusted R-square
ACTUAL	Dependent series value
ADJRSQ	Adjusted R-square
AGGCHILDPREDICT	Aggregated prediction of child nodes
AIC	Akaike Information Criterion
AICC	Finite sample corrected Akaike Information Criterion

Variable Name	Description
APC	Amemiya's Prediction Criterion
DFE	Degrees of freedom error
END	Ending value of the time ID
ENDOBS	Number of the last observation in the data
ERROR	Prediction errors
FINALPREDICT	Predicted value for the parent node
GMAPE	Geometric mean percent error
GMAPPE	Geometric mean predictive percent error
GMAPES	Geometric mean absolute error percent of standard deviation
GMASPE	Geometric mean symmetric percent error
GMRAE	Geometric mean relative absolute error
ISRECONCILED	If the node is reconciled, then this variable is 1; if the node is not reconciled, then this variable is 0.
LEAF	Keyword used in model generation
LLOCK	Lock level for lower bound on the forecast
LOWBROVR	Lower confidence limits before override reconciliation
LOWER	Lower confidence limits
LOWERBD	Lower bound on the forecast
MAE	Mean absolute error
MAPE	Mean absolute percent error
MAPPE	Symmetric mean absolute predictive percent error
MASE	Mean absolute scaled error
MAX	Maximum value
MAXAPES	Maximum absolute error percent of standard deviation
MAXERR	Maximum error
MAXPE	Maximum percent error
MAXPPE	Maximum predictive percent error
MAXRE	Maximum relative error
MAXSPE	Maximum symmetric percent error
MDAPE	Median percent error
MDAPES	Median absolute error percent of standard deviation
MDAPPE	Median predictive percent error
MDASPE	Median symmetric percent error
MDRAE	Median relative absolute error
ME	Mean error
MEAN	Mean value
MIN	Minimum value
MINAPES	Minimum absolute error percent of standard deviation
MINERR	Minimum error
MINPE	Minimum percent error
MINPPE	Minimum predictive percent error
MINRE	Minimum relative error
MINSPE	Minimum symmetric percent error
MPE	Mean percent error
MPPE	Mean predictive percent error

Variable Name	Description
MRAE	Mean relative absolute error
MRE	Mean relative error
MSE	Mean square error
MSPE	Mean symmetric percent error
N	Number of nonmissing observations or number of variance products
NAME	Variable name
NMISS	Number of missing observations
NMISSA	Number of missing actuals
NMISSP	Number of missing predicted
NOBS	Number of observations
NONMISSCHLD	Number of nonmissing children in the current AGGBY group
NOTE	Text that user specifies for a series
NPARMS	Number of model parameters
OLOCK	Lock level for equality constraint
OVERRIDE	Equality constraint on the forecast
PREBFOVR	Predicted values before override reconciliation
PREDICT	Predicted values
RECDIFF	Reconciliation difference
RMSE	Root mean square error
RSQUARE	R-square
RWRSQ	Random walk R-square
SBC	Schwarz Bayesian information criterion
SMAPE	Symmetric mean absolute percent error
SSE	Sum of squares error
SST	Corrected total sum of squares
START	Beginning value of the time ID
STARTOBS	Number of the first observation
STD	Prediction standard errors
STDBFOVR	Standard deviation before override reconciliation
STDDEV	Standard deviation
SUM	Summation value
TOP	Keyword used in model generation
TSS	Total sum of squares
ULOCK	Lock level for the upper bound on the forecast
UMSE	Unbiased mean square error
UNLOCK	For locked overrides, the value of this variable is 0. For unlocked overrides, the value is 1.
UPPEFOVR	Upper confidence limits before override reconciliation
UPPER	Upper confidence limits
UPPERBD	Upper bound on the forecast
URMSE	Unbiased root mean square error
XML	User preferences for a series in XML format

Reserved Model Names

SAS Forecast Studio uses the following keywords in the names of the automatically generated models:

- LEAF
- TOP
- HP

These words are reserved for SAS Forecast Studio models. If the name of the model that you are creating contains these keywords, then a warning message appears, and you are prompted to select a different name.

NOTE: If you are copying a system-generated model that contains one of these keywords, then the copied name can contain one of these keywords. For example, if you copy a model named LEAF_*n*, then SAS Forecast Studio names the copied model LEAF_Copy*n*, where *n* is the lowest available integer value. Because this name is unique, SAS Forecast Studio allows this name.

Appendix C

Output Data Sets in the Project Directory

Contents

Overview of Output Data Sets

SAS Forecast Studio creates several output data sets. In general, these data sets contain the variables that you assign as BY variables. These output data sets are created for each level of the hierarchy and are saved in the following locations:

Windows operating environments
: **!ROOT\SAS\Forecast Server\3.1*environment*\Projects***project-name***hierarchy**

UNIX operating environments
: **!ROOT/SAS/ForecastServer/3.1/***environment***/Projects/***project-name***/hierarchy**

Descriptions of Output Data Sets

From the project directory, you can view the following output data sets. If the forecast fails for one series, then the forecasts in the horizon appear as .F in the data. If the series is at or above the reconciliation level, then all the parent nodes have missing reconciled forecasts, and these missing values appear as .F in the data. If the series is at or below the reconciliation level, then all the children nodes have missing reconciled forecasts, and these missing values appear as .F in the data.

FINALFOR= data set
: The FINALFOR= data set contains the observations that have a time ID value in the forecast horizon. If the RECFOR= data set is created for a level in the hierarchy, then the FINALFOR= data set is a subset of the RECFOR= data set; otherwise, the FINALFOR= data set is a subset of the OUTFOR= data set.

