



SAS[®] Forecast Studio 14.1: User's Guide

The correct bibliographic citation for this manual is as follows: SAS Institute Inc. 2015. *SAS® Forecast Studio 14.1: User's Guide*. Cary, NC: SAS Institute Inc.

SAS® Forecast Studio 14.1: User's Guide

Copyright © 2015, SAS Institute Inc., Cary, NC, USA

All Rights Reserved. Produced in the United States of America.

For a hard copy book: No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, or otherwise, without the prior written permission of the publisher, SAS Institute Inc.

For a web download or e-book: Your use of this publication shall be governed by the terms established by the vendor at the time you acquire this publication.

The scanning, uploading, and distribution of this book via the Internet or any other means without the permission of the publisher is illegal and punishable by law. Please purchase only authorized electronic editions and do not participate in or encourage electronic piracy of copyrighted materials. Your support of others' rights is appreciated.

U.S. Government License Rights; Restricted Rights: The Software and its documentation is commercial computer software developed at private expense and is provided with RESTRICTED RIGHTS to the United States Government. Use, duplication, or disclosure of the Software by the United States Government is subject to the license terms of this Agreement pursuant to, as applicable, FAR 12.212, DFAR 227.7202-1(a), DFAR 227.7202-3(a), and DFAR 227.7202-4, and, to the extent required under U.S. federal law, the minimum restricted rights as set out in FAR 52.227-19 (DEC 2007). If FAR 52.227-19 is applicable, this provision serves as notice under clause (c) thereof and no other notice is required to be affixed to the Software or documentation. The Government's rights in Software and documentation shall be only those set forth in this Agreement.

SAS Institute Inc., SAS Campus Drive, Cary, NC 27513-2414

January 2020

SAS® and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.

Other brand and product names are trademarks of their respective companies.

14.1-P2:fsug

Contents

<i>About This Book</i>	<i>ix</i>
<i>Accessibility Features of SAS Forecast Server</i>	<i>xi</i>

PART 1 Introduction to SAS Forecast Studio 1

Chapter 1 • About SAS Forecast Studio	3
What Is SAS Forecast Server?	3
Benefits of Using SAS Forecast Server	4
How SAS Forecast Server Works	4
How SAS Forecast Server Relates to Other SAS Software	5
What Is SAS Forecast Studio?	6
Understanding Role-Based Settings in SAS Forecast Studio	6
Chapter 2 • About SAS Forecast Studio for Desktop	7
What Is SAS Forecasting for Desktop?	7
Key Differences between SAS Forecasting for Desktop and SAS Forecast Server	8
Benefits of Using SAS Forecasting for Desktop	8
How SAS Forecasting for Desktop Works	8
How SAS Forecasting for Desktop Relates to Other SAS Software	9
What Is SAS Forecast Studio for Desktop?	10
Documentation for SAS Forecast Studio for Desktop	10
Chapter 3 • Starting SAS Forecast Studio	13
Understanding SAS Environments and Forecasting Environments	13
Starting SAS Forecast Studio	14
How You Can Run SAS Forecast Studio for Desktop	14
Chapter 4 • Using the Workspace	17
Overview of the Workspace	18
The Hierarchy View	19
The Table View	20
Failed Forecasts	21
The Forecasting View	22
The Modeling View	24
The Series View	26
The Scenario Analysis View	28
Customizing Your Workspace	29

PART 2 Getting Started in SAS Forecast Studio 31

Chapter 5 • Preparing an Input Data Set	33
Preprocessing Your Data	33
Understanding Time Series Data	34
Requirements for the Input Data Set	35
Working with Missing Values	36

Including Future Input Values in the Input Data Source	38
Chapter 6 • Understanding Variable Roles	39
Overview of the Types of Roles	39
Time ID Variables	40
Classification (BY) Variables	49
Dependent Variables	49
Independent Variables	50
Reporting Variables	50
Adjustment Variables	51
Chapter 7 • Forecasting Hierarchically	55
Understanding Hierarchies	55
Configure the Hierarchy	56
Understanding Reconciliation Methods	56
Understanding Aggregation and Accumulation	57
Example: How Forecasts Are Generated	60
Chapter 8 • Working with Projects	61
What Is a Project?	61
Creating a New Project	62
Migrating SAS Forecast Studio Projects	65
Open an Existing Project	66
Sharing a Project	67
Updating Your Project Data	67
Changing the Input Data Source	68
Export the Project Settings	68
Viewing Properties	68
Save a Copy of the Project	69
Create a Shortcut for a Project	70
 PART 3 Analyzing Your Results 71	
Chapter 9 • Generating Forecasts	73
Process for Generating a Forecast	73
Set the Diagnose Properties for Each Series	74
Specify the Number of Periods to Forecast	75
Create a Forecast Data Set for Independent Variables	76
Reviewing a Summary of the Forecast Run	76
Specifying an Inactive Series	77
Reforecasting a Series or Project	78
Chapter 10 • Working with Filters	79
What Is a Filter?	79
Create a Filter	80
View the Results from a Filter	80
Examples of Filters	81
Chapter 11 • Running Reports	91
Overview of Reports	91
Sample Reports in SAS Forecast Studio	92
Create a Report	92
Generating a Report	92
View Your Report History	95

Save the Output from a Report	96
-------------------------------------	----

PART 4 Improving Your Forecasts 97

Chapter 12 • Working with Events	99
Understanding Events	99
Creating New Events	101
Understanding Combination Events	105
Edit an Event	109
Copy an Event	109
Sharing Events	109
How Events Are Used in Model Generation	110
View the Effect of Events	115
Chapter 13 • Working with Overrides	117
Understanding Overrides	117
Adding Overrides	120
Lock and Unlock Overrides	125
Add a Note to a Series	125
Remove Overrides	126
Working with Override Conflicts	126
Chapter 14 • Understanding Hierarchy Reconciliation	129
Understanding Reconciliation	129
Reconcile the Hierarchy	131
Specify the Advanced Reconciliation Options	132
Working with Reconciliation Failures	132
Examples of Reconciling a Hierarchy with Overrides	136
Chapter 15 • Viewing the Results of a Cross-Series Analysis	163
View the Plots for the Time Series Analysis	163
View the Tables for the Time Series Analysis	163
Transform a Time Series for Analysis	164
Viewing the Results of a Time Series Decomposition	164
Chapter 16 • Working with Models	169
Overview of Model Repositories	170
How Models Are Fit	171
Select the Models to Fit to Each Series	173
Specifying the Forecast Model	174
Assessing the Accuracy of the Models	175
Compare Models	179
Add a Model to a Series	180
Importing and Exporting Models	180
Create a New Model Based on a Default Model	181
Edit a Model	182
Copy Models	183
Combining Models	184
Remove Models from the Model Selection List	187
Chapter 17 • Creating User-Defined Models	189
ARIMA Models	190
Subset (Factored) ARIMA Models	194
Unobserved Components Models	197

Exponential Smoothing Models	201
Intermittent Demand Models	203
Multiple Regression Models	205
Moving Average Models	207
Curve Fitting Models	208
Random Walk Models	210
Chapter 18 • Performing Scenario Analysis	213
What Is Scenario Analysis?	213
Create a Scenario	214
Editing the Values of the Input Variables	214
Run a Scenario	215
Compare Scenarios	215
Set Scenario Forecast Values as Overrides	215
Copy a Scenario	216
Reset Input Values in a Scenario	216
Example: Performing a Scenario Analysis	216
 PART 5 Advanced Topics 223	
Chapter 19 • Managing Projects, Environments, and Libraries	225
About the SAS Forecast Project Manager	225
Managing Projects	225
Managing Forecasting Environments	226
Managing Environment Libraries	226
How You Can Run the SAS Forecast Project Manager	226
Additional Documentation	227
 Chapter 20 • Working in Batch Mode	229
Working with the SAS Code	229
Running Projects in Batch Mode	230
 Chapter 21 • Working with SAS Forecast Studio Tasks	237
Overview of the SAS Forecast Studio Tasks	237
The Forecast Studio Create Project Task	239
The Forecast Studio Open Project Task	240
The Forecast Studio Override Project Task	240
Additional Information	241
 PART 6 Appendixes 243	
Appendix 1 • Troubleshooting Tips	245
Viewing the SAS Log	245
Troubleshooting the Time ID Variable	245
Troubleshooting Other Types of Variables	247
Performance	247
 Appendix 2 • Reserved Names for Variables and Models	249
Reserved Variable Names	249
Reserved Model Names	254

Appendix 3 • Output Data Sets in the Project Directory	255
Overview of Output Data Sets	255
FINALFOR= Data Set	256
OUT= Data Set	257
OUTCOMPONENT= Data Set	257
OUTINDEP= Data Set	258
OUTEST= Data Set	258
OUTFOR= Data Set	259
OUTOVRD= Data Set	260
OUTSTAT= Data Set	260
OUTSTATSELECT= Data Set	262
OUTSUM= Data Set	262
RECFOR= Data Set	263
RECSTAT= Data Set	265
RECSUM= Data Set	265
SCENARIO= Data Set	266
Appendix 4 • Statistics of Fit	267
What Are Statistics of Fit?	267
Understanding the Formulas and Definitions	268
Descriptions of Statistics of Fit	268
Additional Information	271
Descriptive Statistics	271
Appendix 5 • Default Model Selection Lists	273
Overview of the Default Model Selection Lists	273
Best Smoothing Model	273
Best Nonseasonal Smoothing Model	274
Best Seasonal Smoothing Model	274
Time Series Forecasting Models	274
Appendix 6 • Autoregressive, Differencing, and Moving Average Orders	277
Which Models Use Autoregressive, Differencing, and Moving Average Orders?	277
Autoregressive Orders	277
Differencing Orders	278
Moving Average Orders	279
Example: Specifying the Orders for the Airline Model	280
Appendix 7 • Sample Reports in SAS Forecast Studio	281
About the Sample Reports	281
Getting Started Reports	281
Time Series Reports	283
Statistical Model Reports	291
Joint Reports	296
Import and Export Reports	296
Quality Reports	297
Reconciliation Reports	299
Transactional Statistics Reports	301
Overrides Reports	302
Final Forecast Reports	302
Reporting Variables Reports	303
Recommended Reading	305
Glossary	307
Index	313

About This Book

Audience

This documentation is designed for users of SAS Forecast Studio (which is shipped with SAS Forecast Server) and users of SAS Forecast Studio for Desktop (which is shipped with SAS Forecasting for Desktop). Any functionality that is not available in SAS Forecast Studio for Desktop is noted in the SAS Forecast Studio Help and in this document.

Requirements

SAS Forecast Server 14.1 or SAS Forecasting for Desktop 14.1 installed by the SAS administrator at your site.

If you are running SAS Forecast Server 14.1, then the following items are also required:

- a valid SAS user account that is configured to access SAS Forecast Server.
- Access to data sources that contains historical data. This access must be provided using SAS libraries that are associated with SAS Workspace Servers.

If you are running SAS Forecasting for Desktop 14.1, you must have access to the data sources that contain your historical data. This access must be provided by SAS libraries that are associated with local SAS sessions.

Accessibility Features of SAS Forecast Server

Overview

SAS Forecast Server has been tested with assistive technology tools. It includes accessibility and compatibility features that improve the usability of the product for users with disabilities. (Some accessibility issues remain and are 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 (2008 draft proposal initiative update). For detailed information about the accessibility of this product, send e-mail to accessibility@sas.com or call SAS Technical Support.

Documentation Format

Please contact accessibility@sas.com if you need this document in an alternative digital format.

Exceptions to Accessibility Standards

Exceptions to accessibility standards are documented in the following table.

Exceptions to Accessibility Standards

Accessibility Issue	Workaround
Tables and items within tables are not correctly announced by screen readers (such as JAWS).	None
In the Hierarchy and Variables Settings dialog box, the drop-down list in the Accumulation column cannot be opened or altered by pressing CTRL+DOWN ARROW.	Press F2+DOWN ARROW to change the selection in the drop-down list.

Accessibility Issue	Workaround
In some areas, two or more icons might appear to have focus at the same time. For example, you give focus to the Forecast Series icon () in the Forecasting View and then press TAB. Now, the focus is on the Copy Table icon (). If you press SHIFT+TAB (to go in reverse order), the focus appears to be on both icons.	None
Some drop-down lists do not have a visual indication of focus. For example, if you bring focus to the Interval drop-down list in the Forecasting Settings dialog box, the focus is not visible.	Press ALT+DOWN ARROW to verify that focus is on the drop-down list.
In the Modeling View and Series View, focus cannot be placed on the Tile icon (). When the icon receives focus, the focus is immediately removed.	When the Modeling View or Series View is selected, select Series ⇒ Tile to display your graphs and tables in a tiled view.
Focus is lost after selecting a value from the Environment drop-down list in the Project dialog box.	Use a mnemonic key combination to return focus to the items in the Project dialog box.
If you use the arrow keys to move through items in the Environment drop-down list in the Projects dialog box, the item that was highlighted using the arrow keys is automatically selected and displayed.	None
The high contrast settings for some operating environments are not fully accepted. Some tables (such as the Active Series table) might not accept the settings, and the focus remains black in some places.	None

Part 1

Introduction to SAS Forecast Studio

<i>Chapter 1</i>	
About SAS Forecast Studio	3
<i>Chapter 2</i>	
About SAS Forecast Studio for Desktop	7
<i>Chapter 3</i>	
Starting SAS Forecast Studio	13
<i>Chapter 4</i>	
Using the Workspace	17

Chapter 1

About SAS Forecast Studio

What Is SAS Forecast Server?	3
Benefits of Using SAS Forecast Server	4
How SAS Forecast Server Works	4
How SAS Forecast Server Relates to Other SAS Software	5
What Is SAS Forecast Studio?	6
Understanding Role-Based Settings in SAS Forecast Studio	6

What Is SAS Forecast Server?

SAS Forecast Server provides automation and analytical sophistication to the forecasting process. By using SAS Forecast Server, you can generate millions of forecasts in the turnaround time that is necessary to run your business. You can also uncover previously undetected trends, and you can predict future seasonal fluctuations. These capabilities create ample opportunities for you to reduce costs and increase revenues.

SAS Forecast Server has these components:

- SAS Forecast Server middle tier, which is supported by the SAS Web Infrastructure Platform.
- SAS Forecast Studio, which is the graphical interface (based on Java) to the forecasting and time series analysis procedures contained in SAS.
- SAS Forecast Server plug-in for SAS Management Console, which enables a SAS Forecast Server administrator to manage projects and environments.
- SAS Forecast Project Manager, which is an administrative component that enables you to manage projects, environments, and libraries.
- a batch interface that enables you to use macros to manage projects and environments. For more information about these macros, see the *SAS Forecast Server: Administrator's Guide*.
- SAS Forecast Studio tasks that are available in SAS Enterprise Guide and the SAS Add-In for Microsoft Office. For more information, see [Chapter 21, “Working with SAS Forecast Studio Tasks,”](#) on page 237.

Benefits of Using SAS Forecast Server

SAS Forecast Server 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 Server recognizes business drivers, holidays, and 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 Server Works

Given a time-stamped data set, the software provides the following automatic forecasting process:

1. accumulates the time-stamped data to form a fixed-interval time series
2. aggregates the time series to form a hierarchical time series.

For each time series in the hierarchy, SAS Forecast Server performs these steps:

1. diagnoses the time series using time series analysis techniques.
2. creates a list of candidate model specifications based on the diagnostics.
3. fits each candidate model specification to the time series.
4. generates forecasts for each candidate fitted model.
5. selects the most appropriate model specification based on either in-sample or holdout sample evaluation using a model selection criterion.
6. refits the selected model specification to the entire range of the time series.
7. creates a forecast score from the selected fitted model.
8. generates forecasts from the forecast score.
9. evaluates the forecast using in-sample analysis and provides for out-of-sample analysis of forecast performance.

For each forecast generated by the time series model, SAS Forecast Server performs these steps:

1. reconciles the model forecasts to form consistent reconciled forecasts.
2. evaluates the forecast using in-sample analysis and provides for out-of-sample analysis of forecast performance.

For each reconciled forecast, SAS Forecast Server performs these steps:

1. applies manual overrides to the reconciled forecasts.
2. reconciles the overrides to form consistent override forecasts.

How SAS Forecast Server Relates to Other SAS Software

SAS Forecast Server uses many of the procedures and options from other SAS products (such as SAS Forecast Server procedures, SAS/ETS, Base SAS, and SAS/GRAPH). Many features not found in this solution software are available in other SAS solutions or in SAS products that are used with this SAS solution. If you do not find a feature that you need in this software, you might find it in one of the following SAS solutions or products.

SAS Forecast Server Procedures

SAS Forecast Server uses several procedures as the basis for the automatic forecasting capabilities. The SAS Forecast Server procedures automatically select 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, see the *SAS Forecast Server Procedures: User's Guide*.

SAS/ETS

SAS/ETS software provides SAS procedures that perform econometric and time series analysis and forecasting, as well as financial analysis and reporting. The software also provides an interactive environment for time series forecast and investment analysis. For more information, see the *SAS/ETS User's Guide*.

Base SAS

Base SAS delivers a highly flexible and extensible fourth-generation programming language that is specially designed for data access, transformation, and reporting. It includes a rich library of encapsulated programming procedures for data manipulation, information storage and retrieval, descriptive statistics, and report writing. The output for the SAS Forecast Server is generated by the Output Delivery System that is part of Base SAS.

SAS/GRAPH

SAS/GRAPH software provides high-impact visuals for all levels of your organization, enabling customers to readily understand complex information and empowering them to make informed, timely decisions. SAS/GRAPH software extends the power of SAS data management, business intelligence, and analytic tools, enabling customers to turn data into full-color graphs and charts.

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. SAS Forecast Studio is shipped as part of SAS Forecast Server, which is an enterprise solution.

Using this application, you can do the following tasks:

- Generate forecasts automatically by using models that are shipped with SAS Forecast Studio. You can generate forecasts by using a model selection list that you have selected.
- Create your own forecasting models or combine generated 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.
- Add events to a model or to the entire project.
- 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.
- Perform rolling simulations to determine the best model for your data.
- Generate reports to share forecasting results with other people at your site.

Understanding Role-Based Settings in SAS Forecast Studio

Because SAS Forecast Studio is used by a variety of people and groups within an organization, your administrator might choose to restrict the menu items and functionality that each user can access. By default, SAS Forecast Server is shipped with several default roles. These roles manage the availability of product features, called *capabilities*.

By using the User Manager plug-in in SAS Management Console, your administrator for SAS Forecast Server can modify these roles and specify the capabilities that are available to users at your site. For example, your administrator might want to restrict who can modify forecasts or who can specify overrides. As a result, all of the capabilities that are documented might not be available to you. If you need access to an application feature, contact your site administrator.

For more information, see the *SAS Forecast Server: Administrator's Guide*.

Chapter 2

About SAS Forecast Studio for Desktop

What Is SAS Forecasting for Desktop?	7
Key Differences between SAS Forecasting for Desktop and SAS Forecast Server	8
Benefits of Using SAS Forecasting for Desktop	8
How SAS Forecasting for Desktop Works	8
How SAS Forecasting for Desktop Relates to Other SAS Software	9
What Is SAS Forecast Studio for Desktop?	10
Documentation for SAS Forecast Studio for Desktop	10

What Is SAS Forecasting for Desktop?

SAS Forecasting for Desktop is a desktop product that provides automation and productivity gains for small-scale forecasting tasks. (SAS Forecast Server is recommended for large-scale forecasting projects.) No servers are required when you use SAS Forecasting for Desktop, so all work is performed and saved locally.

SAS Forecasting for Desktop has these components:

- SAS Forecast Studio for Desktop, which is the graphical interface (based on Java) to the forecasting and time series analysis procedures in SAS.
- SAS Forecast Project Manager, which is an administrative component that enables you to manage projects, environments, and libraries. For more information, see [Chapter 19, “Managing Projects, Environments, and Libraries,”](#) on page 225.
- a batch interface that enables you to use macros to manage projects and environments. For more information about these macros, see the *SAS Forecast Server: Administrator’s Guide*.

Key Differences between SAS Forecasting for Desktop and SAS Forecast Server

SAS Forecasting for Desktop is designed for smaller data and small-scale forecasting tasks. SAS Forecast Server is recommended if you have large amounts of data and large-scale forecasting projects.

Due to a limited license, SAS Forecasting for Desktop differs from SAS Forecast Server in these ways:

- Your input data sets cannot exceed 1,000 time series.
- Because there is no middle tier, you cannot launch SAS Forecast Studio for Desktop from a Java Web Start client.
- You cannot run the code that is generated by SAS Forecast Studio for Desktop in batch mode.
- Custom reports and stored processes are not available in SAS Forecast Studio for Desktop.
- SAS Time Series Studio is not included with SAS Forecasting for Desktop.
- You do not have access to the SAS Forecast Server procedures outside of SAS Forecast Studio for Desktop.
- If you license SAS Forecasting for Desktop, you have access to all of the macros in the batch interface, except %FSCREATE and %FSRUNRPT. For more information about these macros, see the *SAS Forecast Server: Administrator's Guide*.

Benefits of Using SAS Forecasting for Desktop

Here are some of the benefits of using SAS Forecasting for Desktop.

- The entire application can be installed on one machine.
- You can quickly generate forecasts for up to 1000 time series.
- Results (graphs and tables) can be easily exported to Microsoft Office.
- Tabular results can be saved as SAS data sets.

How SAS Forecasting for Desktop Works

Given a time-stamped data set, the software provides the following automatic forecasting process:

1. accumulates the time-stamped data to form a fixed-interval time series
2. aggregates the time series to form a hierarchical time series.

For each time series in the hierarchy, SAS Forecasting for Desktop performs these steps:

1. diagnoses the time series using time series analysis techniques.
2. creates a list of candidate model specifications based on the diagnostics.
3. fits each candidate model specification to the time series.
4. generates forecasts for each candidate fitted model.
5. selects the most appropriate model specification based on either in-sample or holdout sample evaluation using a model selection criterion.
6. refits the selected model specification to the entire range of the time series.
7. creates a forecast score from the selected fitted model.
8. generates forecasts from the forecast score.
9. evaluates the forecast using in-sample analysis and provides for out-of-sample analysis of forecast performance.

For each forecast generated by the time series model, SAS Forecasting for Desktop performs these steps:

1. reconciles the model forecasts to form consistent reconciled forecasts.
2. evaluates the forecast using in-sample analysis and provides for out-of-sample analysis of forecast performance.

For each reconciled forecast, SAS Forecasting for Desktop performs these steps:

1. applies manual overrides to the reconciled forecasts.
2. reconciles the overrides to form consistent override forecasts.

How SAS Forecasting for Desktop Relates to Other SAS Software

SAS Forecasting for Desktop uses many of the procedures and options from other SAS products.

SAS Forecast Server Procedures

SAS Forecast Studio for Desktop uses several procedures as the basis for the automatic forecasting capabilities. These procedures automatically select the appropriate model for each item being forecast, which is 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. At the maximum level, 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, see the *SAS Forecast Server Procedures User's Guide*.

Note: These procedures are not available to you outside of SAS Forecast Studio for Desktop.

SAS/ETS

SAS/ETS software provides SAS procedures that perform econometric and time series analysis and forecasting, as well as financial analysis and reporting. The software also provides an interactive environment for time series forecast and investment analysis. For more information, see the *SAS/ETS User's Guide*.

Base SAS

Base SAS delivers a highly flexible and extensible fourth-generation programming language that is specially designed for data access, transformation, and reporting. It includes a rich library of encapsulated programming procedures for data manipulation, information storage and retrieval, descriptive statistics, and report writing. The output for SAS Forecast Studio for Desktop is generated by the Output Delivery System that is part of Base SAS.

SAS/GRAPH

SAS/GRAPH software provides high-impact visuals for all levels of your organization, enabling customers to readily understand complex information and empowering them to make informed, timely decisions. SAS/GRAPH software extends the power of SAS data management, business intelligence, and analytic tools, enabling customers to turn data into full-color graphs and charts.

What Is SAS Forecast Studio for Desktop?

SAS Forecast Studio for Desktop is a forecasting application that provides most of the same forecasting functionality as SAS Forecast Studio. Here is some of the key functionality in SAS Forecast Studio for Desktop:

- Generate forecasts automatically by using models that are shipped with SAS Forecast Studio for Desktop. You can generate forecasts by using a model selection list that you have created.
- Create your own forecasting models or combine generated 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.
- Add events to a model or to the entire project.
- Analyze the effect on the forecasts of various future values for the input series.
- Perform rolling simulations to determine the best model for your data.

However, SAS Forecast Studio for Desktop is a desktop-only application. All work is performed and saved locally. As a result, some functionality that is part of SAS Forecast Studio is not available from the desktop.

Documentation for SAS Forecast Studio for Desktop

Because the user interface for SAS Forecast Studio for Desktop is virtually identical to SAS Forecast Studio (which is a component of SAS Forecast Server), you can use the *SAS Forecast Studio: User's Guide* to work with this product. Online Help is also

available within SAS Forecast Studio for Desktop. (This help is shared with SAS Forecast Studio.)

Any functionality that is not available in SAS Forecast Studio for Desktop is noted in the product documentation. If you have additional questions about SAS Forecast Studio for Desktop, contact your SAS account representative.

Chapter 3

Starting SAS Forecast Studio

Understanding SAS Environments and Forecasting Environments	13
Starting SAS Forecast Studio	14
How You Can Run SAS Forecast Studio	14
Log On To SAS Forecast Studio	14
How You Can Run SAS Forecast Studio for Desktop	14

Understanding SAS Environments and Forecasting Environments

When you log on to SAS Forecast Studio and you license SAS Forecast Server, you must select the SAS environment (also known as the SAS deployment) that you want to use. A SAS environment is different from a forecasting environment. In the Log On dialog box, each SAS environment is identified by a name that is assigned by your site administrator. Although it is possible for a SAS environment and a forecasting environment to have the same name, the two environments are independent of each other.

Note: If your site licenses SAS Forecasting for Desktop, you do not need to log on and specify a SAS environment. For more information, see [Chapter 2, “About SAS Forecast Studio for Desktop,”](#) on page 7.

A SAS environment represents a deployment. A forecasting environment represents a product workspace for product sessions. Forecasting environments are created by your SAS Forecast Server administrator, or if you license SAS Forecasting for Desktop, you can create these environments by using the SAS Forecast Project Manager. (For more information, see [“About the SAS Forecast Project Manager”](#) on page 225.) These forecasting environments are used only by SAS Forecast Server and SAS Forecasting for Desktop.

SAS environments are available across products. Typically, your site administrator decides which SAS environments are available to you. The administrator also configures these environments. After an environment is configured, the Log On dialog box identifies the SAS environments that are available, applies any product-specific filtering to the list of environments, and then presents the list of available environments in the Log On dialog box. If your site administrator provided a description of a SAS environment, then this description is available as well. If you need access to an environment that does not appear in the drop-down list or you receive an error message instead of the Log On dialog box, contact your site administrator.

For more information about environments, see *SAS Forecast Server: Administrator's Guide*.

Starting SAS Forecast Studio

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** ⇒ **All Programs** ⇒ **SAS** ⇒ **SAS Forecast Studio** ⇒ **SAS Forecast Studio 14.1**.
- 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 SAS Forecast Studio.

To log on to SAS Forecast Studio:

1. Select your SAS environment. The list of environments is managed by your site administrator. For more information, see *SAS Forecast Server: Administrator's Guide*.
2. Specify your user name and password for the SAS Forecast Server.
3. (Optional) If security policies at your site allow it, select the **Remember password** check box for SAS Forecast Server to remember your password.

This option saves the password for the selected SAS environment and automatically uses that password for subsequent logon. Each SAS deployment has an independent policy setting to control whether passwords can be stored. This policy setting is decided by your organization. For more information about client-side storage of passwords, see the *SAS Intelligence Platform: Security Administration Guide*.

4. Click **Log On** to start SAS Forecast Studio.

How You Can Run SAS Forecast Studio for Desktop

SAS Forecast Studio for Desktop runs on Windows. To start SAS Forecast Studio for Desktop, select **Start** ⇒ **All Programs** ⇒ **SAS** ⇒ **SAS Forecast Studio** ⇒ **SAS Forecast Studio for Desktop 14.1**.

SAS Forecast Studio for Desktop appears. You do not need to log on to start the application.

Note: The Java Web Start client is not available if you are running SAS Forecast Studio for Desktop. For more information, see [Chapter 2, “About SAS Forecast Studio for Desktop,”](#) on page 7.

Chapter 4


Using the Workspace

Overview of the Workspace	18
The Hierarchy View	19
The Table View	20
Failed Forecasts	21
The Forecasting View	22
About the Forecasting View	22
Copy Data and Graphs from the Forecasting View to the Clipboard	24
The Modeling View	24
The Series View	26
The Scenario Analysis View	28
Customizing Your Workspace	29
Customize the Table View	29
Customize the Contents of the Forecast Plot or Data Table	29
Edit the Properties in the Modeling View and Series View	29

Overview of the Workspace

The workspaces for SAS Forecast Studio and SAS Forecast Studio for Desktop consist of the following components:

Figure 4.1 SAS Forecast Studio and SAS Forecast Studio for Desktop Workspace



- 1 The toolbar displays some of the most commonly used 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, compare and combine models, and perform running simulations.
 - **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 or SAS Forecast Studio for Desktop. 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 for the project. To change the order of the BY variables in the hierarchy, see [“Configure the Hierarchy” on page 56](#).

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.

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 and SAS Forecast Studio for Desktop automatically create filters for all of the series and for each level of the hierarchy. In the following display, **All** is selected as the filter, so all of the series appear in the Table View.

You can create additional filters. These user-defined filters are added to the **Filter** drop-down list. For more information, see [“Create a Filter” on page 80](#).

After selecting a filter, click **Refresh** to update the contents in the Table View.

- 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).

Figure 4.3 Table View



Hierarchy View | Table View | Failed Forecasts

Statistics of fit calculated over an in-sample range.

Filter: All

(26 of 26 series) Refresh

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

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 by only one column at a time.

Note: 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.

If SAS Forecast Studio or SAS Forecast Studio for Desktop 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 29](#).

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 or SAS Forecast Studio for Desktop could not fit a model from the model selection list. And, instead, the system used a model from the default model selection lists. In the **Problems** column, these series are identified by "Model

list problem." For more information about these model selection lists, see [Appendix 5, "Default Model Selection Lists,"](#) on page 273.


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

About 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**.

Figure 4.4 Forecasting View



The Forecasting View is divided into the following components:

- 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 example, you can specify whether to show the historical data and the forecasts, only the historical data, or only the forecast values on the plot. Additional customizations can be made by using the Forecasting View Properties dialog box. For more information, see [“Customize the Contents of the Forecast Plot or Data Table” on page 29](#).

- 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 77](#).

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 forecasts.
- overrides. You can specify overrides only for future time periods. The cells that you can edit have a white background. You can copy and paste values into the cells that you can edit. For more information, see [“Copying and Pasting Overrides” on page 120](#).
- final forecasts.

To customize the content of the data table, see [“Customize the Contents of the Forecast Plot or Data Table” on page 29](#).

Copy Data and Graphs from the Forecasting View to the Clipboard

You can copy data or graphs from the Forecasting View to the Windows clipboard. After the data or graph is in the clipboard, you can paste it into another software application, such as Microsoft Excel.

Note: The data in the data table is formatted using SAS formats. These formats might not be supported by Microsoft Excel. Before copying the data, you should apply a format to the data that is supported by the application where you plan to paste the results.

To copy the data from the data table:

1. Open the Forecasting View and select **Series** ⇒ **Copy Table**. The Copy To Clipboard dialog box appears.
2. Select the data that you want to copy. You can copy the historical data, the forecasts generated by SAS Forecast Studio, or both. Click **OK**. The content is now available from the Windows clipboard.
3. Paste the results in another software application, such as Microsoft Excel.

To copy the graph:

1. Right-click the graph and select **Copy** from the pop-up menu. The graph is now available from the Windows clipboard.
2. Paste the results in another software application, such as Microsoft Excel.

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 model selection list also shows when a model has been rejected. The name of the current forecast model is in bold. For more information, see [“How Models Are Fit” on page 171](#).
- 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 in the log scale of the white noise probability test for predicted errors

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

You can open additional plots or tables. SAS Forecast Studio and SAS Forecast Studio for Desktop save your final configuration and use that same configuration the next time you open the Modeling View while working in the current project. When you close the current project, 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.

The Series View

The Series View provides diagnostics to help you identify model components that might help 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. You can also add a model to 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 series variable.

- plots of the autocorrelation function, partial autocorrelation function, inverse autocorrelation function, and white noise probability test.

You can open additional 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 you open the Series View while working in the current project. When you close the current project, these settings are lost.


Note: For some series or time series transformations, some of the plots or tables might not be available.

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:

- 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 which columns to display in the Table View:

1. Select **View** ⇒ **Select Table View Columns**. The Table View Columns dialog box appears.
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 which values to display in the forecast plot or in the data table:

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

Edit the Properties in the Modeling View and Series View

To specify the number of lags to plot in the Modeling View or the Series View:

1. Select the Modeling View or the Series View in the workspace.
2. Select **View** ⇒ **Edit Series and Modeling View Properties**. The Edit Series and Modeling View Properties dialog box appears.
3. Select the **Modeling View Properties** or **Series View Properties** tab.

4. Deselect the **Use default** check box and specify the number of lags in the text box. The default value is 24 or three times the length of the seasonal cycle, whichever is smaller.
5. Click **OK**. SAS Forecast Studio re-creates the plots and tables in the Modeling View or Series View.

Part 2

Getting Started in SAS Forecast Studio

<i>Chapter 5</i>	
Preparing an Input Data Set	33
<i>Chapter 6</i>	
Understanding Variable Roles	39
<i>Chapter 7</i>	
Forecasting Hierarchically	55
<i>Chapter 8</i>	
Working with Projects	61

Chapter 5

Preparing an Input Data Set

Preprocessing Your Data	33
Understanding Time Series Data	34
Requirements for Time Series Data	34
How SAS Forecast Studio Creates Time Series Data	34
Requirements for the Input Data Set	35
Required Format	35
Size Limitations of the Input Data Set	35
Examples of Input Data Sets	35
Working with Missing Values	36
How SAS Forecast Studio Interprets Missing Values	36
Specify How to Interpret Missing Values	37
Including Future Input Values in the Input Data Source	38

Preprocessing Your Data

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 Microsoft 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*.

If your site licenses SAS Forecast Server, you can use SAS Time Series Studio to preprocess your data. For more information, see the *SAS Time Series Studio: User's Guide*.

Note: If your site licenses SAS Forecasting for Desktop, SAS Time Series Studio is not available. You will need to preprocess your data using one of the SAS procedures

previously described. For more information, see [Chapter 2, “About SAS Forecast Studio for Desktop,”](#) on page 7.

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.

Your time series data must meet the following requirements:

- The data set contains one variable for each dependent variable.
- The data set contains a time ID variable that identifies the time period for each observation. When SAS Forecast Studio creates a project, 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.

How SAS Forecast Studio Creates Time Series Data

SAS Forecast Studio and SAS Forecast Studio for Desktop create the time series data through the following process:

1. The data is sorted by the BY variables and the time ID variable.
2. The data is accumulated to the appropriate time interval if the input is one of the following types:
 - timestamped data that is recorded at no particular frequency (also called transactional data)
 - data recorded at a smaller time interval than needed for forecasting
3. Your input data set might contain BY variables. If you do not use all of the BY variables in your project, the observations are aggregated.
4. 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 36.

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

Requirements for the Input Data Set

Required Format

SAS Forecast Studio and SAS Forecast Studio for Desktop require a SAS date or datetime variable in the data set to generate forecasts. SAS Forecast Studio and SAS Forecast Studio for Desktop generate 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 and SAS Forecast Studio for Desktop use 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 SAS date or datetime value of each observation.
- BY variables enable you to group observations into time series that can be used to create 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 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 [“Reserved Variable Names” on page 249](#).

For more information, see [Chapter 6, “Understanding Variable Roles,” on page 39](#).

Size Limitations of the Input Data Set

If your site licenses SAS Forecast Server (which includes the SAS Forecast Studio client), you can work with large input data sets. If your site licenses SAS Forecasting for Desktop (which includes the SAS Forecast Studio for Desktop client), your input data sets are limited to 1,000 time series. For more information, see [Chapter 2, “About SAS Forecast Studio for Desktop,” on page 7](#).

Examples of Input Data Sets

Here are two examples of input data sets:

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

Date	Revenue	Avg. Price	Holiday
JAN2011	18817	26.3	0

Date	Revenue	Avg. Price	Holiday
FEB2011	52573	25.3	0
...	...	...	...
DEC2011	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
JAN2011	355	25.3	0	Region 1	Product 1
FEB2011	398	25.3	0	Region 1	Product 1
...	...	...	...	...	...
JAN2011	555	19.8	0	Region 1	Product 2
FEB2011	390	25.3	0	Region 1	Product 2
...	...	...	...	...	...
JAN2011	301	27.1	0	Region 2	Product 1
FEB2011	350	25.3	0	Region 2	Product 1
...	...	...	...	...	...
JAN2011	314	27.2	0	Region 2	Product 2
FEB2011	388	25.3	0	Region 2	Product 2
...	...	...	...	...	...
DEC2011	518	20.3	1	Region 2	Product 2

Note: 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

How SAS Forecast Studio Interprets Missing Values

If your data contains missing values in variables other than the time ID variable (such as the dependent, independent, and reporting variables), you can specify how to interpret missing values (regardless of the variable role).

- You can specify how to replace missing values in the data. You can specify that SAS Forecast Studio or SAS Forecast Studio for Desktop should set missing values to one of the following values:
 - 0.
 - the accumulated average value.
 - the accumulated first nonmissing value.
 - the accumulated last nonmissing value.
 - the accumulated maximum value.
 - the accumulated median value.
 - the accumulated minimum value.
 - missing.
 - the next accumulated nonmissing value. Missing values at the end of the accumulated series remain missing.
 - the previous accumulated nonmissing value. Missing values at the beginning of the accumulated series remain missing.
- You can specify how missing values are removed from the accumulated time series. You can choose to keep all of the missing values, remove the beginning missing values, remove the ending missing values, or remove both the beginning and ending missing values.
- You can also specify how beginning and ending zero values are interpreted in the accumulated time series. You can choose to keep the beginning and ending zeros, set the beginning zero values to missing, set the ending zero values to missing, or set both the beginning and ending zero values to missing.

Specify How to Interpret Missing Values

When you create a project, you can specify how to treat missing values from the data preparation step in the New Project wizard. For more information, see [“Creating a New Project” on page 62](#).

After the project is created, you can change how SAS Forecast Studio or SAS Forecast Studio for Desktop interprets missing values from the Forecasting Settings dialog box.

To specify how to interpret missing values:

1. Select **Project** ⇒ **Forecasting Settings**. The Forecasting Settings dialog box appears.
2. In the selection pane, click **Data Preparation**.
3. Specify how to interpret embedded missing values. By default, missing values are interpreted as missing.
4. Specify whether to remove the leading or trailing missing values in the data set.
5. Specify whether to interpret any leading or trailing zero values as missing.
6. Click **OK**.

The project is automatically diagnosed using the new values for these options. For more information, see [“Process for Generating a Forecast” on page 73](#).

Including Future Input Values in the Input Data Source

To include future input (independent) values for the time series in your project, include these future values in the input data set. For example, suppose that you are running a book store. For the next three months, the publisher will sell you copies of a book for \$20. Because you know this future input value for the series, you want to use the values from your input data set rather than have SAS Forecast Studio or SAS Forecast Studio for Desktop generate a forecast.