OUT= data set

The OUT= data set contains the variables that you assign as the time ID, BY variables, and dependent variables. The values for the time ID are based on the interval of the time series. The values of the dependent variables are accumulated, and missing values are interpreted using the data preparation options that you specified.

If any of the forecasting steps fail for a particular variable, the variable values are extended by missing values.

OUTCOMPONENT= data set

The OUTCOMPONENT= data set contains the variables that you assign as BY variables. This data set can also contain the variables listed below:

- _NAME_ - variable name or the name of the dependent variable
- _COMP_ - the name of the component
- _TIMEID_ - time ID values or the time ID variables
- _ACTUAL_ - the value of the dependent series
- _PREDICT_ - the forecast value of the component
- _LOWER_ - lower confidence limit
- _UPPER_ - upper confidence limit
- _STD_ - prediction standard error

OUTEST= data set

The OUTEST= data set contains the variables that you assign as BY variables. This data set can also contain the variables listed below if a forecasting model has been selected to forecast the actual time series. The following variables contain data that is related to the parameter estimation step:

- _NAME_ - variable name or the name of the dependent variable
- _SELECT_ - the name of the model selection list
- _MODEL_ - the name of the forecasting model
- _MODELVAR_ - the model variable that is used for mapping
- _DSVAR_ - the data set variable that is used for mapping
- _VARTYPE_ - the type of the variable to be forecast. In this case, the value is DEPENDENT.
- _TRANSFORM_ - the transformation for the model that was applied
- _COMPONENT_ - the model component. Examples of model components include autoregressive (AR), moving average (MA), and trend.
- _COMPMODEL_ - the model portion of an intermittent demand component
- _FACTOR_ - the model factor
- _LAG_ - the lag for the input
- _SHIFT_ - the shift
- _PARM_ - the parameter name
- _LABEL_ - the parameter label

- _EST_ - the parameter estimate
- _STDERR_ - standard error of the parameter estimate
- _TVALUE_ - the *t*-values of the parameter estimates
- _PVALUE_ - the *p*-values of the parameter estimates
- _STATUS_ - indicates the success or failure of estimating the parameter

OUTFOR= data set

The OUTFOR= data set contains the variables that you assign as BY variables. This data set can also contain the variables listed below:

- _NAME_ - variable name or the name of the dependent variable
- _TIMEID_ - time ID values or the time ID variables
- ACTUAL - values of the dependent series
- PREDICT - predicted values
- STD - prediction standard errors
- LOWER - lower confidence limits
- UPPER - upper confidence limits
- ERROR - prediction errors

If the forecasting step fails for a particular variable, no observations are recorded. If you specify to transform the time series, then the values in the variables listed above are the inverse transform forecasts. When you transform the time series, you can specify whether the median or mean forecasts are recorded.

OUTOVRD= data set

The OUTOVRD= data set contains any overrides that you specify in SAS Forecast Studio. This data set contains the variables that you assign as BY variables and the following variables:

- _NAME_ - variable name or the name of the dependent variable
- _TIMEID_ - time ID values or the time ID variables
- LOWER - lower bound on the forecast
- LLOCK - lock level for the lower bound
- 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.
- UPPER - upper bound on the forecast
- ULOCK - lock level for the upper bound

This data set is created only if you have specified overrides. If you did not specify any overrides, then SAS Forecast Studio does not create this data set.

OUTSTAT= data set

The OUTSTAT= data set contains the variables that you assign as BY variables. This data set can also contain the variables listed below. The following variables contain observations related to the statistics-of-fit step:

- _NAME_ - variable name or the name of the dependent variable
- _REGION_ - the region in which the statistics are calculated. Statistics that are calculated in the fit region are indicated by FIT. Statistics that are calculated in the forecast region are indicated by FORECAST.
- DFE - degrees of freedom error
- N - number of observations
- NOBS - number of observations that were used in calculating the statistics
- NMISSA - number of missing actuals
- NMISSP - number of missing predicted values
- NPARMS - number of parameters
- TSS - total sum of squares
- SST - corrected total sum of squares
- SSE - sum of squares error
- MSE - mean square error
- UMSE - unbiased mean square error
- RMSE - root mean square error
- URMSE - unbiased root mean square error
- MAPE - mean absolute percent error
- MAE - mean absolute error
- MASE - mean absolute scaled error
- RSQUARE - R-square
- ADJRSQ - adjusted R-square
- AADJRSQ - Amemiya's adjusted R-square
- RWRSQ - random walk R-square
- AIC - Akaike information criterion
- AICC - finite sample corrected AIC
- SBC - Schwarz Bayesian information criterion
- APC - Amemiya's prediction criterion
- MAXERR - maximum error
- MINERR - minimum error
- MINPE - minimum percent error
- MAXPE - maximum percent error
- ME - mean error
- MPE - mean percent error

- MDAPE - median absolute percent error
- GMAPE - geometric mean absolute percent error
- MINPPE - minimum predictive percent error
- MAXPPE - maximum predictive percent error
- MSPPE - mean predictive percent error
- MAPPE - symmetric mean absolute predictive percent error
- MDAPPE - median absolute predictive percent error
- GMAPPE - geometric mean absolute predictive percent error
- MINSPE - minimum symmetric percent error
- MAXSPE - maximum symmetric percent error
- MSPE - mean symmetric percent error
- SMAPE - symmetric mean absolute percent error
- MDASPE - mean absolute symmetric percent error
- GMASPE - geometric mean absolute symmetric percent error
- MINRE - minimum relative error
- MAXRE - maximum relative error
- MRE - mean relative error
- MRAE - mean relative absolute error
- MDRAE - median relative absolute error
- GMRAE - geometric mean relative absolute error
- MAPES - mean absolute error percent of standard deviation
- MDAPES - median absolute error percent of standard deviation
- GMAPES - geometric mean absolute error percent of standard deviation