Chapter 6

Understanding Variable Roles

Overview of the Types of Roles	39
Time ID Variables	40
What Is the Time ID Variable?	40
Basics on Time Intervals	40
Supported Time Intervals	40
Understanding SAS Time Intervals	42
Custom Time Intervals	45
Modify the Interval and Seasonal Cycle Length	46
Examples of Changing the Time Interval	46
Edit the Format of the Time ID Variable	48
Classification (BY) Variables	49
Dependent Variables	49
Independent Variables	50
Reporting Variables	50
Adjustment Variables	51
What Is an Adjustment Variable?	51
Add an Adjustment Variable	52
Examples of Adjustment Variables	52

Overview of the Types of Roles

When you create a project, 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 and SAS Forecast Studio for Desktop do not support time values. When you assign a variable to the time ID role, SAS Forecast Studio and SAS Forecast Studio for Desktop recognize the data's format. If your data uses a time format, a message stating that time data is not valid for the time ID variable appears. If your data has no format, SAS Forecast Studio and SAS Forecast Studio for Desktop assume that the values are dates, and errors occur when your project is created.

Basics on Time Intervals

All time intervals (whether they are shipped with SAS or custom intervals that you create) must meet the following criteria:

- A discrete time interval has a beginning and an ending SAS date or SAS datetime.
- For SAS date intervals, the ending date is defined as 1 day before the beginning of the next interval.
- For SAS datetime intervals, the ending time is 1 second before the beginning of the next interval.
- All observations with an identifying SAS date or SAS datetime that is between the beginning and the end of the interval t_i correspond to the interval it .
- Alignment refers to the identifying date of the interval and does not affect the definition of the interval.

Supported Time Intervals

SAS Forecast Studio and SAS Forecast Studio for Desktop support 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 4. 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 (or subperiods) is in days (DAY). Note that WEEKV differs from WEEK in that WEEKV.1 starts on Monday, WEEKV.2 starts on Tuesday, and so on.

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 (or subperiods) is in retail 4-4-5 months.

Retail 4-5-4 Year

specifies ISO 8601 weekly interval, except that the starting subperiod (or subperiods) is in retail 4-5-4 months.

Retail 5-4-4 Year

specifies ISO 8601 weekly interval, except that the starting subperiod (or subperiods) 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 12th 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 12th 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 contain a leap week. The starting subperiod (or subperiods) 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 contain a leap week. The starting subperiod (or subperiods) 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 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 10-day intervals. Each month consists of three periods. The first period is the 1st through the 10th day of the month. The second period is the 11th through the 20th day of the month. The third period is the 21st 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 and SAS Forecast Studio for Desktop analyze 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 datetime 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, SAS Forecast Studio and SAS Forecast Studio for Desktop accumulate 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 and SAS Forecast Studio for Desktop supply 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.

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 generates 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 or SAS Forecast Studio for Desktop 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 or SAS Forecast Studio for Desktop.

Here is the syntax for an interval:

name<*multiplier*><*.starting-point*>

Here is an explanation of each of the user-supplied values:

name

is the name of the interval.

multiplier

specifies the multiplier 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	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 16th of each month	Semimonthly periods	SEMIMONTH2.2 - intervals from the 16th of one month through the 15th 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.

Interval Name (in SAS code format)	Default	Shift Period	Example
TENDAYm.s	First, 11th, and 21st 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.
HOURLm.s	Start of the day (midnight)	Hours	HOURL8.7 specifies eight-hour intervals starting at 6:00 a.m., 2:00 p.m., and 10:00 p.m.

Custom Time Intervals

What Is a Custom Time Interval?

Depending on your data, the standard time intervals that are available in SAS Forecast Studio or SAS Forecast Studio for Desktop might not be appropriate. Here are some examples of when you might need to create a custom interval:

- In SAS, the MONTH interval begins on the 1st of every month, but your fiscal months begin on a different day (such as the 10th) of every month.
- Using a standard interval results in gaps in your data. For example, you have retail data that is collected hourly. However, the business is closed at night. You do not want to include the hours when the business is closed in your forecasts.
- You have irregular gaps in the data, such as holidays.
- You need to transform a function to improve stationarity.

A custom time interval is a base time interval that you define by using a SAS data set. For more information about custom intervals, see *SAS Language Reference: Concepts* and the *SAS/ETS User's Guide*.

Creating a Custom Time Interval

In order to create a custom time interval, you must be familiar with SAS programming and have access to the autoexec.sas or sasv9_usermods.cfg file for the SAS server. Therefore, it is recommended that you ask your site administrator to create the custom interval that you need. For more information about creating a custom time interval, see *SAS Forecast Server: Administrator's Guide*.

How to Select a Custom Time Interval

Custom intervals should appear in the same drop-down lists as the time intervals that are shipped with SAS Forecast Studio and SAS Forecast Studio for Desktop. For example, you can select a custom interval when creating a new project in the New Project wizard. After the project is created, you can select the time interval from the Forecasting Settings dialog box. If you do not see the interval that you need or you need more information about an interval, contact your site administrator.

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 created. For more information, see [“Time ID Variables” on page 40](#).

When you assign a variable to the time ID role, SAS Forecast Studio and SAS Forecast Studio for Desktop try 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 or SAS Forecast Studio for Desktop uses an interval of one month, and a seasonal cycle length that corresponds 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:

1. Select **Project** ⇒ **Forecasting Settings**. The Forecasting Settings dialog box appears.
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 that is 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 or SAS Forecast Studio for Desktop 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 and SAS Forecast Studio for Desktop, 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 (or SAS Forecast Studio for Desktop) needs to use an interval of WEEK.2.

To change the first day of the week to Wednesday:

1. Select **Project** ⇒ **Forecasting Settings**. The Forecasting Settings dialog box appears.
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 and SAS Forecast Studio for Desktop assign to each day of the week, see [“Supported Time Intervals” on page 40](#).


Example 2: Specify a Yearly Interval

In SAS Forecast Studio and SAS Forecast Studio for Desktop, 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 (or SAS Forecast Studio for Desktop) needs to use an interval of YEAR2.6.

To specify an interval of YEAR2.6:

1. Select **Project** ⇒ **Forecasting Settings**. The Forecasting Settings dialog box appears.
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 6.1 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 and SAS Forecast Studio for Desktop assume 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 (or SAS Forecast Studio for Desktop).

To specify the days in the weekend:

1. Select **Project** ⇒ **Forecasting Settings**. The Forecasting Settings dialog box appears.
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 appears.
6. Select the check boxes for **Sunday**, **Friday**, and **Saturday**. Click **OK**.

Figure 6.2 Weekend Days Are Friday, Saturday, and Sunday



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

After SAS Forecast Studio (or SAS Forecast Studio for Desktop) 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 and SAS Forecast Studio for Desktop automatically determine 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 Formats and Informats: Reference*.

To edit the format of the time ID variable:

1. Select **Project** ⇒ **Forecasting Settings**. The Forecasting Settings dialog box appears.
2. In the selection pane, click **Time ID**.
3. Click **Edit**. The Edit Time ID Format dialog box appears.
4. Select the format that you want to use.

Note: The WEEKUw., WEEKVw., and WEEKWw. formats are not supported.

5. (Optional) Specify the format width.
6. Click **OK** in the Edit Time ID Format dialog box.
7. Click **OK** in the Forecasting Settings dialog box.

Classification (BY) Variables

SAS Forecast Studio and SAS Forecast Studio for Desktop group 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.


For more information about hierarchies, see [Chapter 7, “Forecasting Hierarchically,”](#) on [page 55](#).

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, you assign the Sales variable to the dependent variable role. When the project is created, the forecasting plot shows the sales forecast, as shown in the following display.

Figure 6.3 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 and SAS Forecast Studio for Desktop try to use the independent variables in the model generation. SAS Forecast Server will forecast future values for your independent variables if future values do not already exist.

For example, in your initial forecasts, you assign the 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 Cost variable has on your forecasts, so you assign the Cost variable to the independent variable role. SAS Forecast Studio and SAS Forecast Studio for Desktop use this independent variable when generating the models, but in this case, the forecasts do not change.

Figure 6.4 Forecast Values When Cost 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 for the Cost variable do not influence the forecasts.

Figure 6.5 Cost as a Reporting Variable



Node Report Series Table

Sales Region=Region1 Name of product line=Line1 Product Name=Product3

Order Date	sale	Unit Cost
JAN1998	300.000	44.00
FEB1998	301.000	20.13
MAR1998	359.000	20.15
APR1998	397.000	20.10
MAY1998	344.000	20.39
JUN1998	345.000	38.00
JUL1998	344.000	20.15
AUG1998	338.000	20.15
SEP1998	343.000	20.15

Adjustment Variables

What Is an Adjustment Variable?

Systematic variations and deterministic components are included in time series data. By adjusting the time series for known systematic variations or deterministic components, the underlying stochastic (unknown) time series process can be more readily identified and modeled. For more information about adjustment variables, see the “Adjustment Step” topic in the *SAS Forecast Server Procedures: User’s Guide*.

You can specify adjustments in the following ways:

- before generating the forecasts. After the timestamped 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 the inverse of each other.

Add an Adjustment Variable

To add an adjustment variable:

1. Select **Project** ⇒ **Hierarchy and Variable Settings**. The Hierarchy and Variable Settings dialog box appears.
2. Click **Adjustments**. The Adjustments dialog box appears.
3. Click **New**. The Adjustment Properties dialog box appears.
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 (or SAS Forecast Studio for Desktop) adjusts the forecasts by multiplying the sales values by the exchange rate.

To create this adjustment variable:

1. Select **Project** ⇒ **Hierarchy and Variables Settings**. The Hierarchy and Variables Settings dialog box appears.
2. Click **Adjustments**. The Adjustments dialog box appears.
3. Click **New**. The Adjustment Properties dialog box appears.
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 **Multiply**.

Figure 6.6 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 (or SAS Forecast Studio for Desktop) subtracts the advanced bookings from the time series. After the forecasts are generated, SAS Forecast Studio (or SAS Forecast Studio for Desktop) adjusts the forecasts by adding the advanced bookings.

To create this adjustment variable:

1. Select **Project** ⇒ **Hierarchy and Variable Settings**. The Hierarchy and Variable Settings dialog box appears.
2. Click **Adjustments**. The Adjustments dialog box appears.
3. Click **New**. The Adjustment Properties dialog box appears.
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 6.7 Adjustment Properties Dialog Box for the Advanced Bookings



The dialog box is titled "Adjustment Properties" and contains the following fields:

- Adjusted variable:** A list box containing "Sales" with a checkmark.
- Adjusting variable:** A dropdown menu showing "Advanced Bookings".
- Description:** A text area containing "Adjusting a time series for advanced bookings".
- Accumulation:** A dropdown menu showing "Sum of values".
- Aggregation:** A dropdown menu showing "Sum of values".
- Pre-operation:** A dropdown menu showing "Subtract".
- Post-operation:** A dropdown menu showing "Add".

At the bottom right, there are three buttons: "OK", "Cancel", and "Help".

9. Click **OK**.

Chapter 7

Forecasting Hierarchically


Understanding Hierarchies	55
Configure the Hierarchy	56
Understanding Reconciliation Methods	56
What Are the Reconciliation Methods?	56
Select the Reconciliation Method	57
Understanding Aggregation and Accumulation	57
Difference between Aggregation and Accumulation	57
The Accumulation and Aggregation Methods	58
Specify the Aggregation and Accumulation Methods	60
Example: How Forecasts Are Generated	60

Understanding Hierarchies

To forecast hierarchically, you must assign at least one classification (BY) variable. When you assign a classification variable, you can specify whether SAS Forecast Studio or SAS Forecast Studio for Desktop should forecast your data hierarchically. If you choose to forecast your data hierarchically, then SAS Forecast Studio and SAS Forecast Studio for Desktop sort 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.

An example of a hierarchy is Region > Product Line > Product Name. You can preview the hierarchy for your project.

Figure 7.1 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.

When you forecast hierarchically, you must specify whether SAS Forecast Studio or SAS Forecast Studio for Desktop should reconcile the hierarchy and whether to resolve any override conflicts. You can also select the reconciliation method to use. For more information, see [“Understanding Reconciliation Methods” on page 56](#).

For the dependent, independent, adjustment, and reporting variables in your project, you can specify how to aggregate the data across the levels of the hierarchy. For more information, see [“Understanding Aggregation and Accumulation” on page 57](#).

Configure the Hierarchy

When you create a project, you specify the BY variables to include in the hierarchy and the order of the hierarchy in the New Project wizard. After the project is created, you cannot remove BY variables from the hierarchy or add BY variables to the hierarchy. However, you can change the order of the hierarchy.

If the source data set is very large, consider creating more than one project and grouping time series with similar properties. This type of grouping is a recommended practice.

Note: Changing any of the hierarchy settings results in rebuilding the entire project.

To change the order of the hierarchy:

1. Select **Project** ⇒ **Hierarchy and Variable Settings**. The Hierarchy and Variable Settings dialog box appears.
2. Select the variable that you want to move and click the up or down arrow. To see your changes, click **Preview**.
3. When you are finished, click **OK** to apply the changes to your project.

Understanding Reconciliation Methods

What Are the Reconciliation Methods?

When forecasting hierarchically, 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 and SAS Forecast Studio for Desktop then use 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 some components of 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 using the lowest level of the hierarchy as a reference level for the forecasts, 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 and SAS Forecast Studio for Desktop use 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 or SAS Forecast Studio for Desktop, 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**.

Select the Reconciliation Method

When you create a project, you can specify the reconciliation method to use. By default, SAS Forecast Studio and SAS Forecast Studio for Desktop use the top-down reconciliation method.

Note: Changing any of the hierarchy settings results rebuilding the entire project.

To change the reconciliation method:

1. Select **Project** ⇒ **Hierarchy and Variable Settings**. The Hierarchy and Variable Settings dialog box appears.
2. From the **Reconcile the hierarchy** drop-down list, select the reconciliation method to use, and click **OK**.

You can specify additional reconciliation options for your project. For more information, see [“Specify the Advanced Reconciliation Options” on page 132](#).

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. Accumulation combines data within the same time interval. For example, you can aggregate all of the series of each product in the Electronic group into a single series that represents the group as a whole. In the New Project wizard, you can specify the aggregation and accumulation options for the dependent, independent, adjustment, and reporting variables in your project.

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 refers to data aggregation in the time domain. It 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.

Aggregation is restricted to the total and average methods, but the same equations apply.

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

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

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

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

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

Average

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

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

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

Minimum

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

$$a = \min(\{r_q\}_{q=1}^Q)$$

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

Maximum

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

$$a = \max(\{r_q\}_{q=1}^Q)$$

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

Median

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

$$a = \text{median}(\{r_q\}_{q=1}^Q)$$

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

Number of Nonmissing Observations

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

$$a = Q_N$$

Number of Observations

accumulates the vector values based on the number of values.

$$a = Q$$

Number of Missing Observations

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

$$a = Q_{NMISS}$$

First Occurrence

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

$$a = r_1$$

Last Occurrence

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

$$a = r_Q$$

Standard Deviation

accumulates the vector values based on their standard deviation.

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

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

Uncorrected Sum of Squares

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

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

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

Corrected Sum of Squares

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

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

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

Specify the Aggregation and Accumulation Methods

When you create a project, you can specify the aggregation and accumulation method to use for the dependent, independent, adjustment, and reporting variables in the New Project wizard.

Note: Changing any of the hierarchy settings results rebuilding the entire project.

To specify an aggregation or accumulation method for a variable in your project:

1. Select **Project** ⇒ **Hierarchy and Variable Settings**. The Hierarchy and Variable Settings dialog box appears.
2. From the **Reconcile the hierarchy** drop-down list, select the reconciliation method to use, and click **OK**.

Example: How Forecasts Are Generated

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.

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

Chapter 8

Working with Projects

What Is a Project?	61
Creating a New Project	62
Overview of the New Project Wizard	62
Step 1: Name Your Project	62
Step 2: Select the Data Set That You Want to Forecast	63
Step 3: Specify Classification (BY) Variables and Whether to Forecast Hierarchically	63
Step 4: Specify the Properties of the Time Dimension of Your Data	63
Step 5: Assign Roles to Variables in Your Data	64
Step 6: Specify How to Prepare the Data for Each Forecast	64
Step 7: Finish Creating Project	64
Step 8: Optional Project Actions	65
Migrating SAS Forecast Studio Projects	65
Open an Existing Project	66
Sharing a Project	67
Updating Your Project Data	67
Changing the Input Data Source	68
Export the Project Settings	68
Viewing Properties	68
View Project Properties	68
View Series Properties	69
Save a Copy of the Project	69
Create a Shortcut for a Project	70

What Is a Project?

SAS Forecast Studio and SAS Forecast Studio for Desktop organize 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.)

The default location for the project files depends on your SAS configuration directory.

- If you are running SAS Forecast Studio (which is a component of SAS Forecast Server), an example of a default directory path in a Windows operating environment is `C:\SAS\Forecast\Lev1\AppData\SASForecastServer14.1`.
- If you are running SAS Forecast Studio for Desktop (which is a component of SAS Forecasting for Desktop), an example of a default directory in a Windows operating environment is `C:\SAS\ForecastStudioDesktop\Environments\Default`.

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*.

Note: For a project to work correctly, any custom formats, custom time intervals, external model repositories, and external model selection lists must be available. If you move, migrate, or copy a project to a different environment, verify that these items are also available in that environment.

If the source data set is very large, consider creating more than one project and grouping time series with similar properties. This type of grouping is always a recommended practice for statistical modeling concerns as well as for reliability.

Creating a New Project

Overview of the New Project Wizard

You can have only one project open in SAS Forecast Studio or SAS Forecast Studio for Desktop at a time. You can create a new project either when you first open SAS Forecast Studio (or SAS Forecast Studio for Desktop) 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 or SAS Forecast Studio for Desktop, click **New** in the Projects dialog box.
- To create a new project when SAS Forecast Studio or SAS Forecast Studio for Desktop is already running, select **File** ⇒ **New Project** or click  in the toolbar.

Step 1: Name Your Project

First, select the forecasting environment for your project. For more information, see [“Understanding SAS Environments and Forecasting Environments”](#) on page 13.

Specify a name and description for your project. The project name must be a valid SAS name. For more information about SAS naming conventions, see *SAS Language Reference: Concepts*.

In this step, you can also add any start-up and shutdown code that you want to run. For more information, see “Adding Start-Up and Shutdown Code to a Project” in the *SAS Forecast Server: Administrator's Guide*.

Click **Next**.

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

When selecting a data set, remember the following conditions:

- Before you can use a data set in SAS Forecast Studio or SAS Forecast Studio for Desktop, it must be in the required format. For more information, see [Chapter 5, “Preparing an Input Data Set,”](#) on page 33.

You can also use SAS Time Series to export a data set that you can import into SAS Forecast Studio. For more information, see *SAS Time Series Studio: User’s Guide*.

- 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 or SAS Forecast Studio for Desktop. If you create a project using a data set in the Work library, then SAS Forecast Studio or SAS Forecast Studio for Desktop will be unable to find this data the next time you open the project. Also, no files in the Work library should be modified while SAS Forecast Studio or SAS Forecast Studio for Desktop is running.

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** check box.
3. To reconcile the hierarchy after any changes, select the **Reconcile the hierarchy** check box and then select the reconciliation method to use.
4. To select a disaggregation method, to specify how to resolve conflicts between locked overrides, or whether to restrict the direction of the reconciliation, 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. To specify a custom interval, select **Custom** from the **Interval** drop-down list. Select the name of the custom interval. For more information, see [“Custom Time Intervals”](#) on page 45.

Note: By default, SAS Forecast Studio and SAS Forecast Studio for Desktop determine the time interval automatically. For more information about when the time interval cannot be detected and other common errors, see [“Troubleshooting the Time ID Variable”](#) on page 245.

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. If you are forecasting hierarchically, you can use same method for accumulation and aggregation.
3. Specify how the independent variables should be used in a system-generated model. You can choose from these options:
 - **Try to use** – specifies that the independent variable be included in the model as long as its parameters are significant and the increment of the value of criterion exceeds a threshold.
 - **Use if significant** – specifies that the independent variable be included in the model as long as the parameters of the independent variable are significant.
 - **Force to use** – specifies that the independent variable be included in the model as long as the model does not fail to be diagnosed. When you select this option, the delay and orders for the numerators and denominators are set to zero.

These options are also available with either the positive or negative option in parentheses. For example, selecting **Force to use (positive)** drops the input variable from the model if its coefficient is negative, and selecting **Force to use (negative)** implies the opposite. The specification of positive or negative does not mean that constraints are imposed during the estimation of the variable's coefficient in the model.
4. (Optional) To specify any adjustment variables, click **Adjustments**.
5. Click **Next**.

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

1. Specify how SAS Forecast Studio or SAS Forecast Studio for Desktop should interpret missing values in the data, whether to remove leading or trailing missing values, and which leading or trailing zero values to interpret as missing.
2. (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.
3. Click **Next**.

Step 7: Finish Creating Project**General Steps**

1. Specify the number of periods to forecast (also called the horizon). By default, the horizon is 12.
2. To use the project settings from an external file, click **Import project settings**.
3. To change the diagnostics or model settings, click **Change other forecasting settings**. The Forecasting Settings dialog box appears. Most of these forecasting options can be set after the project is created. An exception is the option for

suppressing the component series data set. For more information, see [“Suppress the Component Series Data Set” on page 65](#).

4. To save the project settings that are specified in the Forecasting Settings dialog box, click **Export project settings**. The project settings are saved to an external file and can be imported later to specify the project settings for another project.
5. To add an event to your project, click **Create and import events**. For more information, see [“Creating New Events” on page 101](#).
6. Click **Next**.

Suppress the Component Series Data Set

By default, the component series data set is automatically created when you create a project. A component data set is created for each BY variable in the project. Because a component data set can also contain several other variables, these data sets could be quite large. As a result, a component data set could require additional disk space and computing time. For more information, see [“OUTCOMPONENT= Data Set” on page 257](#).

If you want to suppress the creation of the component data set, you must specify this option when you create the project. However, if you suppress this data set, then you will not see the effects of events or outliers in the plot and data table in the Forecasting View. Suppressing the component data set does not affect the component plots and tables that are available in the Modeling View.

To suppress the creation of the component series data set:

1. In the selection pane of the Forecasting Settings dialog box, click **Forecast**.
2. Deselect the **Create the component series data set** check box.
3. Click **OK** in the Forecasting Settings dialog box to return to the New Project wizard.

Step 8: Optional Project Actions

Specify whether to generate the forecasts when you finish or to save the SAS code to run later. Click **Finish**.

You can monitor the progress of the forecasts from the status bar. If you click **Cancel** in the progress indicator, the project is not created, and SAS Forecast Studio deletes any files that were created.

After the project is created, the Forecast Summary dialog box appears. For more information, see [“Reviewing a Summary of the Forecast Run” on page 76](#).

Migrating SAS Forecast Studio Projects

Note: Migration is specific to customers who are running SAS Forecast Studio, which is a component of SAS Forecast Server.

Before you can open a SAS Forecast Studio project from a previous release, your site administrator must migrate your project to the current environment. For more information about migrating projects, see the *SAS Forecast Server: Migration Guide*.

To update a SAS Forecast Studio project to the current release:

1. Open the project in the current release of SAS Forecast Studio. The Update Project Version dialog box appears.
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. If it has been updated, SAS Forecast Studio updates the project data and reforecasts the existing series based on the option that you selected.
5. Click **OK**. SAS Forecast Studio automatically converts the project to be compatible with the current release.


Open an Existing Project

You can open a project in the following ways:

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

If the input data set has changed since the last time you opened the project, the Update Project Data dialog box appears. By default, the **Cancel: do not update project data** check box is selected, so the project data is not updated. If you select a different option, SAS Forecast Studio or SAS Forecast Studio for Desktop updates the project data and reforecasts any existing series.



For more information about the options that are available when you reforecast a project, see [“Reforecast a Project” on page 78](#).

Sharing a Project

By default, only the creator of the project and administrators have permission to open or edit an existing project. Any projects that are not available to you appear dimmed in the Projects dialog box. After initially creating a project, you decide to share it with other users in the same forecasting environment.

Note: A project can be opened by only one user at a time.

To allow others to view and edit one of your projects:

1. Select **File** ⇒ **Project Properties**. The Project Properties dialog box appears.
2. Select the **Allow other users to view and edit this project** check box. Only the project creator or an administrator can change this option.
3. Click **OK** to save your changes and to close the Project Properties dialog box.


Note: Although the project is shared, the functionality that is available depends on the capabilities that are available to that user. For more information, see [“Understanding Role-Based Settings in SAS Forecast Studio”](#) on page 6.

Updating Your Project Data

The data for the project might have changed since the last time you forecast your data. For example, the data might include a new level of a BY variable (such as a new product) or new observations (such as sales data for the past three months). If the input data source for your project is a data view (rather than a SAS data set) and the data view is updated, then you must update the data in your project to generate accurate forecasts.

Note: You cannot change the structure of an existing project. For example, you cannot add a new BY variable to the project hierarchy.

When you reforecast a project, SAS Forecast Studio or SAS Forecast Studio for Desktop recognizes if the project data has changed. In the Forecast Project dialog box, you can choose whether to use the updated version of the data.



For more information about how to update your project data, see [“Reforecast a Project”](#) on page 78.

Note: 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 the *SAS Forecast Server: Administrator's Guide*.

Changing the Input Data Source

To change the input data source after the project has been created:

1. Select **File** ⇒ **Change Data Source**. The Data Set Selector dialog box appears.
2. Select the data set that you want to use and click **OK**.
3. To apply the changes immediately, click **Yes** in the confirmation dialog box.

Export the Project Settings

You might want to use the same project settings for every project. Instead of specifying these settings every time you create a project, you can import the project settings from an external file. The project settings are specified in the Forecasting Settings dialog box.

To export the settings for the current project:

1. Select **File** ⇒ **Export Project Settings**. The Export Project Settings dialog box appears.
2. Specify the name of the file and click **Save**.

Viewing Properties

View Project Properties

To view the properties for a project:

1. Select **File** ⇒ **Project Properties**. The Project Properties dialog box appears.
2. (Optional) Specify the code to run when the project starts or when you close the project. Click **Modify** to add this code.
3. (Optional) To add a project description, enter the description in the **Description** box.
4. (Optional) To share this project with other users at your site, select the **Allow other users to view and edit this project** check box. This option is not selected by default. If this check box is not selected, only the person who created the project can open it. Only the project creator or a SAS Forecast Server administrator can change this option.
5. Click **OK** to save your changes and to close the Project Properties dialog box.


View Series Properties

To view the properties for a series:

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. The series properties show the statistics for the raw data.

Figure 8.1 Series Properties for Node in Hierarchy View



Save a Copy of the Project

To save a copy of the project:

1. Select **File** ⇒ **Projects**. The Projects dialog box appears.
2. Specify a name for the project. This name must be a valid SAS name. By default, the name is *project-name_Copy_n*, where **n** is the lowest available integer.
3. Click **OK**.

Create a Shortcut for a Project

For local installations of SAS Forecast Studio and SAS Forecast Studio for Desktop, you can create a shortcut to your project on your Windows desktop.

To create a shortcut to a project:

1. Select **File** ⇒ **Save Project Link**. The Save dialog box appears.
2. Specify a name for the project link. By default, the name is *project-name.fs*.
3. Click **Save**.

Part 3

Analyzing Your Results

<i>Chapter 9</i>	
Generating Forecasts	73
<i>Chapter 10</i>	
Working with Filters	79
<i>Chapter 11</i>	
Running Reports	91

Chapter 9

Generating Forecasts

Process for Generating a Forecast	73
Set the Diagnose Properties for Each Series	74
Specify the Number of Periods to Forecast	75
Create a Forecast Data Set for Independent Variables	76
Reviewing a Summary of the Forecast Run	76
Specifying an Inactive Series	77
What Is an Inactive Series?	77
Set a Series as Inactive	77
Reforecasting a Series or Project	78
Reforecast a Series	78
Reforecast a Project	78

Process for Generating a Forecast

SAS Forecast Studio and SAS Forecast Studio for Desktop use 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 and SAS Forecast Studio for Desktop perform the following diagnostic tests:

- 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.
- 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 74](#).

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 and SAS Forecast Studio for Desktop select 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 and SAS Forecast Studio for Desktop use 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 and SAS Forecast Studio for Desktop automatically regenerate 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 or SAS Forecast Studio for Desktop 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:

1. Select **Project** ⇒ **Forecasting Settings**. The Forecasting Settings dialog box appears.
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, enter 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, enter 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 or SAS Forecast Studio for Desktop 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. Specify whether to refine insignificant parameters of the final model, identify the factors to retire, and identify the order of factors.
13. Click **OK**.

Specify the Number of Periods to Forecast

To specify how to fit the forecast model for each series:

1. Select **Project** ⇒ **Forecasting Settings**. The Forecasting Settings dialog box appears.
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.

Create a Forecast Data Set for Independent Variables

In some situations, you might not know the future values of the independent variables. Depending on the model for your project, SAS Forecast Studio might use these independent variables to create the forecasts for your dependent variables. SAS Forecast Server will forecast future values for your independent variables if future values do not already exist. Using SAS Forecast Studio, you can create a data set (called OUTINDEP=) that contains the future values of the independent variables. You can review this data set to verify whether these values are reasonable. For more information, see [“OUTINDEP= Data Set” on page 258](#).

To create a forecast data set for the independent variables:

1. Select **Project** ⇒ **Forecasting Settings**. The Forecasting Settings dialog box appears.
2. In the selection pane, click **Forecast**.
3. Select the **Create forecast data set for independent variables** check box.

Note: For more information about the model used to generate the forecasts, see [“OUTEST= Data Set” on page 258](#).

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:

- the weighted MAPE.
- a graph of the distribution of the specified statistic of fit.
- a graph of the different types of model families. This graph shows how each model type (ARIMA, Combined, Exponential Smoothing, Intermittent Demand, and Unobserved Components) fits the series.
- a graph of the model type. 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 table of the number of errors. This table shows the number of failed forecasts, the number of series where SAS Forecast Studio or SAS Forecast Studio for Desktop 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 you specify a series as inactive, the following occurs:

- No baseline forecast is selected in the Modeling View. As a result, no forecasts are produced for the series.
- Any child series must also be inactive. SAS Forecast Studio and SAS Forecast Studio for Desktop prompt 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.

Set a Series as Inactive

To specify an inactive series:

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

A series remains inactive even after SAS Forecast Studio or SAS Forecast Studio for Desktop 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:

1. Select the series in the hierarchy or table view and select **Series** ⇒ **Forecast Series**. The Forecast Series dialog box appears.
2. Specify how to forecast the series.
3. (Optional) To have SAS Forecast Studio or SAS Forecast Studio for Desktop 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 to select a forecast model, see [“Specifying the Forecast Model” on page 174](#).
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:

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

Note: If you choose any option other than **Diagnose: create new system-generated models, refit all other models, and select a forecast model**, then any new series are forecast using the model selection list that is shipped with SAS Forecast Studio and SAS Forecast Studio for Desktop. For more information, see [Appendix 5, “Default Model Selection Lists,” on page 273](#).

3. (Optional) To have SAS Forecast Studio or SAS Forecast Studio for Desktop 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 to select a forecast model, see [“Specifying the Forecast Model” on page 174](#).
4. (Optional) Select the **Use updated data if available** check box. SAS Forecast Studio or SAS Forecast Studio for Desktop 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

What Is a Filter?	79
Create a Filter	80
View the Results from a Filter	80
Examples of Filters	81
Example 1: Creating a Simple Filter	81
Example 2: Creating a Filter with Multiple Expressions	83
Example 3: Creating a Filter for Missing Values	85
Example 4: Create a Filter for a Combined Model	86
Example 5: Create a Filter for an Input Variable	87

What Is a Filter?

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

- 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 or SAS Forecast Studio for Desktop, you can create additional filters. By default, these filters are created using the dependent variable. However, you can also create filters using the input or report variables. If you create a filter using an input or report variable, you can specify only series properties (such as the first observation, minimum value, number of missing values, and standard deviation) in the filter criteria. For an example of how to create a filter, see [“Examples of Filters” on page 81](#).

To view the results from a filter, click the **Table View** tab in the overview panel. If the filter is for an input or report variable, SAS Forecast Studio or SAS Forecast Studio for Desktop returns only one dependent series if the project contains multiple dependent variables.

Create a Filter

To create a filter:

1. Select **Tools** ⇒ **Filters**. The Filters dialog box appears.
2. Click **New**. The Filter Properties dialog box appears.
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 whether to create the filter for a dependent variable or for an input or report variable. If you select **Input or Report variables**, you must specify the variable that you want to use.
6. Specify the filter by using the columns in the table.

Note: If you are creating a filter for an input or report variable, you can specify only series properties as the filter criteria. So when you select an input or report variable, the **Category** column is not available from the table.

7. (Optional) To add additional filter expressions, click .

Note: If you are creating a filter for an input or report variable, you must click  to add the filter criteria.

8. (Optional) If you are creating a filter for a dependent variable, specify whether the series must have overrides, notes, or both.
9. Click **OK**. The new filter now appears in the Filters dialog box.

View the Results from a Filter

To view the series that meet the filter criterion:

1. In the overview panel, click the **Table View** tab.
2. From the **Filter** drop-down list, select the filter that you want to apply. SAS Forecast Studio or SAS Forecast Studio for Desktop 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:


1. Select **Tools** ⇒ **Filters**. The Filters dialog box appears.
2. Click **New**. The Filter Properties dialog box appears.
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. For the **Filter by** option, select **Dependent variables**.
7. In the table, specify the following values for each column:
 - a. In the **Category** column, select **Reconciled (Rec.) statistic of fit**.
 - b. In the **Property** column, select **Mean absolute percent error**.
 - c. In the **Condition** column, select **>**.
 - d. In the **Value** column, type **3** and press ENTER.

The Filter Properties dialog box should look like the following display:

Figure 10.1 Completed Filter Properties Dialog Box for Filter Example



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

Figure 10.2 New Filter in Filters Dialog Box

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

To view the series that meet this filter criterion:

1. In the overview panel, click 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

Hierarchy View Table View Failed Forecasts				
Statistics of fit calculated over an in-sample range.				
Filter: MAPE				
(21 of 26 series)			Refresh	
REGION...	PRODUC...	PRODUC...	MAPE /	Rec. MAPE
Region3	Line4	*	2.42	3.02
Region1	*	*	3.18	3.08
Region1	Line1	*	3.18	3.08
Region3	Line5	*	3.19	3.62
Region2	Line2	Product7	3.67	4.29
Region1	Line1	Product1	3.93	3.97
Region2	Line3	Product11	3.97	3.96
Region3	Line4	Product13	4.05	5.47
Region3	Line5	Product17	4.07	4.17
Region2	Line2	Product5	4.21	4.32
Region3	Line5	Product16	4.26	4.90
Region1	Line1	Product3	4.48	5.15
Region2	Line3	Product9	4.52	4.32
Region3	Line4	Product12	4.54	4.86
Region1	Line1	Product2	4.64	4.93
Region2	Line3	Product8	4.66	4.66
Region2	Line3	Product10	4.74	5.00
Region3	Line4	Product15	4.81	4.30
Region3	Line4	Product14	4.99	5.08
Region2	Line2	Product4	5.29	4.95
Region2	Line2	Product6	5.84	5.65

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:

1. Select **Tools** ⇒ **Filters**. The Filters dialog box appears.
2. Click **New**. The Filter Properties dialog box appears.
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. For the **Filter by** option, select **Dependent variables**.
7. In the table, specify the following values for each column:
 - a. In the **Category** column, select **Reconciled (Rec.) statistic of fit**.
 - b. In the **Property** column, select **Mean absolute percent error**.
 - c. In the **Condition** column, select **>**.
 - d. In the **Value** column, type **3** and press ENTER.
8. Click . A new row appears in the Filter Properties dialog box. SAS Forecast Studio or SAS Forecast Studio for Desktop automatically joins the filter conditions by AND.
9. In the new row, specify the following values for each column:
 - a. In the **Category** column, select **Series property**.
 - b. In the **Property** column, select **Minimum value**.
 - c. In the **Condition** column, select **<**.
 - d. In the **Value** column, type **300** and press ENTER.

The Filter Properties dialog box should look like the following display:


Figure 10.4 Completed Filter Properties Dialog Box for Filter Example



10. Click **OK** to add the filter to the list in the Filters dialog box.
11. 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



REGION...	PRODUC...	PRODUC...	MAPE /	Rec. MAPE
Region2	Line3	Product11	3.97	3.96
Region1	Line1	Product3	4.48	5.15
Region2	Line2	Product6	5.84	5.65

Example 3: Creating a Filter for Missing Values


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

To create this filter:

1. Select **Tools** ⇒ **Filters**. The Filters dialog box appears.
2. Click **New**. The Filter Properties dialog box appears.
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. For the **Filter by** option, select **Dependent variables**.
7. In the table, specify the following values for each column:
 - a. In the **Category** column, select **Reconciled (Rec.) statistics of fit**.
 - b. In the **Property** column, select **Mean absolute percent error**.
 - c. In the **Condition** column, select **=**.
 - d. In the **Value** column, type **.** and press ENTER.

The Filter Properties dialog box should look like the following display:

Figure 10.6 Completed Filter Properties Dialog Box for Filter Example



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

Example 4: Create a Filter for a Combined Model

In your project, you have several combined models. In this example, you want to know which combined models have a Mean Absolute Percent Error (MAPE) greater than 3.


To create this filter:

1. Select **Tools** ⇒ **Filters**. The Filters dialog box appears.
2. Click **New**. The Filter Properties dialog box appears.
3. In the **Name** box, type **CombinedMAPE3**.
4. In the **Description** box, type **Filter to identify which series in the combined models have a MAPE greater than 3**.
5. From the **Hierarchy level** drop-down list, select **All**.
6. For the **Filter by** option, select **Dependent variables**.
7. In the table, specify the following values for each column:
 - a. In the **Category** column, select **Model property**.
 - b. In the **Property** column, select **Model family**.
 - c. In the **Condition** column, select **=**.
 - d. In the **Value** column, select **COMBINED**.

8. Click . A new row appears in the Filter Properties dialog box. SAS Forecast Studio or SAS Forecast Studio for Desktop automatically joins the filter conditions by AND.
9. In the new row, specify the following values for each column:
 - a. In the **Category** column, select **Reconciled (Rec.) statistic of fit**.
 - b. In the **Property** column, select **Mean absolute percent error**.
 - c. In the **Condition** column, select **>**.
 - d. In the **Value** column, type **3** and press ENTER.

The Filter Properties dialog box should look like the following display:

Figure 10.7 Completed Filter Properties Dialog Box for Filter Example



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

Example 5: Create a Filter for an Input Variable

In this example, you want to create a filter using a report variable from the project.


To create this filter:

1. Select **Tools** ⇨ **Filters**. The Filters dialog box appears.
2. Click **New**. The Filter Properties dialog box appears.
3. In the **Name** box, type **Cost**.
4. In the **Description** box, type **Filter to identify which series have a cost greater than 30**.
5. From the **Hierarchy level** drop-down list, select **All**.

6. For the **Filter by** option, select **Input or Report variables**.
7. From the drop-down lists, select **Report only variable**, **=**, and **cost**.
8. Click  to add the filter criteria.
9. In the table, specify the following values for each column:
 - a. In the **Property** column, select **Maximum value**.
 - b. In the **Condition** column, select **>**.
 - c. In the **Value** column, type **30**.

The Filter Properties dialog box should look like the following display:

Figure 10.8 Completed Filter Properties Dialog Box for Filter Example



10. Click **OK** to add the filter to the list in the Filters dialog box.
11. 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 **Cost** to view the series that meet these filter criteria.