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

OUTSTATSELECT= data set
The OUTSTATSELECT= data set contains the same variables as the OUTSTAT= data set with the addition of the following:

- _MODEL_ - the name of the forecasting model
- _SELECT_ - the name of the model selection list
- _SELECTED_ - whether the model was chosen to forecast the dependent series

OUTSUM= data set
The OUTSUM= data set contains the variables that you assign as BY variables. The OUTSUM= data set records the summary statistics for each dependent variable. If you select a model for the series, then the forecast values can include the prediction values, the upper confidence limits, or the lower confidence limits, depending on the options that you specify. If you do not select a model, then the forecast values are set to missing.

The following variables that are related to summary statistics can also appear in this data set. These variables depend on the options that you set for accumulation and for the interpretation of missing values.

- _NAME_ - variable name or the name of the dependent variable
- _STATUS_ - forecasting status. Nonzero values imply that no forecast was generated for the series.
- NOBS - number of observations
- N - number of nonmissing observations or number of observations
- NMISS - number of missing observations
- MIN - minimum value
- MAX - maximum value
- MEAN - mean value
- STDDEV - standard deviation

The following variables that are related to forecast summation are based on the values that you specify for the lead and for starting the summation:

- PREDICT - forecast summation of the predicted values
- STD - forecast summation of the prediction standard errors
- LOWER - forecast summation of the lower confidence limits
- UPPER - forecast summation of the upper confidence limits

Variance-related computations are computed only when no transformation is specified for the time series.

Your data set can also contain the _LEAD_ variable if you have a multistep forecast. A multistep forecast ranges from 1 to the lead number that you specify. This forecast can also contain the lower confidence limits, the upper confidence limits, or the predicted values.

If the forecast step fails for a particular variable, the variables related to forecasting are set to missing. The OUTSUM= data set contains a summary of the (accumulated) time series and might contain the forecasts for all series.

RECFOR= data set

The RECFOR= data set contains the values of the reconciled forecasts. This data set can also contain the variables listed below. The following variables contain data that is related to the forecasting step:

- _NAME_ - variable name or the name of the dependent variable
- _TIMEID_ - time ID values or the time ID variables
- PREDICT - predicted values
- STD - prediction standard errors
- LOWER - lower confidence limits
- UPPER - upper confidence limits
- ERROR - prediction errors
- _RECONSTATUS_ - the reconciliation status. This variable contains a code that specifies whether the reconciliation is successful or not. A corresponding message is also printed to the log.

  The _RECONSTATUS_ variable can have the following values:

0	Success.
500	A locked equality constraint has been imposed.
1000	ID value out of the range with respect to the START= and END= interval.
2000	Insufficient data to reconcile.
3000	Reconciliation failed for the predicted value. This implies that it also failed for the confidence limits and standard error.
4000	Reconciliation failed for the standard error.
5000	Reconciliation failed for the confidence limits.
6000	The constrained optimization problem is infeasible.
7000	The option DISAGGREGATION=PROPORTION has been changed to DISAGGREGATION=DIFFERENCE for this observation due to discordant signs in the input.
8000	The option STDMETHOD= provided by the user has been changed for this observation.
9000	The option CLMETHOD= provided by the user has been changed for this observation.
10000	The standard error hit the limits imposed by the STDDIFBD= option.
11000	Multiple warnings have been printed to the log for this observation.

NOTE: The RECFOR= data set is not created if you chose not to forecast your data hierarchically or when the reconciliation functionality is disabled. If you chose to forecast the data hierarchically and reconciliation functionality is enabled, then SAS Forecast Studio does not create the RECFOR= data set for the reconciliation level in the hierarchy. In this case, the reconciled forecasts are identical to the forecasts generated by the statistical model, so the content of the RECFOR= data set will be the same as the OUTFOR= data set.

RECSTAT= data set

The RECSTAT= data set contains the values of the reconciled forecasts. This data set also contains the same variables as the OUTSTAT= data set.

If the statistics-of-fit step fails for a particular variable, no observations are recorded. If you specified to transform the time series, the values in the variables listed above are computed based on the inverse transform forecasts. When you transform the time series, you can specify whether the median or mean forecasts are recorded.

NOTE: The RECSTAT= data set is not created if you chose not to forecast your data hierarchically or when the reconciliation functionality is disabled. If you chose to forecast the data hierarchically and reconciliation functionality is enabled, then SAS Forecast Studio does not create the RECSTAT= data set for the reconciliation level in the hierarchy. In this case, the reconciled forecasts are identical to the forecasts generated by the statistical model, so the content of the RECSTAT= data set will be the same as the OUTSTAT= data set.

RECSUM= data set

The RECSUM= data set contains the variables that you assigned to the BY variables role. The RECSUM= data set records the summary statistics for each dependent variable. If you selected a model for the series, then the reconciled forecasts can include the prediction values, the upper confidence limits, or the lower confidence limits, depending on the options that you specify. If you did not select a model, then the reconciled forecasts are set to missing.

The following variables that are related to summary statistics can also appear in this data set. These variables depend on the options that you set for accumulation and for the interpretation of missing values.