Figure 10.9 Results in the Table View After Applying the Cost Filter

Hierarchy View Table View Failed Forecasts					
Statistics of fit calculated over an in-sample range.					
Filter: Cost					
(16 of 26 series)					
REGION...	PRODUC...	PRODUC...	MAPE /	Rec. MAPE	MAX
*	*	*	1.29	1.29	36.685294
Region2	*	*	1.89	2.10	31.131250
Region3	*	*	2.13	2.82	46.366667
Region3	Line4	*	2.42	3.02	50.937500
Region2	Line3	*	2.89	2.66	40.112500
Region1	*	*	3.18	3.08	32.133333
Region1	Line1	*	3.18	3.08	32.133333
Region3	Line5	*	3.19	3.62	37.225000
Region3	Line4	Product13	4.05	5.47	55.550000
Region3	Line5	Product17	4.07	4.17	36.650000
Region3	Line5	Product16	4.26	4.90	37.800000
Region2	Line3	Product9	4.52	4.32	78.000000
Region3	Line4	Product12	4.54	4.86	66.900000
Region1	Line1	Product2	4.64	4.93	52.350000
Region3	Line4	Product15	4.81	4.30	54.700000
Region2	Line2	Product4	5.29	4.95	30.900000

Chapter 11

Running Reports

Overview of Reports	91
Sample Reports in SAS Forecast Studio	92
Create a Report	92
Generating a Report	92
Run a Report	92
Example: Generating a Report on the Final Forecasts for a Selected Node	93
View Your Report History	95
Save the Output from a Report	96

Overview of Reports

Note: Reports are available only if you are running SAS Forecast Studio. These reports are not available if you are running SAS Forecast Studio for Desktop. For more information, see [Chapter 2, “About SAS Forecast Studio for Desktop,”](#) on page 7.

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, XML, or other ODS output. 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 is shipped with a collection of sample reports. When deployed, these reports are visible from the Reports and Stored Processes dialog box. For a description of each of the sample reports, see [Appendix 7, “Sample Reports in SAS Forecast Studio,”](#) on page 281.

In addition to these sample reports, you might want to create your own reports. For more information, see [“Create a Report”](#) on page 92.

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

SAS Forecast Studio is shipped with sample reports, and 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 the *SAS Forecast Server: Administrator's Guide*.

Generating a Report

Run a Report

To run a report for the selected series:

1. Select **Tools** ⇒ **Reports and Stored Processes**. The Reports and Stored Processes dialog box appears.
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 on which the report was created.
3. Select the format and style of the output.
4. (Optional) Specify any ODS options for the selected format. For example, if you select CSV as the format of the output, you might want to specify TAB as the delimiter. In the **Options** box, enter **options (Delimiter='09'x)**.

Note: Any special characters that must be encoded due to the SAS Macro language (such as TAB) must be entered using their code rather than the keyboard character.

For more information about the options that are available for each format, see *SAS Output Delivery System: User's Guide*.

5. (Optional) Select the execution server for the report.

Note: The availability of this option depends on how your administrator has configured your site. For more information, see *SAS Forecast Server: Administrator's Guide*.

6. Specify whether to view the output immediately or to save the output to a file for later viewing.
7. 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 appears 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 or is saved to the output location that you specified.


Note: If the report that 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:

1. Select **Tools** ⇒ **Reports and Stored Processes**. The Reports and Stored Processes dialog box appears.
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



SAS Output - Internet Explorer provided by SAS Institute

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
JAN03	6680.61	6450.36	6910.87	117.477	6680.61	6450.36	6910.87	117.477	0
FEB03	6954.81	6724.25	7185.38	117.638	6954.81	6724.25	7185.38	117.638	0
MAR03	7015.61	6784.73	7246.49	117.798	7015.61	6784.73	7246.49	117.798	0
APR03	7058.41	6827.22	7289.61	117.959	7058.41	6827.22	7289.61	117.959	0
MAY03	7270.82	7039.30	7502.33	118.120	7270.82	7039.30	7502.33	118.120	0
JUN03	7384.22	7152.39	7616.05	118.281	7384.22	7152.39	7616.05	118.281	0
JUL03	7094.63	6862.48	7326.77	118.442	7094.63	6862.48	7326.77	118.442	0
AUG03	7009.24	6776.78	7241.69	118.603	7009.24	6776.78	7241.69	118.603	0
SEP03	7042.85	6810.07	7275.62	118.765	7042.85	6810.07	7275.62	118.765	0
OCT03	6730.06	6496.97	6963.15	118.926	6730.06	6496.97	6963.15	118.926	0
NOV03	6567.48	6334.07	6800.88	119.088	6567.48	6334.07	6800.88	119.088	0
DEC03	6462.69	6228.97	6696.42	119.249	6462.69	6228.97	6696.42	119.249	0


View Your Report History

SAS Forecast Studio logs when you try to run a report and if the report ran successfully. After running several reports, you might want to compare the results from different reports or compare the results of the same report on different days.

To view your report history:

1. Select **Tools** ⇒ **Reports and Stored Processes**. The Reports and Stored Processes dialog box appears.
2. Click the **History** tab.

Here is an example of the contents in the **History** tab:



For any reports that failed, you can view the log to determine why the report failed. For any reports that ran successfully, you can view the log, view the output, or save the output.

Save the Output from a Report

You might want to save the output for a report for future reference or to share the results with people in your organization who do not use SAS Forecast Studio. You can save the report output when you initially run the report. (For more information about how to run a report, see [“Run a Report” on page 92](#).) You can also save the output from the **History** tab in the Reports and Stored Processes dialog box.

To save the output for a report that you previously ran:

1. Select **Tools** ⇒ **Reports and Stored Processes**. The Reports and Stored Processes dialog box appears.
2. Click the **History** tab.
3. Select the report from the table and click **Save Output**. The Save dialog box appears.
4. Specify the location and filename for the output, and click **Save**.

If the output format uses multiple files (for example, HTML output generates several files), all files are saved to the same location. The filenames of any dependent files (such as image files) are assigned by the server. Do not change these filenames to ensure the referential integrity of the output.

Because of the number of output files, you might want to save your reports to an empty directory, so the files are associated with the report output.

Part 4

Improving Your Forecasts

<i>Chapter 12</i>	
Working with Events	99
<i>Chapter 13</i>	
Working with Overrides	117
<i>Chapter 14</i>	
Understanding Hierarchy Reconciliation	129
<i>Chapter 15</i>	
Viewing the Results of a Cross-Series Analysis	163
<i>Chapter 16</i>	
Working with Models	169
<i>Chapter 17</i>	
Creating User-Defined Models	189
<i>Chapter 18</i>	
Performing Scenario Analysis	213

Chapter 12

Working with Events

Understanding Events	99
What Is an Event?	99
Event Types	100
Event Occurrence	101
Event Recurrence	101
Creating New Events	101
Create an Event	101
Example 1: Creating an Event for Annual Promotion	102
Example 2: Creating an Event for a Holiday	104
Understanding Combination Events	105
Combine Events	105
Example 3: Creating an Event for Third-Quarter Sales	106
Edit an Event	109
Copy an Event	109
Sharing Events	109
Add Events from a SAS Data Set	109
Save Events to a SAS Data Set	110
How Events Are Used in Model Generation	110
Understanding How Events Are Interpreted	110
Specify the Event Status	111
Adding an Event to a Custom Model	111
View the Effect of Events	115

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 and SAS Forecast Studio for Desktop 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 and SAS Forecast Studio for Desktop create an indicator variable (also called a dummy variable). This variable indicates the occurrence of the event at any time period. This indicator variable is used as a regressor variable for time series modeling and forecasting.

For example, 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 or SAS Forecast Studio for Desktop 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 and SAS Forecast Studio for Desktop must have data that explains how retail sales were impacted by previous Christmas events. Without this data, SAS Forecast Studio or SAS Forecast Studio for Desktop cannot predict the effect of the next Christmas event on retail sales.

SAS Forecast Server and SAS Forecasting for Desktop are shipped with a list of defined events. However, you can customize the list of events that appear in SAS Forecast Studio or SAS Forecasting for Desktop. For more information about how to create custom events, see the *SAS Forecast Server: Administrator's Guide*.


The following components can define an event:

- the event type
- the occurrence
- the recurrence

Event Types

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:

1. Select **Project** ⇒ **Event Repository**. The Event Repository dialog box appears.
2. Click **New**. The Event Properties dialog box appears.
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 99](#).
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 or SAS Forecast Studio for Desktop generates. For more information, see [“How Events Are Used in Model Generation” on page 110](#).
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.

- 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.
- To specify the recurrence, click **Edit** next to the **Recurrence** box. From the Event Recurrence dialog box, you can specify how frequently the event recurs 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.

- 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 (or SAS Forecast Studio for Desktop) to try to include this event in model generation. The recurrence option enables you to create an event and specify how frequently this event occurs. The following example shows how to create a pulse event for this annual promotion.

To create an event for the June promotion:

- Select **Project** ⇒ **Event Repository**. The Event Repository dialog box appears.
- Click **New**. The Event Properties dialog box appears.
- In the **Name** box, type **Annual_Sale**.
- In the **Description** box, type **Event for annual sale in June**.
- To add an occurrence, click **Add**. The Select Occurrences dialog box appears.
- 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.

- To specify the recurrence, click **Edit**. The Event Recurrence dialog box appears.

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 (or SAS Forecast Studio for Desktop) will try to use this event in all system-generated models.

The Event Properties dialog box looks similar to the following display:

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:

1. Select **Project** ⇒ **Event Repository**. The Event Repository dialog box appears.
2. Click **New**. The Event Properties dialog box appears.
3. In the **Name** box, type **Holiday_Sale**.
4. To add an occurrence, click **Add**. The Select Occurrences dialog box appears.
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 appears.
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.

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 display:

Figure 12.4 Event Properties for a Holiday Event




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

Understanding Combination Events

Combine Events

To combine events:

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

Figure 12.5 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 which 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:


1. Select **Project** ⇒ **Event Repository**. The Event Repository dialog box appears.

Figure 12.6 Events in the Events Repository

2. Click **Combine**. The Select Events dialog box appears.
3. Select the check boxes for the FALL_SALES and SCHOOL_SALES events and click **OK**. The Combination Event Properties dialog box appears.


Figure 12.7 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** drop-down list, select **Add the values**.

Figure 12.8 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. The Event Repository Changed dialog box appears, so you can quickly reforecast the project.

Figure 12.9 Event Repository Changed Dialog Box

Edit an Event

To edit an event:

1. Select **Project** ⇒ **Event Repository**. The Event Repository dialog box appears.
2. Select the event to modify and click **Edit**. The Event Properties dialog box appears. 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:

1. Select **Project** ⇒ **Event Repository**. The Event Repository dialog box appears.
2. Select an event and click **Copy**. The Event Properties dialog box appears. By default, the copied event is named **event-nameCOPYn** 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 or SAS Forecast Studio for Desktop. These event repositories can be created by using the HPFEVENTS procedure. When writing your SAS code, do not specify the CONDENSE option in the EVENTDATA statement. If you do, then the output data set does not include all of the variables that are required by SAS Forecast Studio or SAS Forecast Studio for Desktop. For more information about the HPFEVENTS procedure, see the *SAS Forecast Server Procedures: User's Guide*.

After the event repository data set is created, you can import these events into SAS Forecast Studio or SAS Forecast Studio for Desktop.

To import an event:

1. Select **Project** ⇒ **Event Repository**. The Event Repository dialog box appears.
2. Click **Add events from a data set**. The Import Events dialog box appears.
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 or SAS Forecast Studio for Desktop users at your site. This SAS data set is saved to an events repository. Other 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:

1. Select **Project** ⇒ **Event Repository**. The Event Repository dialog box appears.
2. Click **Save events to a data set**. The Save As dialog box appears.
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-name2*) 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 or SAS Forecast Studio for Desktop 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 109](#).

How Events Are Used in Model Generation

Understanding How Events Are Interpreted

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

Do not use

does not include the event in the model.

Try to use

specifies that the events be included in the model as long as the parameters of the events are significant and the increment of the value of the criterion exceeds a specified threshold.

Use if significant

specifies that the events be included in the model as long as the parameters of the events are significant.

Force use

specifies that the events be included in the model as long as the model does not fail to be diagnosed.

This status is used for all the models that SAS Forecast Studio or SAS Forecast Studio for Desktop 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, see [“Adding an Event to a Custom Model” on page 111](#).

Specify the Event Status

To specify the status of an event for all the models that SAS Forecast Studio or SAS Forecast Studio for Desktop generates in the project:

1. Select **Project** ⇒ **Event Repository**. The Event Repository dialog box appears.
2. Select the event (or events) 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 Events Are Interpreted” on page 110](#). For more information about these options, see [“Understanding How Events Are Interpreted” on page 110](#).

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:

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 each event 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 12.10 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 area. The following example shows a model with the WINTER_SALES event with simple and seasonal differencing.

Figure 12.11 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:


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 each event 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 area. The following example shows a model with the WINTER_SALES event that includes a lag of one period.

Figure 12.12 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:

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 each event 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 area. 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 12.13 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:

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 each event that you want to include in the model.

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

Figure 12.14 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:

1. Open the Forecasting View.
2. Select **View** ⇒ **Edit Forecasting View Properties**. The Forecasting View Properties dialog box appears.
3. For the forecast plot and data table, select the **Effect of events** check box.

Chapter 13

Working with Overrides

Understanding Overrides	117
What Is an Override?	117
When Can You Specify an Override?	118
How Are Overrides Reconciled in a Hierarchy?	118
Adding Overrides	120
About Adding Overrides	120
Adding Overrides in the Forecasting View	120
Using the Overrides Calculator	122
Specify Overrides for a Branch or Filter	124
Lock and Unlock Overrides	125
Add a Note to a Series	125
Remove Overrides	126
Working with Override Conflicts	126
What Is an Override Conflict?	126
Resolve Override Conflicts	128

Understanding Overrides

What Is an Override?

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

- If you have forecast your data hierarchically, then the override is an adjustment that is done with respect to the reconciled statistical forecast.
- If you have not forecast 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 and SAS Forecast Studio for Desktop:

- **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 or SAS Forecast Studio for Desktop 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 or SAS Forecast Studio for Desktop 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 125](#).

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 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, see [“Specify the Number of Periods to Forecast” on page 75](#).

How Are Overrides Reconciled in a Hierarchy?

How SAS Forecast Studio and SAS Forecast Studio for Desktop process overrides in a hierarchy depends on whether any override conflicts are detected and how SAS Forecast Studio and SAS Forecast Studio for Desktop resolve 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, see [“Understanding Reconciliation” on page 129](#).

If you require that conflicts be resolved before reconciliation, then SAS Forecast Studio and SAS Forecast Studio for Desktop use 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 and SAS Forecast Studio for Desktop check for conflicts in the locked overrides across the hierarchy.

- If no conflicts are detected, SAS Forecast Studio and SAS Forecast Studio for Desktop perform 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 or SAS Forecast Studio for Desktop 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 132](#).

Note: Unreconciled nodes occur only if you restrict the direction of the reconciliation. For more information, see [“Specify the Advanced Reconciliation Options” on page 132](#).

- If conflicts are detected, the Override Conflicts dialog box appears. You view and resolve any conflicts.

After all the conflicts are resolved, SAS Forecast Studio and SAS Forecast Studio for Desktop 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 or SAS Forecast Studio for Desktop 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 132](#).

Note: Unreconciled nodes occur only if you restrict the direction of the reconciliation. For more information, see [“Specify the Advanced Reconciliation Options” on page 132](#).

If you choose to ignore the conflicts, then SAS Forecast Studio and SAS Forecast Studio for Desktop use the following process:

- a. When you add an override, a "Reconciliation needed" message appears at the top of the Forecasting View.
- b. When you click **Reconcile**, SAS Forecast Studio and SAS Forecast Studio for Desktop try 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 and SAS Forecast Studio for Desktop 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 132](#).

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, see [“Specify the Advanced Reconciliation Options” on page 132](#).

Adding Overrides

About Adding Overrides

You can add an override in the following ways:

- from the data table of the Forecasting View. To override a single value, enter the override value in the Overrides row of the data table. For more information, see [“Adding Overrides in the Forecasting View” on page 120](#).

To apply overrides to several time periods, use the Override Calculator. For more information, see [“Using the Overrides Calculator” on page 122](#).

- by selecting a filter or branch of the hierarchy in the Branch/Filter Overrides dialog box. For more information, see [“Specify Overrides for a Branch or Filter” on page 124](#).

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 75](#).

For each override, you can specify whether the override is locked or unlocked. This locking determines how SAS Forecast Studio or SAS Forecast Studio for Desktop treats the override during the reconciliation process. For more information, see [“Lock and Unlock Overrides” on page 125](#).

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 125](#).

Adding Overrides in the Forecasting View

Entering Overrides in the Data Table

To add an override from the data table in the Forecasting View:

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 entering the new value.

Note: If you cannot add an override to a future value, see [“When Can You Specify an Override?” on page 118](#).

Copying and Pasting Overrides

You can update overrides by using the copy and paste functionality.

- You can copy overrides from one time period to another time period in the same series.
- You can copy overrides from one time period to another time period in different series.
- You can also copy override values from a Microsoft Excel worksheet. Use tab (t) as the separator between override values.

The copied values are pasted irrespective of time IDs, starting from the selected cell. If the copied values contain an invalid value, then the invalid value is dropped and only the valid values are pasted.

Here are some examples of how the copy and paste functionality works.

- This first example shows what happens if your data includes an invalid value.

First, you copy the following values from Microsoft Excel:

Table 13.1 Example 1: Override Values in Excel

Column 1	Column 2	Column 3	Column 4
1250	1350	abc	1650

The copied values are copied to the Clipboard. In SAS Forecast Studio or SAS Forecast Studio for Desktop, you select the cell in the data table where you want to paste the override from Column 1, and then press **CTRL+V**. The following values appear in the data table:

Table 13.2 Example 1: Updated Override Values

Pasted Value 1	Pasted Value 2	Pasted Value 3	Pasted Value 4
1250	1350		1650

The value for Column 3 (abc) is an invalid value, so it is not pasted into the data table. Instead, this value is dropped.

- This second example shows how blank values are retained when you copy values from Excel to SAS Forecast Studio or SAS Forecast Studio for Desktop.

First, you copy the following values from Microsoft Excel:

Table 13.3 Example 2: Override Values in Excel

Column 1	Column 2	Column 3	Column 4
1250			1650

The copied values are placed on the Windows clipboard. In SAS Forecast Studio or SAS Forecast Studio for Desktop, you select the cell in the data table where you want to paste the override from Column 1, and then press **CTRL+V**. The following values appear in the data table:

Table 13.4 Example 2: Updated Override Values

Pasted Value 1	Pasted Value 2	Pasted Value 3	Pasted Value 4
1250			1650

Because you do not have to specify an override value for each value in the series, blank values are valid values for overrides in SAS Forecast Studio and SAS Forecast Studio for Desktop. As a result, the blank values from Excel appear in the data table.

Using the Overrides Calculator

Add an Override By Using the Override Calculator

To create an override by using the Override Calculator:

1. Select **Series** ⇒ **Override Calculator**. The Override Calculator appears.
If the Override Calculator is not available, see [“When Can You Specify an Override?” on page 118](#).
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

To create overrides for the Feb2003 and Mar2003 time periods that are 15% less than the value of the final forecast:

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 appears.
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 display:

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:

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 appears.
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 display:

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

Active series											
		Dec2002	Jan2003	Feb2003	Mar2003	Apr2003	May2003	Jun2003	Jul2003	Aug2003	Sep2003
Historical Data	D0	1486.000000									
Forecast Model	E5	1485.442688	1526.634067	1532.040381	1534.650914	1535.893216	1553.147548	1536.711729	1536.797485	1536.805348	1536.7754
Reconciled Forecast	P6	1442.073448	1511.758099	1586.895879	1595.632198	1554.749545	1566.250825	1608.897974	1565.409997	1519.783401	1553.6555
Override				1246.568640	1253.431360						
Lock Override				☒	☒						
Final Forecast			1416.790236	1586.895879	1595.632198	1554.749545	1566.250825	1608.897974	1565.409997	1519.783401	1553.6555

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

Time Period	Proportion of Total Forecast	Override Value
Feb2003	$\frac{1586.895879}{3182.527859} = .498627$	$.498627 * 2500 = 1246.568540$
Mar2003	$\frac{1595.632198}{3182.527859} = .501373$	$.501373 * 2500 = 1253.431360$

Specify Overrides for a Branch or Filter

To apply overrides to a series in a branch of the hierarchy or to a series that meets the criteria for a project filter:

1. Select **Project** ⇒ **Branch/Filter Overrides**. The Branch or Filter Overrides dialog box appears.
2. Select the location of the series. You can choose from these options:
 - **Select from filter**. This option is selected by default, and the results of the project's default filter are displayed in the Branch/Filter Overrides dialog box.