- _NAME_ - variable name or the name of the dependent variable
- _STATUS_ - forecasting status. Nonzero values imply that no forecast was generated for the series.
- NOBS - number of observations
- N - number of nonmissing observations or number of observations
- NMISS - number of missing observations
- MIN - minimum value
- MAX - maximum value
- MEAN - mean value
- STDDEV - standard deviation

The following variables that are related to forecast summation are based on the values that you specify for the lead and for starting the summation:

- PREDICT - forecast summation of the predicted values
- STD - forecast summation of the prediction standard errors
- LOWER - forecast summation of the lower confidence limits
- UPPER - forecast summation of the upper confidence limits

Variance-related computations are computed only when no transformation is specified for the time series.

Your data set can also contain the _LEAD_ variable if you have a multistep forecast. A multistep forecast ranges from one to the lead number that you specified. This reconciled forecast can also contain the lower confidence limits, the upper confidence limits, or the predicted values.

If the forecast step fails for a particular variable, the variables related to forecasting are set to missing. The RECSUM= data set contains a summary of the (accumulated) time series and might contain the 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.

SCENARIO data set

The SCENARIO data set contains information about any scenarios that you created in the Scenario Analysis View. This data set is saved in the directory for the series that is used in the scenario. If no scenarios are created for a series, then no data set is created. Modifications to the values of the input variables are stored with the scenario, so the values of the time ID variable are not continuous. For each series, the names of the scenarios must be unique. Although this data set is not indexed, a combination of the BY variables and the _NAME_ variable can be used as a unique key for the data set.

The SCENARIO data set contains the following variables:

- _NAME_ - name of the scenario
- _DESC_ - description of the scenario
- _MODEL_ - name of the model that is used in the scenario
- _STATUS_ - a system note for the scenario. This column is not currently used by SAS Forecast Studio.
- columns for each BY variable. The values in these columns identify the series that is associated with the scenario.
- a column for the time ID variable. This column stores the modifications to the input variable.
- a column for each input variable. This column stores the modification to the input variable.

Appendix D

Statistics of Fit

Contents

What Are Statistics of Fit?

The statistics of fit are statistical values that are used to evaluate how well a forecasting model performs by comparing the actual data to the predictions. For a given forecast model that has been fitted to the time series data, the model should be checked or evaluated to see how well it fits or forecasts the data. Commonly used statistics of fit are Root Mean Square Error (RMSE), Mean Absolute Percent Error (MAPE), Akaike Information Criteria (AIC), and many others. The statistics of fit can be computed from the model residuals or the prediction errors.

When the full range of data is used to both fit and evaluate the model, this is referred to as in-sample evaluation. When the most recent data is excluded for parameter estimation (holdout) and this holdout sample is used for evaluation, this is referred to as holdout sample evaluation. Holdout sample analysis is similar to the training and testing of neural networks. A portion of the data is withheld from training (fit) and the withheld data (holdout) is used to test performance.

When a particular statistic of fit is used for forecast model selection, it is referred to as the model selection criterion. For example, if the MAPE (an often recommended choice) is used as a model selection criterion, the forecast model with the smallest MAPE in the evaluation region (in-sample or holdout-sample) is chosen as the best model.

When a particular statistic of fit is used to judge how well the forecasting process is predicting the future, it is referred to as the performance statistic.

Descriptions of Statistics of Fit

You can use statistics of fit to measure how well different models fit the data. The statistics of fit for the various forecasting models can be printed or stored in a data set.

The definitions and formulas for the statistics of fit that are available in SAS Forecast Studio are described below. In these formulas, n is the number of nonmissing observations and k is the number of fitted parameters in the model. $APE = |100 * (y_t - \hat{y}_t)/y_t|$ is the absolute percent error. $ASPE = |100 * (y_t - \hat{y}_t)/0.5(y_t + \hat{y}_t)|$ is the absolute symmetric percent error. $APPE = |100 * (y_t - \hat{y}_t)/\hat{y}_t|$ is the absolute predictive percent error. $RAE = |(y_t - \hat{y}_t)/(y_t - y_{t-1})|$ is the relative absolute error. The errors are ignored in the statistical computations when the denominator is zero.

Adjusted R-square (ADJRSQ)
: The adjusted R^2 statistic, $1 - (\frac{n-1}{n-k})(1 - R^2)$.

Akaike Information Criterion (AIC)
: Akaike's information criterion, $n \ln(SSE/n) + 2k$.

Akaike Information Criterion, finite sample size corrected (AICC)
: Akaike's information criterion with an empirical correction for small sample sizes, $AIC + (\frac{2k(k+1)}{n-k-1})$.

Amemiya's Adjusted R-square (AADJRSQ)
: Amemiya's adjusted R^2, $1 - (\frac{n+k}{n-k})(1 - R^2)$.

Amemiya's Prediction Criterion (APC)
: Amemiya's prediction criterion, $\frac{1}{n}SST(\frac{n+k}{n-k})(1 - R^2) = (\frac{n+k}{n-k})\frac{1}{n}SSE$.

Geometric Mean Absolute Error Percent of Standard Deviation (GMAPES)
: The geometric mean of the absolute error as a percentage of the standard deviation.

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

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

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

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

Maximum Absolute Error Percent of Standard Deviation (MAXAPES)
: The maximum of the absolute error as a percentage of the standard deviation.

Maximum Error (MAXERR)
: The largest prediction error.

Maximum Percent Error (MAXPE)
: The largest percent prediction error, $100 \max((y_t - \hat{y}_t)/y_t)$. The summation ignores observations where $y_t = 0$.

Maximum Predictive Percent Error (MAXPPE)
: The maximum of the predictive percent errors.

Maximum Relative Error (MAXRE)
: The maximum of the relative errors.

Maximum Symmetric Percent Error (MAXSPE)
: The maximum of the symmetric percent errors.

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

Mean Absolute Error Percent of Standard Deviation (MAPES)
: The mean of the absolute error as a percentage of the standard deviation.

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

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

Mean Absolute Scaled Error (MASE)
: The mean of the absolute scaled errors.

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

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

Mean Percent Error (MPE)
: The mean percent prediction error, $\frac{100}{n}\sum_{t=1}^{n}\frac{(y_t-\hat{y}_t)}{y_t}$. The summation ignores observations where $y_t = 0$.

Mean Predictive Percent Error (MPPE)
: The mean of the predictive percent error.

Mean Relative Absolute Error (MRAE)
: The mean of the relative absolute errors.

Mean Relative Error (MRE)
: The mean of the relative errors.

Mean Square Error (MSE)
: The mean squared prediction error calculated from the one-step-ahead forecasts. $MSE = \frac{1}{n}SSE$. This formula enables you to evaluate small holdout samples.

Mean Symmetric Percent Error (MSPE)
: The mean of the symmetric percent errors.

Median Absolute Error Percent of Standard Deviation (MDAPES)
: The median of the absolute error as a percentage of the standard deviation.

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

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

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

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

Minimum Absolute Error Percent of Standard Deviation (MINAPES)
: The minimum of the absolute error as a percentage of the standard deviation.

Minimum Error (MINERR)
: The smallest prediction error.

Minimum Percent Error (MINPE)
: The smallest percent prediction error, $100\ \min((y_t - \hat{y}_t)/y_t)$. The summation ignores observations where $y_t = 0$.

Minimum Predictive Percent Error (MINPPE)
: The smallest predictive percent error.

Minimum Relative Error (MINRE)
: The smallest relative error.

Minimum Symmetric Percent Error (MINSPE)
: The smallest symmetric percent error.

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, might 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 $\mu = \frac{1}{n-1} \sum_{t=2}^{n} (y_t - y_{t-1})$.

Root Mean Square Error (RMSE)
: The root mean square error, $\sqrt{MSE}$.

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.

Total Corrected Sum of Squares for the dependent variable
: The total sum of squares for the series corrected for the mean: $\sum_{t=1}^{n} (y_t - \overline{y})^2$, where $\overline{y}$ is the series mean.

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

Unbiased Mean Square Error (UMSE)
: The unbiased mean square error.

Unbiased Root Mean Square Error (URMSE)
: The unbiased root mean square error.

Additional Information

For more information about statistics of fit, see the "Forecasting Process Details" chapter in the *SAS High-Performance Forecasting User's Guide*.

Descriptive Statistics

You can view the descriptive statistics for each series. These descriptive statistics are available from the Series Properties pane. You can also create a filter by using these descriptive statistics.

First observation
: The time period of the first observation in the series.

Last observation
: The time period of the last observation in the series.

Number of values
: The number of values in the series.

Number of missing values
: The number of series with missing values.

Number of nonmissing values
: The number of series with nonmissing values.

Sum of values
: The sum of the values in the series.

Mean of values
: The mean of the values in the series.

Standard deviation
: The standard deviation for the values in the series.

Minimum value
: The minimum value in the series.

Maximum value
: The maximum value in the series.

Appendix E

Default Model Selection Lists

Contents

Overview of the Default Model Selection Lists

Although you can create your own model, SAS Forecast Studio ships with several lists of candidate models that can be added to the project during the project setup. These model selection lists are saved in several SAS catalogs in the SASHELP library.

When you use a model, that model is automatically added to the model repository for the project. After a model is added to the repository, you can add a model to any series in the project. To view the complete list of models that are available for the current project, select **Project → Model Repository**.

Best Smoothing Model

The model selection list for the Best Smoothing Model is saved in the SASHELP.HPFDFLT.BEST catalog. The following models are in this selection list:

- Damped Trend Exponential Smoothing
- Double (Brown) Exponential Smoothing
- Linear (Holt) Exponential Smoothing

- Seasonal Exponential Smoothing
- Simple Exponential Smoothing
- Winters (Additive) Method
- Winters (Multiplicative) Method

Best Nonseasonal Smoothing Model

The model selection list for the Best Nonseasonal Model is saved in the SASHELP.HPFDFLT.BESTN catalog. The following models are in this selection list:

- Seasonal Exponential Smoothing
- Winters (Additive) Method
- Winters (Multiplicative) Method

Best Seasonal Smoothing Model

The model selection list for the Best Seasonal Smoothing Model is saved in the SASHELP.HPFDFLT.BESTS catalog. The following models are in this selection list:

- Damped Trend Exponential Smoothing
- Double (Brown) Exponential Smoothing
- Linear (Holt) Exponential Smoothing
- Simple Exponential Smoothing

Time Series Forecasting Models

The model selection list for the time series forecasting models is saved in the SASHELP.HPFDFLT.TSFSELECT catalog. The following models are in this selection list:

- Airline Model
- ARIMA(0,1,1)s NOINT

- ARIMA(0,1,1)(1,0,0)s NOINT
- ARIMA(0,1,2)(0,1,1)s NOINT
- ARIMA(0,2,2)(0,1,1)s NOINT
- ARIMA(2,0,0)(1,0,0)s
- ARIMA(2,1,0)(0,1,1)s NOINT
- ARIMA(2,1,2)(0,1,1)s NOINT
- Damped Trend Exponential Smoothing
- Double (Brown) Exponential Smoothing
- Linear (Holt) Exponential Smoothing
- Linear Trend
- Linear Trend with Autoregressive Errors
- Linear Trend with Seasonal Terms
- Log Airline Model
- Log ARIMA(0,1,1)s NOINT
- Log ARIMA(0,1,1)(1,0,0)s NOINT
- Log ARIMA(0,1,2)(0,1,1)s NOINT
- Log ARIMA(0,2,2)(0,1,1)s NOINT
- Log ARIMA(2,0,0)(1,0,0)s
- Log ARIMA(2,1,0)(0,1,1)s NOINT
- Log ARIMA(2,1,2)(0,1,1)s NOINT
- Log Damped Trend Exponential Smoothing
- Log Double (Brown) Exponential Smoothing
- Log Linear (Holt) Exponential Smoothing
- Log Linear Trend
- Log Linear Trend with Autoregressive Errors
- Log Linear Trend with Seasonal Terms
- Log Mean
- Log Random Walk with Drift
- Log Seasonal Dummy

- Log Seasonal Exponential Smoothing
- Log Simple Exponential Smoothing
- Log Winters Method – Additive
- Log Winters Method – Multiplicative
- Mean
- Random Walk with Drift
- Seasonal Dummy
- Seasonal Exponential Smoothing
- Simple Exponential Smoothing
- Winters (Additive) Method
- Winters (Multiplicative) Method

Appendix F

Autoregressive, Differencing, and Moving Average Orders

Contents

What Models Use Autoregressive, Differencing, and Moving Average Orders?

For ARIMA and Subset (factored) ARIMA models, you can specify the autoregressive, differencing, and moving average orders for the model. How you specify these orders differs slightly between the models. For the ARIMA model, you can specify integer values for the orders that make up the nonseasonal and seasonal components. For Subset (factored) ARIMA models, you have more flexibility. SAS Forecast Studio uses the syntax from PROC HPFARIMASPEC to specify this order. For more information about the underlying code, see the HPFARIMASPEC procedure in the *SAS High-Performance Forecasting User's Guide*.

For ARIMA, Subset (factored) ARIMA, and Unobserved Components models, you can also specify the differencing order for inputs, such as events.

Autoregressive Orders

The **Autoregressive (p)** option specifies the autoregressive part of the model.

To specify an autoregressive order, use the following syntax:

- For an autoregressive model with degrees from 1 to an integer value, use p = *integer*. This syntax creates a model of P = (1, . . . ,*integer*).

- For an autoregressive model with parameters at specified lags, use p = (*integer*). This syntax creates a model of P = (*integer*).
- For an autoregressive model with a seasonal component, use p = (*integer, . . . , integer*)(*integer, . . . , integer*)s, where *s* is a placeholder for the seasonal component. This placeholder value is supplied later. You can also specify an integer value for this seasonal component. This syntax creates a model of P = (*integer, . . . , integer*)(*integer, . . . , integer*)s.
- For an autoregressive model with multiple factors, use p = (*integer, . . . , integer*)...(*integer, . . . , integer*). This sytanx creates a model of P = (*integer, . . . , integer*)...(*integer, . . . , integer*).

The following table shows some examples of this syntax:

Model description	**Value of p**	**Result in Details pane**
Autoregressive model with degrees from 1 to an integer value	p = 5	P = (1 2 3 4 5)
Autoregressive model with parameters at specified lags	p = (5)	P = (5)
Autoregressive model with a seasonal component with "s" as a placeholder	p = (1 2 3)(2 7)s	P = (1 2 3)(2 7)s
Autoregressive model with a seasonal component of 12	p = (1 2 3)(2 7)12	P = (1 2 3)(24 84)
Autoregressive model with three factors	p = (1 2)(1)(1)s	P = (1 2)(1)(1)s

Differencing Orders

The **Differencing (d)** option specifies the differencing orders for the dependent series.

For differencing orders, you must enclose the value or values that you specify in parentheses. Unlike the autoregressive and moving average orders, you must include all orders that you specify in a single list, and you can specify an integer value multiple times in the same list. You can use placeholders (such as "s") to specify seasonal differencing.

To specify a differencing order, use the following syntax:

- For a single differencing order, use D = (*integer*). This syntax creates a model of D = (*integer*).
- For multiple differencing orders, use d = (*integer, . . . , integer*). This syntax creates a model of D = (*integer . . . integer*).
- For a differencing order with a seasonal component, use d = (*integer, . . . , s*). This syntax creates a model of D = (*integer, . . . , s*), where *s* is a placeholder for the seasonal order. This placeholder value is supplied later.

The following table includes examples of this syntax:

Model description	Value of d	Result in Details pane
Single differencing order	d = 5	Incorrect syntax; no order specified
Single differencing order	d = (5)	D = (5)
Multiple differencing orders	d = (1 1 3 5)	D = (1 1 3 5)
Differencing order with a seasonal component	d = (1 s s)	D = (1 s s)

Moving Average Orders

The **Moving average (q)** option specifies the moving average part of the model. The syntax of the moving average order is identical to the syntax of the autoregressive order.