You can select a different filter or create a new filter. After you have selected the filter, select the series where you want to apply the overrides.

- **Select from branch.** If you select this option, then select the branch of the hierarchy that contains the series where you want to apply the overrides. All series in the selected branch are selected. This option is available only if you forecast hierarchically.
3. Select the cells that you want to override and specify the new override value. You can apply individual overrides by entering a value directly into the table. You can also select multiple cells and specify override values by using the Override Calculator.

Lock and Unlock Overrides

To lock an individual override:

1. 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 or SAS Forecast Studio for Desktop 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 forecast value, you might want to document why you specified this override. In SAS Forecast Studio and SAS Forecast Studio for Desktop, 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:

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, enter 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**.



Note: After you remove an override, you must manually re-enter that override. There is no other way to restore any overrides that you previously specified.

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 another locked override at the same time interval.

The following figure shows an example of this 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:

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:
 - Deselect 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 118](#).

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

Understanding Reconciliation	129
Reconcile the Hierarchy	131
Specify the Advanced Reconciliation Options	132
Working with Reconciliation Failures	132
About the Reconciliation Failure Report	132
How to View Unreconciled Nodes	132
Example: Resolving Unreconciled Nodes	133
Examples of Reconciling a Hierarchy with Overrides	136
Overview of the Examples	136
Example 1: Top-Down Method of Reconciliation with No Conflicts	137
Example 2: Top Down with Implicit Bottom-Up Method of Reconciliation	140
Example 3: Unreconciled Node in the Hierarchy	143
Example 4: Bottom-Up Method of Reconciliation with Locked and Unlocked Overrides in the Lowest Level	146
Example 5: Implicit Top-Down Reconciliation Process in the Bottom-Up Method of Reconciliation	148
Example 6: Implied Limits Have Precedence over an Unlocked Override	152
Example 7: Middle-Out Method of Reconciliation with No Conflicts	156
Example 8: Middle-Out Method of Reconciliation with an Unreconciled Node	159

Understanding Reconciliation

If you add an override to data that is forecast hierarchically, then you must reconcile the hierarchy before SAS Forecast Studio or SAS Forecast Studio for Desktop can calculate the final forecast. SAS Forecast Studio or SAS Forecast Studio for Desktop 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 and SAS Forecast Studio for Desktop reconcile 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, see [“Specify the Advanced Reconciliation Options” on page 132](#).

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

How Conflicts Are Resolved	Reconciliation Direction	Result
Require that conflicts be resolved before reconciliation	No restriction	When reconciling the hierarchy, SAS Forecast Studio and SAS Forecast Studio for Desktop determine 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 and SAS Forecast Studio for Desktop reconcile the hierarchy. If you select this combination of options and there are no override conflicts (or those conflicts are resolved), no failure occurs.
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 and SAS Forecast Studio for Desktop determine whether there are any conflicts. However, because any conflicts should be ignored, the Override Conflicts dialog box does not open if conflicts are detected. If a failure occurs, the Reconciliation Failure dialog box appears. 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 Conflicts Are Resolved	Reconciliation Direction	Result
Require that conflicts be resolved before reconciliation	Restrict direction	Before reconciling the hierarchy, SAS Forecast Studio and SAS Forecast Studio for Desktop determine 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 and SAS Forecast Studio for Desktop reconcile 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 appears and shows the reconciliation failures.
Ignore conflicts	Restrict direction	You might want to select these options if you want SAS Forecast Studio or SAS Forecast Studio for Desktop to reconcile the hierarchy with restrictions but without intervention. Before reconciling the hierarchy, SAS Forecast Studio and SAS Forecast Studio for Desktop determine whether there are any conflicts. However, because any conflicts should be ignored, the Override Conflicts dialog box does not open if conflicts are detected. If a failure occurs, the Reconciliation Failure dialog box appears. 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:

- Click  in the toolbar.
- Select **Project** ⇒ **Reconcile Hierarchy**.
- 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:

1. Select **Project** ⇒ **Hierarchy and Variable Settings**. The Hierarchy and Variable Settings dialog box appears.
2. Click **Advanced**. The Advanced Reconciliation Settings dialog box appears.
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 Forecast Server Procedures: 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 for 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 direction of the reconciliation, see [“Specify the Advanced Reconciliation Options” on page 132](#).

How to View Unreconciled Nodes

To view and resolve any reconciliation failures:

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 appears.
 3. Resolve any conflicts in the Overrides Conflict dialog box and click **Reconcile**. For more information, see [“Resolve Override Conflicts” on page 128](#).

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 and SAS Forecast Studio for Desktop generate 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:


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




2. Click **Resolve Conflicts** to determine how many of these unreconciled nodes are due to override conflicts. The Override Conflicts dialog box appears. 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 and SAS Forecast Studio for Desktop reconcile 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 **Project** ⇨ **Hierarchy and Variable Settings**. The Hierarchy and Variable Settings dialog box appears.
6. Click **Advanced**. The Advanced Reconciliation Settings dialog box appears.

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 and SAS Forecast Studio for Desktop reconcile 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 and SAS Forecast Studio for Desktop reconcile 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

Level in the Hierarchy	Possible Values
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 Dec2012.
- **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 75](#).
- 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 132](#).

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

Purpose of the Example

This example shows the following:

- how SAS Forecast Studio and SAS Forecast Studio for Desktop reconcile a hierarchy using the top-down method of reconciliation
- how SAS Forecast Studio and SAS Forecast Studio for Desktop honor 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 Dec2012 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. For more information, see [“Overview of the Examples” on page 136](#).

Explanation of the Reconciliation Process

The following figure shows how SAS Forecast Studio and SAS Forecast Studio for Desktop reconcile the hierarchy for the Dec2012 time interval:



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 and SAS Forecast Studio for Desktop start 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 and SAS Forecast Studio for Desktop honor 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 and SAS Forecast Studio for Desktop do not honor the value of this unlocked override. Instead, SAS Forecast Studio and SAS Forecast Studio for Desktop use 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 and SAS Forecast Studio for Desktop honor this value, so the final forecast for the Trains node is 80. Because the final forecast for the Toys node is 100, the final forecast for 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 and SAS Forecast Studio for Desktop use 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. 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 $.25 \times 100 = 25$). The final forecast for the Pants node is 75% of the final forecast for the Clothes node (or $.75 \times 100 = 75$).

Example 2: Top Down with Implicit Bottom-Up Method of Reconciliation**Purpose of the Example**

This example shows the following:

- how SAS Forecast Studio and SAS Forecast Studio for Desktop reconcile 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 and SAS Forecast Studio for Desktop use 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 Dec2012 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. For more information, see [“Overview of the Examples” on page 136](#).

Explanation of the Reconciliation Process

The following figure shows how SAS Forecast Studio and SAS Forecast Studio for Desktop reconcile the hierarchy for the Dec2012 time interval:



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 and SAS Forecast Studio for Desktop start 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 and SAS Forecast Studio for Desktop honor 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 and SAS Forecast Studio for Desktop do not honor the value of this unlocked override. Instead, SAS Forecast Studio and SAS Forecast Studio for Desktop use 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 and SAS Forecast Studio for Desktop 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 and SAS Forecast Studio for Desktop set the final forecast for 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 for the Southeast node is 120, the final forecast for 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 and SAS Forecast Studio for Desktop honor 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 and SAS Forecast Studio for Desktop use 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. 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 $.25 \times 40 = 10$). The final forecast for the Pants node is 75% of the final forecast for the Clothes node (or $.75 \times 40 = 30$).

Example 3: Unreconciled Node in the Hierarchy

Purpose of the Example

This example shows the following:

- how SAS Forecast Studio and SAS Forecast Studio for Desktop try 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 Dec2012 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. For more information, see [“Overview of the Examples” on page 136](#).

Explanation of the Reconciliation Process

The following figure shows how SAS Forecast Studio and SAS Forecast Studio for Desktop reconcile the hierarchy for the Dec2012 time interval:



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 and SAS Forecast Studio for Desktop start 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 and SAS Forecast Studio for Desktop honor 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 and SAS Forecast Studio for Desktop do not honor the value of this unlocked override. Instead, SAS Forecast Studio and SAS Forecast Studio for Desktop use 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 and SAS Forecast Studio for Desktop do 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.
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 and SAS Forecast Studio for Desktop honor 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 and SAS Forecast Studio for Desktop use 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 $.25 \times 60 = 15$). The final forecast for the Pants node is 75% of the final forecast for the Clothes node (or $.75 \times 60 = 45$).

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 and SAS Forecast Studio for Desktop reconcile 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 restriction on the reconciliation direction


You also specified the following overrides for the Dec2012 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. For more information, see [“Overview of the Examples” on page 136](#).

Explanation of the Reconciliation Process

The following figure shows how SAS Forecast Studio and SAS Forecast Studio for Desktop reconcile the hierarchy for the Dec2012 time interval:



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 and SAS Forecast Studio for Desktop start 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 and SAS Forecast Studio for Desktop honor 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 and SAS Forecast Studio for Desktop honor 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 and SAS Forecast Studio for Desktop use the final forecasts for 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). 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 and SAS Forecast Studio for Desktop do 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 and SAS Forecast Studio for Desktop try 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 and SAS Forecast Studio for Desktop use 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 Dec2012 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. For more information, see [“Overview of the Examples” on page 136](#).

Explanation of the Reconciliation Process

The following figure shows how SAS Forecast Studio and SAS Forecast Studio for Desktop reconcile the hierarchy for the Dec2012 time interval:



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 and SAS Forecast Studio for Desktop start 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 and SAS Forecast Studio for Desktop honor 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 and SAS Forecast Studio for Desktop use the final forecasts for 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). 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 for 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 and SAS Forecast Studio for Desktop perform 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 and SAS Forecast Studio for Desktop are now performing an implicit top-down reconciliation, the final forecast of 170 that was calculated by the bottom-up method of reconciliation is replaced with the 200 that you specified as the locked override.

Level in the Hierarchy	Reconciliation Process for the Nodes in This Level
Product Category	You specified an unlocked override of 60 for the Toys node. SAS Forecast Studio and SAS Forecast Studio for Desktop use 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 and SAS Forecast Studio for Desktop divide 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 and SAS Forecast Studio for Desktop honor 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 and SAS Forecast Studio for Desktop use 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. 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 $.25 \times 100 = 25$). The final forecast for the Pants node is 75% of the final forecast for the Clothes node (or $.75 \times 100 = 75$).

Example 6: Implied Limits Have Precedence over an Unlocked Override

Purpose of the Example

This example shows the following:

- how SAS Forecast Studio and SAS Forecast Studio for Desktop try 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 and SAS Forecast Studio for Desktop use 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 Dec2012 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. For more information, see [“Overview of the Examples” on page 136](#).

Explanation of the Reconciliation Process

The following figure shows how SAS Forecast Studio and SAS Forecast Studio for Desktop reconcile the hierarchy for the Dec2012 time interval:



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 and SAS Forecast Studio for Desktop start 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 and SAS Forecast Studio for Desktop honor 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 and SAS Forecast Studio for Desktop use the final forecasts for 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). 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. 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 and SAS Forecast Studio for Desktop perform 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 and SAS Forecast Studio for Desktop are now performing an implicit top-down reconciliation, the final forecast of 170 that was calculated by the bottom-up method of reconciliation is replaced with the 100 that you specified as the locked override.

Level in the Hierarchy	Reconciliation Process for the Nodes in This Level
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 $.80 \times 100 = 80$), and the final forecast for the Clothes node is 20 (or $.20 \times 100 = 20$).
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 and SAS Forecast Studio for Desktop honor 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 and SAS Forecast Studio for Desktop use 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. 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 $.25 \times 20 = 5$). The final forecast for the Pants node is 75% of the final forecast for the Clothes node (or $.75 \times 20 = 15$).

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

Purpose of This Example

This example shows the following:

- how SAS Forecast Studio and SAS Forecast Studio for Desktop reconcile 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 Dec2012 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. For more information, see [“Overview of the Examples” on page 136](#).

Explanation of the Reconciliation Process

The following figure shows how SAS Forecast Studio and SAS Forecast Studio for Desktop reconcile the hierarchy for the Dec2012 time interval:



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 and SAS Forecast Studio for Desktop start 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 and SAS Forecast Studio for Desktop are 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 and SAS Forecast Studio for Desktop calculate 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.
Product Line	SAS Forecast Studio and SAS Forecast Studio for Desktop calculate 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 and SAS Forecast Studio for Desktop honor locked overrides, so the final forecast for the Trains node is 20. Because SAS Forecast Studio and SAS Forecast Studio for Desktop are 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.

Level in the Hierarchy	Reconciliation Process for the Nodes in This Level
Product Line	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 $40 \times .25 = 10$. - 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 $40 \times .75 = 30$. 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 and SAS Forecast Studio for Desktop use 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 $60 \times .333 = 20$. - 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 $60 \times .666 = 40$.

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

Purpose of This Example

This example shows the following:

- how SAS Forecast Studio and SAS Forecast Studio for Desktop reconcile 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 Dec2012 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. For more information, see [“Overview of the Examples” on page 136](#).

Explanation of the Reconciliation Process

The following figure shows how SAS Forecast Studio and SAS Forecast Studio for Desktop reconcile the hierarchy for the Dec2012 time interval:



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 and SAS Forecast Studio for Desktop start 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 and SAS Forecast Studio for Desktop honor this unlocked override. 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 and SAS Forecast Studio for Desktop leave 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 and SAS Forecast Studio for Desktop calculate 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 and SAS Forecast Studio for Desktop use 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.
Product Line	SAS Forecast Studio and SAS Forecast Studio for Desktop calculate 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 and SAS Forecast Studio for Desktop honor 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 and SAS Forecast Studio for Desktop honor 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

Viewing the Results of a Cross-Series Analysis

View the Plots for the Time Series Analysis	163
View the Tables for the Time Series Analysis	163
Transform a Time Series for Analysis	164
Viewing the Results of a Time Series Decomposition	164
About Time Series Decomposition	164
How to View the Time Series Decomposition	166

View the Plots for the Time Series Analysis

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 to or out of a plot.

Note: Not all plots might be available for the current series. 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:

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

View the Tables for the Time Series Analysis

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

Note: Not all tables might be available for the current series.

To open a time series table in the Series View:

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 for Analysis

To improve your understanding of the statistical properties of the time series, you can transform a time series. The time series analysis is then performed on the transformed series.

To transform a time series:

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 and SAS Forecast Studio for Desktop immediately update 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 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. 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	$O_t = TC_t(S_t + I_t - 1)$

Component	Decomposition Method	Formula
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	$SI_t = S_t I_t = \frac{O_t}{TC_t}$
	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	$SI_t = S_t + I_t - 1 = \frac{O_t}{TC_t}$
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	$I_t = \frac{SI_t}{S_t}$
	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	$TCS_t = TC_t S_t = \frac{O_t}{I_t}$
	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	$TCS_t = TC_t S_t$

Component	Decomposition Method	Formula
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	$SA_t = \frac{O_t}{S_t} = TC_t I_t$
	additive	$SA_t = O_t - S_t = TC_t + I_t$
	log-additive	$SA_t = \frac{O_t}{\exp(S_t)} = \exp(TC_t + I_t)$
	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:

$$TC_t = \frac{\sum_{k=-[s/2]}^{[s/2]} \frac{y_{t+k}}{2}}$$

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

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

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 15.1 Decomposition Plots and Tables in Series View



Chapter 16

Working with Models

Overview of Model Repositories	170
What Is a Model Repository?	170
Types of Models in the Model Repository	170
How Models Are Fit	171
Select the Models to Fit to Each Series	173
Specifying the Forecast Model	174
What Is the Forecast Model?	174
Select the Forecast Model	174
Assessing the Accuracy of the Models	175
The Selection Criterion	175
View the Diagnostics for Each Model	175
Using a Holdout Sample	175
Running an End Zero Test	176
Using an Out-of-Sample Region	176
Using Rolling Simulations to Explore Model Performance	177
View the Model Plots	178
View the Model Tables	179
Compare Models	179
Add a Model to a Series	180
Importing and Exporting Models	180
Add Models from a Catalog	180
Save Models to a Catalog	181
Export a Model Selection List	181
Create a New Model Based on a Default Model	181
Edit a Model	182
About Editing Models	182
Edit a Model from the Model Selection List	182
Edit a Model from the Model Repository	183
Copy Models	183
Combining Models	184
What Is a Combined Model?	184
Understanding the Process for Combining Forecasts	184
Generating a Combined Model	185
Combine Models in the Model Selection List	186

Overview of Model Repositories

What Is a Model Repository?


All of the models that are available to a project are stored in a model repository. The following table describes each of the model repositories that are used by SAS Forecast Studio and SAS Forecast Studio for Desktop:

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 <code>projectmodrep.sas7bcat</code> , which is in the home directory for SAS Forecast Studio or SAS Forecast Studio for Desktop. 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 the SAS Forecast Server procedures. You cannot modify the content in this repository.
model repositories for each level in the hierarchy	SAS Forecast Studio and SAS Forecast Studio for Desktop automatically create 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 and SAS Forecast Studio for Desktop identify these models by using specific model names. You cannot use these model names for any user-created models. For more information, see Appendix 2, “Reserved Names for Variables and Models,” on page 249.

Types of Models in the Model Repository

To open the model repository, select **Project** ⇒ **Model Repository**. SAS Forecast Studio and SAS Forecast Studio for Desktop have 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 and SAS Forecast Studio for Desktop support the following types of models:

default models

models that are shipped with SAS Forecast Studio and SAS Forecast Studio for Desktop. 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 or SAS Forecast Studio for Desktop. 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. In the model repository, models from an external list are listed as type Custom (Read-only). 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 Models Are Fit

For a project, you can specify the classes of models that SAS Forecast Studio or SAS Forecast Studio for Desktop should try to fit to the data. These models make up the model selection list for a series. When SAS Forecast Studio and SAS Forecast Studio for

Desktop diagnose a project, SAS Forecast Studio and SAS Forecast Studio for Desktop attempt to fit all the models in the model selection list. To add or remove models from the model selection lists, see “[Select the Models to Fit to Each Series](#)” on page 173.


SAS Forecast Studio and SAS Forecast Studio for Desktop run a series of diagnostics to determine the characteristics of the data (such as seasonality or intermittency) and avoid models that are inappropriate for the data. Here are some examples of when SAS Forecast Studio and SAS Forecast Studio for Desktop avoid a model:

- 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.

Note: If none of the models in the model selection list fit, then SAS Forecast Studio and SAS Forecast Studio for Desktop attempt to use the model selection list that is shipped with SAS Forecast Studio and SAS Forecast Studio for Desktop. This default model selection list consists of exponential smoothing models.

After SAS Forecast Studio and SAS Forecast Studio for Desktop generate the forecasts, you can view the model selection list in the Modeling View.

Figure 16.2 The Model Selection List in the Modeling View



SAS Forecast Studio and SAS Forecast Studio for Desktop use the value of the selection criterion to determine the best-performing model (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 the statistical properties of the time series are not compatible with the model specification, then a model might be rejected. The term **Rejected** appears in the column for the statistic of fit.

Here are examples of when a model is rejected:

- Seasonal models are rejected if the seasonality is less than two.
- Log transform models are rejected if the series is not strictly positive.

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 and SAS Forecast Studio for Desktop diagnose and forecast 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, see [“Set the Diagnose Properties for Each Series” on page 74](#).

Select the Models to Fit to Each Series

You must select at least one model for SAS Forecast Studio or SAS Forecast Studio for Desktop to use when generating the model selection lists. By default, SAS Forecast Studio and SAS Forecast Studio for Desktop create an ARIMA model and an exponential smoothing model.