To specify a moving average order, use the following syntax:

- For a moving average model with degrees from 1 to an integer value, use q = *integer*. This syntax creates a model of Q = (1, . . . ,*integer*).
- For a moving average model with parameters at specified lags, use q = (*integer*). This syntax creates a model of Q = (*integer*).
- For a moving average model with a seasonal component, use q = (*integer* . . . *integer*)(*integer* . . . *integer*)s, where *s* is a placeholder for the seasonal component. This placeholder value is supplied later. You can also specify an integer value for this seasonal component. This syntax creates a model of Q = (*integer,* . . . *, integer*)(*integer,* . . . *, integer*)s.
- For a moving average model with multiple factors, use q = (*integer,* . . . *, integer*)...(*integer,* . . . *, integer*). This syntax creates a model of Q = (*integer,* . . . *, integer*)...(*integer,* . . . *, integer*).

The following table shows some examples of this syntax:

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

Example: Specifying the Orders for the Airline Model

Using the syntax that is available in SAS Forecast Studio, you can create the familiar Airline model, ARIMA(0,1,1)(0,1,1)s, by specifying the differencing and moving average orders in the Subset (factored) ARIMA model. You can specify this model in any of the following ways:

Model description	Value of d	Value of q	Result
Season length is specified as a wildcard	d = (1 s)	q = (1)(1)s	D = (1 s) Q = (1)(1)s
Monthly seasonality is specified	d = (1 12)	q = (1)(1)12	D = (1 12) Q = (1)(12)
Alterative way to specify monthly seasonality	d = (1 12)	q = (1)(12)	D = (1 12) Q = (1)(12)

Appendix G

Sample Reports in SAS Forecast Studio

Contents

Overview of Reports

The sample reports that ship with SAS Forecast Studio are available from the Reports and Stored Processes dialog box. These reports are organized into the following categories:

- Getting Started
- Time Series Reports
- Statistical Model Reports
- Joint Reports
- Import and Export
- Quality Reports
- Reconciliation Reports
- Transactional Statistics Reports

- Override Reports
- Final Forecast Reports
- Reporting Variables Reports

Description of Sample Reports

Getting Started Reports

Overview

The Getting Started reports show the basic functionality that is available in SAS Forecast Studio. When you run these sample reports, you can see how parameters work and the types of output that you can generate.

Introduction

Name of Report	Description
Introduction to Reports	Creates a 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.
Introduction to Stored Processes	Demonstrates a simple example of a stored process. When you run this report, you can specify the title of the report.

Special Features

Name of Report	Description
Using Dynamic Parameters	Shows all of the dynamic parameters that you can use in SAS Forecast Studio. For more information about dynamic report parameters, see *SAS Forecast Server Administrator's Guide*.
Using Metadata Libraries	Shows how you can access registered metadata libraries. This sample includes a demonstration of the %HPF_LoadMetadataLibraries() macro.
Working with List Variables	Shows how parameters that accept multiple values are passed to the SAS code. This sample also shows how the SAS Forecast Server macros can be used to reformat these parameter values, so you can run code that was written for a previous release of SAS Forecast Server.

Time Series Reports

Exploration

The Exploration 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.

<table>
<tr><th>Name of Report</th><th>Description</th></tr>
<tr><td>Branch Series Plot</td><td>Creates a 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.</td></tr>
<tr><td>Level Descriptive Statistics Table</td><td>Creates a 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</td></tr>
<tr><td>Node Series Plot</td><td>Creates a report that contains a time series plot for the selected node.
When you run this report, you can specify the title of the report.</td></tr>
<tr><td>Path Cross Series Plot</td><td>Creates a report that contains a plot of the time series and independent variables 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. You can also change the title of the report.
NOTE: To run this report, you must have at least one independent variable in the project.</td></tr>
<tr><td>Node Cross Series Plot</td><td>Creates a 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. You can also change the title of the report.
NOTE: To run this report, you must have at least one independent variable in the project.</td></tr>
</table>

Name of Report	Description
Path Series Plot	Creates a 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.
Hierarchical Series Plot	Creates a 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.

Statistical Analysis

The Statistical Analysis 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 reports and lists the parameters that you can customize at run time.

<table>
<tr><th>Name of Report</th><th>Description</th></tr>
<tr><td>Level Seasonal Cycles Plot</td><td>Creates a 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</td></tr>
<tr><td>Node X12-ARIMA Analysis Report</td><td>Creates a 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</td></tr>
</table>

Name of Report	Description
Node Seasonal Cycles Plot	Creates a 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
Node Stationarity Analysis Report	Creates a report that contains the stationarity analysis report for the selected level in the hierarchy. When you run this report, you can customize the following parameters: • the title of the report • the number of dependent variable lags • the series transformation options
Node GARCH Autoregression Analysis Report	Creates a 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

<table>
<tr><th>Name of Report</th><th>Description</th></tr>
<tr><td>Node Autocorrelation Analysis Report</td><td>Creates an autocorrelation analysis report for the selected node in the hierarchy. The output 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</td></tr>
<tr><td>Node Cross Correlation Analysis Report</td><td>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</td></tr>
</table>

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 output 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 • 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.

<table>
<tr><th>Name of Report</th><th>Description</th></tr>
<tr><td>Node Decomposition Analysis Report</td><td>Creates a decomposition analysis report for the selected node in the hierarchy. The output 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</td></tr>
<tr><td>Node Stepwise Regression Analysis Report</td><td>Creates a 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</td></tr>
</table>

Name of Report	Description
Node GARCH Regression Analysis Report	Creates a 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
Node Stepwise Autoregression Analysis Report	Creates a 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

Statistical Model Reports

Evaluation

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
Level Model Statistics Univariate Plot	Creates a 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 Table	Creates a 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.
Node Model Statistics Table	Creates a 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.
Hierarchical Model Statistics Table	Creates a 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.
Level Model Statistic Volatility Plot	Creates a report that contains a plot of the volatility in the selected level.