To select the models to fit to each series:

1. Select **Project** ⇒ **Forecasting Settings**. The Forecasting Settings dialog box appears.
2. In the selection pane, click **Model Generation**.
3. Select the system-generated models to create. By default, SAS Forecast Studio and SAS Forecast Studio for Desktop create 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 appears, 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 and SAS Forecast Studio for Desktop generate 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 and SAS Forecast Studio for Desktop use to generate a numerical prediction of the future values for the time series. By default, SAS Forecast Studio and SAS Forecast Studio for Desktop choose 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 and specify the forecast model. When you reforecast the project, SAS Forecast Studio and SAS Forecast Studio for Desktop select the forecast model based on the current forecast model and how you specified to reforecast the data. In SAS Forecast Studio and SAS Forecast Studio for Desktop, you can specify where to start the process for generating a forecast. For more information about each of these steps, see [“Process for Generating a Forecast” on page 73](#).

- You select a system-generated model as the forecast model. Then reforecast the project and have SAS Forecast Studio and SAS Forecast Studio for Desktop start at the Diagnose step. You also specify that SAS Forecast Studio and SAS Forecast Studio for Desktop should ignore your selection for the forecast model. In this case, SAS Forecast Studio and SAS Forecast Studio for Desktop select the best-performing model.
- You change the forecast model for one or more series. Then you reforecast the project and have SAS Forecast Studio and SAS Forecast Studio for Desktop start at the Select step. SAS Forecast Studio and SAS Forecast Studio for Desktop keep 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 and SAS Forecast Studio for Desktop should ignore your selection for the forecast model, then SAS Forecast Studio and SAS Forecast Studio for Desktop select the best-performing model.

For more information, see [“Reforecasting a Series or Project” on page 78](#).

Select the Forecast Model

By default, SAS Forecast Studio and SAS Forecast Studio for Desktop determine the best-performing model. To select the forecast model:

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 and SAS Forecast Studio for Desktop use 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 and SAS Forecast Studio for Desktop, 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 Square 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.

View the Diagnostics for Each Model

To view the results of the diagnostic tests for all of the models, click  in the toolbar of the Modeling View.

Using a Holdout Sample

About the Holdout Sample

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:

1. Select **Project** ⇒ **Forecasting Settings**. The Forecasting Settings dialog box appears.
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 $\min(\text{holdout-sample}, \text{maximum percentage}T)$, 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 criteria, see [Appendix 4, “Statistics of Fit,” on page 267](#).

Running an End Zero Test

What Is the End Zero Test?

You can specify when SAS Forecast Studio or SAS Forecast Studio for Desktop should run an end zero model test. When this test is run, SAS Forecast Studio and SAS Forecast Studio for Desktop consider a model only if the series has less than or equal to the maximum number of trailing zero values or the maximum percentage of trailing zero values relative to the number of nonzero values in the entire series. If the series ends with numerous zero values, the forecasts are set to zero.

How to Run the End Zero Test

To specify the values for the end zero test:

1. Select **Project** ⇒ **Forecasting Settings**. The Forecasting Settings dialog box appears.
2. In the selection pane, click **Model Selection**.
3. To specify the maximum number or percentage of trailing values, use the **Maximum number of ending zero values for non-zero model** or **Maximum percentage of ending zero values for non-zero model** options. If you specify both options, then a nonzero model is considered only if the series has less than or equal to the maximum number of trailing zero values and if the number of trailing zero values is less than or equal to the specified percentage of the number of nonmissing and nonzero values of the entire series.
4. (Optional) To specify a threshold for the series length that is required to enable the end zero test, set the **Minimum number of observations to perform the end-zeros test** option. By default, this threshold is 0, and no minimum length is required.

Using an Out-of-Sample Region

What Is an Out-of-Sample Region?

The out-of-sample region is the number of observations before the end of the data where the multistep forecasts are to begin. Removing observations from the end of your data (especially outliers) can help improve your forecasting model and as a result your forecasts. Determining the number of observations to remove can be difficult.

Specify an Out-of-Sample Region

To specify an out-of-sample region:

1. Select **Project** ⇒ **Forecasting Settings**. The Forecasting Settings dialog box appears.
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.

Using Rolling Simulations to Explore Model Performance

Rolling simulations is an out-of-sample analysis that enables you to better understand the lead time performance of a model. For example, if you need to make a decision three periods in the future, you can use the rolling simulation functionality to set the lead time to 3 and see how the model performs.


To view a rolling simulation:

1. Open the Modeling View.
2. In the Hierarchy View or Table View, select a series.
3. Select **Series** ⇒ **Rolling Simulation**. The Rolling Simulation *-series-name* dialog box appears.
4. Specify the number of out-of-sample observations for the simulation. This is the value of the BACK= option.

Note: The value that you specify applies only to the rolling simulation. When a rolling simulation is being performed, the BACK= option is applied to an already detected model that has positive results. The BACK= option for the project is set in the Forecasting Settings dialog box.

5. Specify the horizon to use for the rolling simulation. This is the value of the LEAD= option.
6. Click **Run Simulation**.

Here is an example of the results in the Rolling Simulations dialog box:



In this example, you simulated the values for six observations before the end of the data. The **Simulations** tab is displayed in graphical and tabular format and shows the back values and actual values for each time period. The **Simulations Statistics** tab displays the mean, mean error, MAPE, MAD, and MAD to Mean ratio for observations 1 through 6.

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:

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.

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:

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.


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 you want to use.

To compare models:

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 appears.
3. Compare the plots and statistics of fit for each model in the series. Each model is plotted in a different color.

Figure 16.3 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 4, “Statistics of Fit,”](#) on page 267.

Figure 16.4 Statistics of Fit Tab in Compare Models Dialog Box

Model /	AADJRSQ	ADJRSQ	AIC	AICC	APC	GMAPE	GMAPE5	GMAPE10
Generated ARIMA Model (TOP_0)	-0.30050	-0.24741	-13.528045	-13.441088	0.754403	32E-15	2.65686E-12	32E-15
Generated ARIMA Model (TOP_1)	0.17396	0.22687	-45.704004	-45.267640	0.460908	0.44	46.189696	0.44
Generated Smoothing Model (TOP_2)	-0.08011	-0.02785	-31.092036	-30.881510	0.595606	0.50	52.266861	0.50

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 and SAS Forecast Studio for Desktop add the model to the current series.

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

1. Select **Series** ⇒ **Add Model**. The Add Model dialog box appears.
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:

1. Select **Project** ⇒ **Model Repository**. The Model Repository dialog box appears.
2. Click **Add models from a catalog**. The Import Model Repository dialog box appears.
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:

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

Export a Model Selection List

You might want to export a model selection list, so you can share it with other users at your site. When you export a model selection list, you save it to a library. You must have permission to save content to this library.

To export a model selection list:

1. Select **View** ⇒ **Details** ⇒ **Modeling View** to open the Modeling View.
2. Click . The Export Model Selection List dialog box appears.
3. Select the library where you want to save the model selection list and specify a catalog name. Click **OK**.

Create a New Model Based on a Default Model

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

To create a new custom model based on a default model:

1. Select **Project** ⇒ **Model Repository**. The Model Repository dialog box appears.
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-nameCOPYn**, where **n** is the lowest available integer.
3. Edit the model options and click **OK**.

SAS Forecast Studio and SAS Forecast Studio for Desktop save the new model and displays it in the Model Repository dialog box.

Edit a Model


About Editing Models

In SAS Forecast Studio and SAS Forecast Studio for Desktop, 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, SAS Forecast Studio and SAS Forecast Studio for Desktop automatically update the model in all of the series. You can see whether a model is used in multiple series by using the search feature in the model specification dialog box.

Figure 16.5 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:

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**.

Note: You cannot edit the following types of models:

- default models that shipped with SAS Forecast Studio or SAS Forecast Studio for Desktop.
- models that SAS Forecast Studio and SAS Forecast Studio for Desktop created using the HPFDIAGNOSE procedure. You select the models that SAS Forecast Studio and SAS Forecast Studio for Desktop generate from the Models pane in the Forecasting Settings dialog box. For more information, see [“Select the Models to Fit to Each Series” on page 173](#).
- imported models. For more information, see [“Importing and Exporting Models” on page 180](#).

Edit a Model from the Model Repository

To edit a custom model from the model repository:

1. Select **View** ⇒ **Details** ⇒ **Model Repository**. The Model Repository dialog box appears.
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**.

Copy Models

To copy a model:

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-nameCopy_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 182](#).

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

Combining Models

What Is a Combined Model?

Given the response time series $\{y_t: 1 \leq t \leq n\}$ with previously generated forecasts for the m component models, a combined forecast is created from these component forecasts:

$$\text{Predictions: } \hat{y}_t = \sum_{i=1}^m w_i \hat{y}_{i,t}$$

$$\text{Prediction Errors: } \hat{e}_t = y_t - \hat{y}_t$$

Here, $\hat{y}_{i,t}$ are the forecasts of the component models and w_i are the combining weights.

The estimate of the root mean square prediction error and forecast confidence limits for the combined forecast are computed by assuming independence of the prediction errors of the component forecasts.

$$\text{Standard Errors: } \hat{\sigma}_t = \sqrt{\sum_{i=1}^m w_i^2 \hat{\sigma}_{i,t}^2}$$

$$\text{Confidence Limits: } \pm \hat{\sigma}_t Z_{\alpha/2}$$

Here, $\hat{\sigma}_{i,t}$ are the prediction standard errors for the component models, α is the confidence limit width, $1 - \alpha$ is the confidence level, and $Z_{\alpha/2}$ is the $\frac{\alpha}{2}$ quantile of the standard normal distribution.

In practice, there might be a positive correlation between the prediction errors of the component forecasts, so these confidence limits could be too narrow.

Understanding the Process for Combining Forecasts

SAS Forecast Studio and SAS Forecast Studio for Desktop use the following process to create a combined model:

1. Forecast candidate tests. By using tests for seasonality and intermittency, this step removes any models from the set that are not suitable for the targeted forecasting variable. For example, the target series is seasonal, and some models are not seasonal. These non-seasonal models would be removed from the list of candidates.
2. Forecast missing value tests. This step removes any models that have too many missing values. You can specify thresholds for the percentage of missing values in the horizon or in the estimated region.
3. Encompassing tests. This step removes any models that contain redundant information. This test is performed in a pairwise iteration over the ranked set of candidates that remain after the previous step. Performing an encompassing test is optional.
4. Weight estimation. This step assigns weights to the forecasts that were generated after running any encompassing tests. The weights are assigned based on the information in the combined model list.

5. Forecast combination. This step results in the weighted combination of the forecasts.
6. Weighted average. This step calculates the forecasts using a weighted average of the predictions from the models in the table. Models with missing values are not included in the forecast combination. Instead, these missing values are either dynamically rescaled using the weights of the nonmissing forecasts or generated from the combined forecast. The results of this step are the combined forecast, its residual, standard error, and the upper and lower confidence intervals.

For more information, see the “Forecast Combination Computational Details” chapter in the *SAS Forecast Server Procedures: User’s Guide*.

Generating a Combined Model

About Automatically Generating a Combined Model

Using SAS Forecast Studio or SAS Forecast Studio for Desktop, you can automatically generate a combined model for the set of time series models that is generated from the diagnosis of each series. The options for combined models are available in the **Combined Model** pane in the Forecast Settings dialog box. These options correspond to the options for COMBINE statement in the HPFDIAGNOSE procedure. For more information, see the *SAS Forecast Server Procedures: User’s Guide*.

Note: If you change any of the combined model options and you already have a combined model in your project, then any changes could affect the existing models.

Project Requirements for Automatically Generating a Combined Model

In order for SAS Forecast Studio or SAS Forecast Studio for Desktop to create a combined model, you must select either of the following options in the **Model Generation** pane in the Forecasting Settings dialog box.

- You select more than one model family to fit to each series. For example, if you select both the **System-generated ARIMA models** and the **System-generated exponential smoothing models** check boxes, then SAS Forecast Studio and SAS Forecast Studio for Desktop can create a combined model. However, if you select only the **System-generated exponential smoothing models** check box, then no combined models are created.
- You select the **System-generated ARIMA models** check box and the **Create two models, each of which uses a different identification method for model inclusion** option. Because two models are always available, SAS Forecast Studio and SAS Forecast Studio for Desktop can generate a combined model. You do not have to select the check boxes for the exponential smoothing and unobserved components models.

For more information, see [“Select the Models to Fit to Each Series”](#) on page 173.

How to Automatically Include Combined Models in Your Model Selection List

To automatically generate a combined model:

1. Select **Project** ⇒ **Forecasting Settings**. The Forecasting Settings dialog box appears.
2. In the selection pane, select **Combined Model**.

Note: This option is not available if only a single model (ESM, UCM, or ARIMA) is selected. An exception is if you create two ARIMA models.

3. In the Combined model pane, select the **Combine Diagnosed Models** check box.
4. Set the values that you want SAS Forecast Studio or SAS Forecast Studio for Desktop to use when generating a combined model and click **OK**.

Combine Models in the Model Selection List

Often, combining several forecasting models can result in superior forecasts when compared to a single forecasting model. Using SAS Forecast Studio or SAS Forecast Studio for Desktop, you can create a combined model and add it to the model selection list for the current series. Options in SAS Forecast Studio and SAS Forecast Studio for Desktop enable you to control different aspects of the forecast combination process for the candidate models.

To combine a set of models:

1. Select **View** ⇒ **Details** ⇒ **Modeling View** to open the Modeling View.
2. In the Modeling View, click . The Combine Models dialog box appears.
3. Specify a name and a description for the new model.
4. Select the models for the forecast combination process. Only the models that can be used in a combined model are listed. Because you cannot edit system-generated models or default models that were added through the model repository, copies of these models are used. Even models that previously failed are included in this list because the failed model might work when combined with other models.

However, not all of the models might be in the final result. Whether a model is included in the final result depends on the criteria that you specify for the forecast combination process.

5. In the **Method of combination** drop-down list, select the method for determining the combination weights used in the weighted average of the candidate models. If you select **RANKWGT** or **USERDEF** as the method, then you must specify the weights for each model that you selected. For both methods, the weights must sum to 1.
6. Specify the statistic of fit to use when ranking the selected models or for an encompassing test.
7. In the **Treatment of missing values** drop-down list, select the method to use for any missing values in the combined forecast. For example, these missing values can be dynamically rescaled using the weights of the nonmissing forecasts, or these missing values could be generated in the combined forecast.

Note: For the Ordinary Least Squares (**OLS**) and Nonnegative Restricted Least Squares (**NRLS**) methods of combination, the estimated weights do not have to sum to one. Therefore, the option to rescale the missing values is not available.

8. In the **Compute prediction error variance series** drop-down list, select the method to use for computing the prediction standard error. This error is then used to compute the confidence interval for the combined forecast.
9. To remove any models that contain redundant information, select the **Forecast encompassing test** check box.
10. To remove any models that contain too many missing values, select these options:
 - **Percentage of missing forecast values in the combination horizon**
 - **Percentage of missing forecast values in the combination estimation region**

Because these tests are independent of each other, you can select one test or both tests. For each test that you select, you must specify the threshold for determining when a model has too many missing values.

The new combined model is added to the model selection list for the series.

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. If a model is used in a combined model, then you cannot delete the model until the combined model is deleted.

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 or SAS Forecast Studio for Desktop 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

ARIMA Models	190
Overview of ARIMA Models	190
Create an ARIMA Model	190
Detecting Outliers in ARIMA Models	193
Predefined Variables for the ARIMA Model	194
Subset (Factored) ARIMA Models	194
Overview of Subset (Factored) ARIMA Models	194
Create a Subset (Factored) ARIMA Model	195
Unobserved Components Models	197
Overview of Unobserved Components Models	197
Create an Unobserved Components Model	197
Exponential Smoothing Models	201
Overview of Exponential Smoothing Models	201
Create an Exponential Smoothing Model	202
Intermittent Demand Models	203
Overview of Intermittent Demand Models	203
Create an Intermittent Demand Model	204
Multiple Regression Models	205
Overview of the Multiple Regression Model	205
Create a Multiple Regression Model	206
Moving Average Models	207
Overview of Moving Average Models	207
Create a Moving Average Model	207
Curve Fitting Models	208
Overview of Curve Fitting Models	208
Create a Curve Fitting Model	209
Random Walk Models	210
Overview of Random Walk Models	210
Create a Random Walk Model	210

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 Forecast Server Procedures: User's Guide*.

Note: After any differencing is performed, at least six (6) observations are required for ARIMA models.

Create an ARIMA Model

To create an ARIMA model:

1. Select **Project** ⇒ **Model Repository**. The Model Repository dialog box appears.

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 180.

2. Click **New**. The New Model dialog box appears.
3. From the drop-down list, select **ARIMA** as the model type. The ARIMA Model dialog box appears.

Figure 17.1 ARIMA Model Dialog Box

The screenshot shows the 'ARIMA Model' dialog box. At the top, there's a title bar. Below it, the 'Name' field is set to 'ARIMA1'. The 'Description' field is empty. The 'Details' section shows 'Y = CONST'. On the left, a sidebar contains a list of options: 'Specification' (highlighted), 'Independent Variables', 'Predefined Variables', 'Outlier Variables', 'Events', and 'Estimation'. The main content area is divided into several sections. 'Functional transformation' is set to 'None'. 'Box-Cox parameter' is set to '0'. 'Forecast' is set to 'Mean'. There is a checked box for 'Intercept'. Below this, there are two columns of options: 'ARIMA Options' and 'Seasonal ARIMA Options'. Each column has three rows: 'Autoregressive (p):' with a value of '0', 'Differencing (d):' with a value of '0', and 'Moving Average (q):' with a value of '0'. At the bottom right, there are four buttons: 'View Code', 'OK', 'Cancel', and 'Help'.

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

Note: SAS Forecast Studio and SAS Forecast Studio for Desktop have a list of reserved model names. For more information, see [Appendix 2, “Reserved Names for Variables and Models,”](#) on page 249.

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

Note: The Details pane shows a representation of the model specification. SAS Forecast Studio and SAS Forecast Studio for Desktop generate 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 asks whether 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. The seasonal options are available only if the data is seasonal. 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.

For more information about the syntax of these orders or for more examples, see [Appendix 6, “Autoregressive, Differencing, and Moving Average Orders,”](#) on page 277.

7. To include an independent variable, predefined variable, or event in the model, select **Independent Variables**, **Predefined Variables**, or **Event** in the selection pane. (For more information, see [“Predefined Variables for the ARIMA Model”](#) on page 194.) For each independent variable, predefined variable, or event that you want to include, select the check box in the **Include in Model** column.

To specify whether these variables and events enter the model after some modifications (such as lagging or differencing), click **Transfer Function**. 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 [“Differencing Orders”](#) on page 278.

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)<s1>... (lag,..., lag)<sk>

DEN= order

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

DEN= (lag,..., lag)<s1>... (lag,..., lag)<sk>

For more information about how to specify the polynomial order, see [“Autoregressive Orders”](#) on page 277.

- e. Click **OK**.

8. To view the outlier variables that are included in the model, select **Outlier Variables** in the selection pane. For more information, see [“Remove an Outlier from the Model” on page 194](#).
9. To specify the estimation options, select **Estimation** in the selection pane.
10. Click **OK**.

Detecting Outliers in ARIMA Models

About Detecting Outliers

Outliers show you changes in the level of the response series that are not accounted for by the model. In SAS Forecast Studio and SAS Forecast Studio for Desktop, these changes can be additive outliers or level shifts. For each outlier, SAS Forecast Studio and SAS Forecast Studio for Desktop create a dummy regressor or indicator variable. Then the model and the dummy regressors are fit to the data.

SAS Forecast Studio and SAS Forecast Studio for Desktop use 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:

1. Select **Project** ⇒ **Forecasting Settings**. The Forecasting Settings dialog box appears.
2. In the selection pane, click **Diagnostics**.
3. Select the **Detect outliers** check box to specify whether SAS Forecast Studio or SAS Forecast Studio for Desktop 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 $\min(\text{maximum-number}, \text{percentage}T)$, 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:

1. 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 appears.
2. 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.
3. Select the outlier that you want to delete and click **Delete**.

Predefined Variables for the ARIMA Model

For ARIMA models, you can include a predefined trend or seasonal dummy variables. In the following list of options, let t represent the observation count from the start of the period of fit for the model, and let X_t be the value of the time trend variable at observation t .