Forecast

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 these forecast reports and lists the parameters that you can customize at run time.

Name of Report	Description
Branch Model Forecast Chart	Creates a report that contains a bar chart of the BY variable for the selected 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.
Node Model Forecast Plot	Creates a 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.
Path Model Forecast Plot	Creates a 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 Table	Creates a 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.
Branch Model Forecast Plot	Creates a 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.
Node Model Forecast Analysis Report	Creates a 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.
Node Model Forecast Percent Change Table	Creates a 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 Percent Change Plot	Creates a 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.

Parameter Analysis

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
Level Parameter Estimates Univariate Plot	Creates a 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.
Node Parameter Estimates Table	Creates a 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 a 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.

Component

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 these component reports and lists the parameters that you can customize at run time.

Name of Report	Description
Node Component Plot	Creates a 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.
Node Component Table	Creates a 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.

Summary

These reports contain the forecast values for the project. The following table gives a brief explanation of these summary reports and lists the parameters that you can customize at run time.

Name of Report	Description
Level Model Forecast Summary Table	Creates a 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.
Node Model Forecast Summary Table	Creates a 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.

Joint Reports

These reports contain information about the difference between the reconciled and unreconciled forecasts. The following table gives a brief explanation of the Joint reports and lists the parameters that you can customize at run time.

Name of Report	Description
Node Joint Statistics Table	Creates a 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.
Node Joint Forecast Plot	Creates a 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.

Import and Export Reports

These reports demonstrate how project data can be imported to or exported from an external data source. The following table gives a brief explanation of the Import and Export reports and lists the parameters that you can customize at run time.

Name of Report	Description
Export Node Macro Definitions	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 full path and name of this external file.
Export as OLAP Cube	Exports the forecasted data in the project to 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.
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.

Quality Reports

The following table gives a brief explanation of the quality reports and lists the parameters that you can customize at run time.

Name of Report	Description
Level Reconciled Forecast Quality Missing Table	Creates a 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.
Level Model Forecast Quality Extreme Table	Creates a 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 a 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 Reconciled Status Table	Creates a report that contains a table listing the status for the selected level in the hierarchy.
Level Model Forecast Quality Missing Table	Creates a 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.

Reconciliation Reports

Forecast

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 Table	Creates a 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 Analysis Report	Creates an analysis report for the reconciled forecast values for the selected node in the hierarchy. The output 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
Path Reconciled Forecast Plot	Creates a 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 Plot	Creates a 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.
Branch Reconciled Forecast Plot	Creates a 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.

Evaluation

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
Level Reconciled Statistic Volatility Plot	Creates a report that contains a scatter plot of the volatility of the reconciled forecasts for each node in the selected hierarchy. When you run this report, you can customize the title of the report. You can also select the statistic that you want to plot.
Level Reconciled Statistics Table	Creates a 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 a 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.
Hierarchy Reconciled Statistics Table	Creates a 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.
Node Reconciled Statistics Table	Creates a 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.

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 Seasonal Statistics Table	Creates a 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.

Name of Report	Description
Node Trend Statistics Table	Creates a 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.

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 a 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.

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 Table	Creates a 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 Plot	Creates a 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.

Reporting Variables Reports

These reports contain the statistics for the reporting variable. To run this report, you must have at least one reporting variable in the project.

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 a 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. You can also customize the title of the report.
Node Report Series Table	Creates a 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. You can also customize the title of the report and specify whether to include historical data in the output.

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.

ACF plot
: See autocorrelation function plot.

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 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. Short form: ACF plot.

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, and top-down method of reconciliation.

confidence limits
: the upper and lower values of a (usually 95%) confidence interval. In repeated sampling, approximately (1-alpha) 100% of the resulting intervals would contain the true value of the parameter that the interval estimates (where alpha is the confidence level associated with 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 and 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.

environment
: a virtual container of run-time settings for SAS Forecast Server client sessions.

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.

IACF plot
: See inverse autocorrelation function plot.

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

inverse autocorrelation function 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. Short form: IACF plot.

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, and 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.

PACF plot
See partial autocorrelation function plot.

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 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. Short form: PACF.

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.

report
output that is generated by running custom SAS code against the data in your project.

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 fits the historical series by comparing the actual data to the predicted values.

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

top-down method of reconciliation
: a reconciliation method that uses the forecasts at the highest level of the hierarchy to adjust the forecasts for the lower levels. See also bottom-up method of reconciliation, middle-out method of reconciliation, and reconciliation method.

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

unlocked override
: a user-supplied value for a forecast that acts as a guideline for the final forecast value. The final forecast for the level reflects the value of the unlocked override, but the final forecast and the unlocked override are often not identical. Because these overrides can be overridden when the hierarchy is reconciled, unlocked overrides do not generate override conflicts. See also locked override.

unreconciled node
: a node in the hierarchy that cannot be reconciled. A node can be unreconciled if (1) the final forecast of an upper level is not equal to the aggregate of the final forecasts of the lower levels; or (2) a final forecast violates the limits that are implied by a locked override in one of the lower levels in the same branch of the hierarchy.

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

Index

Your Turn

We welcome your feedback.

- If you have comments about this book, please send them to **`yourturn@sas.com`**. Include the full title and page numbers (if applicable).
- If you have comments about the software, please send them to **`suggest@sas.com`**.