LINEAR

a linear trend, with $X_t = t - c$

QUADRATIC

a quadratic trend, with $X_t = (t - c)^2$

CUBIC

a cubic trend, with $X_t = (t - c)^3$

INVERSE

an inverse trend, with $X_t = \frac{1}{t}$

SEASONAL

seasonal dummies. For a seasonal cycle of length s , the seasonal dummy regressors include $X_{i,t}$: $1 \leq i \leq (s - 1)$, $1 \leq t \leq n$ for models that include an intercept term, and $X_{i,t}$: $1 \leq i \leq (s)$, $1 \leq t \leq n$ for models that do not include an intercept.

Each element of a seasonal dummy regressor is either zero or one.

$$X_{i,t} = \begin{cases} 1 & \text{when } i = t \mod s \\ 0 & \text{otherwise} \end{cases}$$

If the model includes an intercept term, the number of seasonal dummy regressors is one less than s to ensure that the linear system is full rank.

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 Forecast Server Procedures: User's Guide*.

Create a Subset (Factored) ARIMA Model

To create a Subset (factored) ARIMA model:

1. Select **Project** ⇒ **Model Repository**. The Model Repository dialog box appears.

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 180](#).

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

The Subset ARIMA Model dialog box appears.

Figure 17.2 Subset ARIMA Model Dialog Box

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

Note: SAS Forecast Studio and SAS Forecast Studio for Desktop have a list of reserved model names. For more information, see [Appendix 2, “Reserved Names for Variables and Models,” on page 249](#).

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

Note: The Details pane shows a representation of the model specification. SAS Forecast Studio and SAS Forecast Studio for Desktop generate 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 asks whether you want to suppress the intercept.

- the autoregressive (p), differencing (d), and moving average (q) orders for your model.

For SAS Forecast Studio or SAS Forecast Studio for Desktop 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 6, “Autoregressive, Differencing, and Moving Average Orders,”](#) on page 277.

7. To include an independent variable, predefined variable, or event in the model, select **Independent Variables**, **Predefined Variables**, or **Event** in the selection pane. (For more information, see [“Predefined Variables for the ARIMA Model”](#) on page 194.) For each independent variable, predefined variable, or event that you want to include, select the check box in the **Include in Model** column.

To specify whether these variables and events enter the model after some modifications (such as lagging or differencing), click **Transfer Function**. 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 [“Differencing Orders”](#) on page 278.

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)<s1>... (lag,..., lag)<sk>

DEN= order

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

DEN= (lag,..., lag)<s1>... (lag,..., lag)<sk>

For more information about how to specify the polynomial order, see [“Autoregressive Orders”](#) on page 277.

- e. Click **OK**.
8. To view the outlier variables in the model, select **Outlier Variables** in the selection pane. For more information, see [“Remove an Outlier from the Model”](#) on page 194.

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

Unobserved Components Models

Overview of Unobserved Components Models

Unobserved components models can include events and independent variables.

For more information about the Unobserved Components model, see the HPFUCMSPEC procedure in the *SAS Forecast Server Procedures: User's Guide*.


Create an Unobserved Components Model

To create an unobserved components model:

1. Select **Project** ⇒ **Model Repository**. The Model Repository dialog box appears.
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 180](#).
2. Click **New**. The New Model dialog box appears.
3. From the drop-down list, select **Unobserved Components** as the model type.

The Unobserved Components Model dialog box appears.

Figure 17.3 Unobserved Components Model Dialog Box



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

Note: SAS Forecast Studio and SAS Forecast Studio for Desktop have a list of reserved model names. For more information, see [Appendix 2, “Reserved Names for Variables and Models,”](#) on page 249.

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

Note: The Details pane shows a representation of the model specification. SAS Forecast Studio and SAS Forecast Studio for Desktop generate 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 Seasonal Component Properties dialog box appears.
- 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 Block-Seasonal Component Properties dialog box appears.
 - 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 Cycle Components Properties dialog box appears.
 - b. To specify an initial cycle period to use during the parameter estimation process, select the **Define initial period** check box. 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 or SAS Forecast Studio for Desktop uses during the parameter estimation process, select the **Define initial variance** check box. 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 as 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 the following syntax for the **Lags** option: (integer, ... ,integer)...(integer, ... ,integer). This syntax creates a lag of (integer, ..., integer)...(integer, ..., integer).

The following table shows some examples of this syntax:

Model Description	Value of 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.
- To specify whether the independent variable enters the model as a simple regressor or it enters the model after some modifications (such as lagging or differencing), click **Transfer Function**. 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, see [“Differencing Orders” on page 278](#).
 - 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 or SAS Forecast Studio for Desktop, you can create many types of exponential smoothing models. For more information about the exponential smoothing model, see the HPFESMSPEC procedure in the *SAS Forecast Server Procedures: User's Guide*.

Create an Exponential Smoothing Model

To create an exponential smoothing model:

1. Select **Project** ⇒ **Model Repository**. The Model Repository dialog box appears.

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 180](#).

2. Click **New**. The New Model dialog box appears.
3. From the drop-down list, select **Exponential smoothing** as the model type.

The Exponential Smoothing Model dialog box appears.

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 and SAS Forecast Studio for Desktop automatically generate a unique name for the model.

Note: SAS Forecast Studio and SAS Forecast Studio for Desktop have a list of reserved model names. For more information, see [Appendix 2, “Reserved Names for Variables and Models,” on page 249](#).

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

Note: The Details pane shows a representation of the model specification. SAS Forecast Studio and SAS Forecast Studio for Desktop generate 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:

- a. Select **Specify** as the weight and click **Edit**. The Edit Weights dialog box appears.
- 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:

1. Select **Project** ⇒ **Model Repository**. The Model Repository dialog box appears.

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 180](#).

2. Click **New**. The New Model dialog box appears.
3. From the drop-down list, select **Intermittent demand** as the model type.

The Intermittent Demand Model dialog box appears.

Figure 17.5 Intermittent Demand Model Dialog Box

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

Note: SAS Forecast Studio and SAS Forecast Studio for Desktop have a list of reserved model names. For more information, see [Appendix 2, “Reserved Names for Variables and Models,” on page 249](#).

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

Note: The Details pane shows a representation of the model specification. SAS Forecast Studio and SAS Forecast Studio for Desktop generate 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 and SAS Forecast Studio for Desktop 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 and SAS Forecast Studio for Desktop 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 Model**, **Demand Intervals Model**, or **Demand Sizes Model** in the selection pane. Then 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:

1. Select **Specify** as the weight and click **Edit**. The Edit Weights dialog box appears.
 2. Restrict the values and set the initial value for the parameters. Only the parameters that are relevant for a model are available.
 3. 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.
 4. 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:


1. Select **Project** ⇒ **Model Repository**. The Model Repository dialog box appears.

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 180](#).

2. Click **New**. The New Model dialog box appears.
3. From the drop-down list, select **Multiple Regression** as the model type.

The Multiple Regression Model dialog box appears.

Figure 17.6 Multiple Regression Model Dialog Box



Note: The Multiple Regression Model dialog box also appears 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 and SAS Forecast Studio for Desktop automatically generate a unique name for the model.

Note: SAS Forecast Studio and SAS Forecast Studio for Desktop have a list of reserved model names. For more information, see [Appendix 2, “Reserved Names for Variables and Models,” on page 249](#).

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

Note: The Details pane shows a representation of the model specification. SAS Forecast Studio and SAS Forecast Studio for Desktop generate 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 or SAS Forecast Studio for Desktop, you can create a moving average model. The formula for the moving average model with width k is

$$y_t = \frac{[y_{t-1} + \dots + y_{t-k}]}{k} + \text{error}.$$

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

Note: For moving average models, the minimum data requirement for n is $p+2$. In other words, the number of observations, n , must be greater than or equal to the number of potential parameters, $p, + 2$.

Create a Moving Average Model

To create a moving average model:


1. Select **Project** $\Rightarrow$ **Model Repository**. The Model Repository dialog box appears.

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 180](#).

2. Click **New**. The New Model dialog box appears.
3. From the drop-down list, select **Moving Average** as the model type.

The Moving Average Model dialog box appears.

Figure 17.7 Moving Average Model Dialog Box



Note: The Moving Average Model dialog box also appears 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 and SAS Forecast Studio for Desktop automatically generate a unique name for the model.

Note: SAS Forecast Studio and SAS Forecast Studio for Desktop have a list of reserved model names. For more information, see [Appendix 2, “Reserved Names for Variables and Models,”](#) on page 249.

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

Note: The Details pane shows a representation of the model specification. SAS Forecast Studio and SAS Forecast Studio for Desktop generate 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 or SAS Forecast Studio for Desktop, you can create a curve fitting model with a linear or quadratic trend.

Create a Curve Fitting Model

To create a curve fitting model:

1. Select **Project** ⇒ **Model Repository**. The Model Repository dialog box appears.

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 180](#).

2. Click **New**. The New Model dialog box appears.
3. From the drop-down list, select **Curve Fitting** as the model type.

The Curve Fitting Model dialog box appears.

Figure 17.8 Curve Fitting Model Dialog Box

Note: The Curve Fitting Model dialog box also appears 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 and SAS Forecast Studio for Desktop automatically generate a unique name for the model.

Note: SAS Forecast Studio and SAS Forecast Studio for Desktop have a list of reserved model names. For more information, see [Appendix 2, “Reserved Names for Variables and Models,” on page 249](#).

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

Note: The Details pane shows a representation of the model specification. SAS Forecast Studio and SAS Forecast Studio for Desktop generate 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 or SAS Forecast Studio for Desktop, 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 $y_t = y_{t-1} + \text{error}$.

You can also create the following random walk models:

Random Walk with Drift

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

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:

1. Select **Project** ⇒ **Model Repository**. The Model Repository dialog box appears.

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 180](#).

2. Click **New**. The New Model dialog box appears.
3. From the drop-down list, select **Random walk** as the model type.

The Random Walk Model dialog box appears.

Figure 17.9 Random Walk Model Dialog Box

Note: The Random Walk Model dialog box also appears 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 and SAS Forecast Studio for Desktop automatically generate a unique name for the model.

Note: SAS Forecast Studio and SAS Forecast Studio for Desktop have a list of reserved model names. For more information, see [Appendix 2, “Reserved Names for Variables and Models,”](#) on page 249.

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

Note: The Details pane shows a representation of the model specification. SAS Forecast Studio and SAS Forecast Studio for Desktop generate 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

What Is Scenario Analysis?	213
Create a Scenario	214
Editing the Values of the Input Variables	214
About Editing Input Variables	214
Edit a Value from the Input Table	214
Edit Values By Using the Input Calculator	215
Run a Scenario	215
Compare Scenarios	215
Set Scenario Forecast Values as Overrides	215
Copy a Scenario	216
Reset Input Values in a Scenario	216
Example: Performing a Scenario Analysis	216
Creating the Price Increase Scenario	216
Creating the Discount Decrease Scenario	219
Applying a Scenario to the Project	220

What Is Scenario Analysis?

SAS Forecast Studio and SAS Forecast Studio for Desktop enable 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. Here are some questions that you can answer by using scenario analysis:

- What happens to the forecasts when the price increases?
- How does a promotion plan affect my sales forecasts for the promoted product?

The Scenario Analysis View enables you to quickly create and run scenarios, compare different scenarios, save scenarios for future use, and apply the forecast values from a scenario in your project. Although the Scenario Analysis View is integrated in the workspace 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 and SAS Forecast Studio for Desktop 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 and SAS Forecast Studio for Desktop, 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 and SAS Forecast Studio for Desktop use the HPFENGINE procedure to generate score files for the forecasts. For more information about the procedures and functions that SAS Forecast Studio and SAS Forecast Studio for Desktop use, see "Using Forecasting Model Score Files and DATA Step Functions" in *SAS Forecast Server Procedures: User's Guide*.

Create a Scenario

To create a scenario:

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 appears.
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 entering 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:

1. In the scenario table, select an input variable. To view all of the input variables, select the **Show All Inputs** check box.
2. In the input table, double-click in the cell for the time period that you want to edit.
3. Type the new value and press ENTER.

Edit Values By Using the Input Calculator

To edit the values for multiple time periods:

1. In the Scenario Analysis View, select **Series** ⇒ **Input Calculator**. The Input Calculator appears.
2. Select the time periods that you want to edit.

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:

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:

1. In the Scenario Analysis View, click **Compare**. The Compare Scenarios dialog box appears.
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:

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:

1. Select the scenario that you want to copy.
2. Click **Save As**. The Edit Scenario dialog box appears.
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:

1. In the hierarchy, select the Region2 node.


Figure 18.1 Region2 in the Hierarchy

2. Select **View** ⇒ **Details** ⇒ **Scenario Analysis View**.
3. In the Scenario Analysis View, click **New**. The Create New Scenario dialog box appears.
4. In the **Name** box, type **Price_Increase**.
5. In the **Description** box, type **Reflects increase in gas prices**.
6. In the Create New Scenario dialog box, click **Quick View** to determine which models use Price as an input variable. The Inputs Used dialog box shows that only the Generated ARIMA Model (HPF0_4) uses Price as an input variable. Click **OK** to close this dialog box.



7. In the Create New Scenario dialog box, select the **Generated ARIMA Model (HPF0_4)** model.

The Create New Scenario dialog box 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



8. 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 (Jan03–Dec03) and select **Series** ⇒ **Input Calculator**. The Input Calculator appears.
 - b. 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, 70.12 now appears as the value of PRICE for each time period.

Figure 18.5 New PRICE Values in the Input Table

Input table:	Jun03	Jul03	Aug03	Sep03	Oct03	Nov03	Dec03
* PRICE	70.12	70.12	70.12	70.12	70.12	70.12	70.12
* DISCOUNT	0.0110104498	0.0109941332	0.010978471	0.0109634369	0.0109490057	0.0109351533	0.0109218564
Model Forecast	3499.2926202	3345.292872	3347.7091668	3220.0446653	3134.8580204	3176.7573196	3104.8530041
Scenario Forecast	3499.2926202	3345.2928719	3347.7091668	3220.0446652	3134.8580203	3176.7573196	3104.8530041

9. 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 3178.2992661. However, the forecast that is generated by the scenario is 3061.4433835. This is a decrease of 116.8492775 units (or a decrease in sales of approximately 3.7%).

Figure 18.6 Difference in Model and Scenario Forecasts for Jan2003

Input table:	Dec02	Jan03	Feb03	Mar03	Apr03	May03	Jun03
* PRICE	68.375	70.12	70.12	70.12	70.12	70.12	70.12
* DISCOUNT	0	0.0111028459	0.0110828239	0.0110636048	0.0110451565	0.011027448	0.0110104498
Model Forecast	53.9306287	3178.2992661	3339.8304369	3274.1385102	3361.3092411	3425.0564651	3499.2926202
Scenario Forecast		3061.4433835	3223.0735467	3157.4770145	3244.7396728	3308.5754831	3382.8970049

Creating the Discount Decrease Scenario

You create another scenario to determine how forecasts would be impacted if you eliminated the current customer discount.

To perform this scenario analysis:

1. In the Scenario Analysis View, click **New**. The Create New Scenario dialog box appears.
2. In the **Name** box, type **Discount_Decrease**.
3. In the **Description** box, type **Eliminates the discount**.
4. In the Create New Scenario dialog box, select the **Generated ARIMA Model (HPF0_4)** model.

Click **OK**. The model now appears in the Scenario Analysis View.

5. In the input table, edit the value of the DISCOUNT variable by completing the following steps:
 - a. Select all time periods in the DISCOUNT variable row (Jan03–Dec03) and select **Series** ⇒ **Input Calculator**. The Input Calculator appears.
 - b. Select **Set to a value** and type 0 in the text box.
 - c. From the **Distribute value among periods** drop-down list, select **Assign specified value to each period**.

Click **OK**. In the input table, 0 now appears as the value of DISCOUNT for each time period.

6. Click **Run Scenario**.

As with the Price Increase scenario, you see that eliminating the discount decreases the generated forecast.

Input table: ☐ Show All Inputs								
	Dec02	Jan03	Feb03	Mar03	Apr03	May03	Jun03	
* PRICE	68.375	67.635649688 ▲	67.637754259 ▲	67.639782341 ▲	67.641736713 ▲	67.643620055 ▲	67.645434947 ▲	67.6471
* DISCOUNT	0	0 ▲	0 ▲	0 ▲	0 ▲	0 ▲	0 ▲	
Model Forecast	306287	3178.2992661	3339.8304369	3274.1385102	3361.3092411	3425.0564651	3499.2926202	3345
Scenario Forecast	.	3155.9806411	3323.0228622	3261.7680845	3352.5107599	3419.1328865	3495.6824021	3343.5

Applying a Scenario to the Project


You want to see how the Price Increase and Discount Decrease scenarios compare. After the comparison, you will apply the values for one of the scenario forecasts so that they become the actual forecast values in the project.

To compare and select a scenario:

1. To compare Price Increase and Discount Decrease scenarios, click **Compare**.

The Compare Scenarios dialog box appears.

Figure 18.7 Comparison Plots of the Scenarios




These plots show that from January 2003 until June 2003 the forecasts for the Price Increase scenario are slightly higher than the forecasts for the Discount Decrease scenario. After June 2003, the two scenarios are virtually identical, but both scenarios forecast show higher sales than the baseline forecast. From this analysis, you decide to apply the values from the Price Increase scenario to the Region2 series.

Click **Close** to close the Compare Scenarios dialog box.

2. In the Scenario Analysis View, select the **Discount_Decrease** 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 or SAS Forecast Studio for Desktop uses the scenario forecasts as overrides for the Region2 series.

Part 5

Advanced Topics

<i>Chapter 19</i>	
Managing Projects, Environments, and Libraries	225
<i>Chapter 20</i>	
Working in Batch Mode	229
<i>Chapter 21</i>	
Working with SAS Forecast Studio Tasks	237

Chapter 19

Managing Projects, Environments, and Libraries

About the SAS Forecast Project Manager	225
Managing Projects	225
Managing Forecasting Environments	226
Managing Environment Libraries	226
How You Can Run the SAS Forecast Project Manager	226
Additional Documentation	227

About the SAS Forecast Project Manager

Whether you license SAS Forecast Server or SAS Forecasting for Desktop, you can use the SAS Forecast Project Manager to manage your forecasting projects and forecasting environments.

Note: The functionality in the SAS Forecast Project Manager is similar to the functionality that is available from the SAS Forecast Server plug-in for SAS Management Console. For more information about the SAS Forecast Server plug-in for SAS Management Console, see the *SAS Forecast Server: Administrator's Guide* and the Help for the SAS Forecast Server plug-in for SAS Management Console.

Managing Projects

Here are some of the tasks that you can perform in the SAS Forecast Project Manager to manage your projects:

- register and unregister a project
- copy or move a project to a different destination environment
- archive and unarchive a project
- specify whether to share a project with other users at your site
- change the project owner

- update a project from a previous release of SAS Forecast Studio to use the current release of SAS Forecast Studio
- rename a project
- delete a project

For more information about projects, see [“What Is a Project?” on page 61](#).

Managing Forecasting Environments

Here are some of the tasks that you can perform in SAS Forecast Project Manager to manage your forecasting environments:

- create an environment
- register and unregister an environment
- delete all of the unregistered content (referred to as purging) in the environment
- rename an environment
- delete an environment
- deploy sample reports and update reports to use the FS_REPORT keyword

Note: Reports are available only in SAS Forecast Studio. Reports are not available if you are running SAS Forecast Studio for Desktop. For more information, see [Chapter 2, “About SAS Forecast Studio for Desktop,” on page 7](#).

- add code that runs when the environment opens or when the environment closes

For more information about environments, see the *SAS Forecast Server: Administrator’s Guide*.

Managing Environment Libraries

Here are some of the tasks that you can perform in SAS Forecast Project Manager to manage libraries that are associated with your forecasting environments:

- create a new environment library
- edit the configuration options for the environment library, such as the library name and the access level (which is either read-only or read-write)
- delete an environment library

For more information about environment libraries, see the *SAS Forecast Server: Administrator’s Guide*.

How You Can Run the SAS Forecast Project Manager

The SAS Forecast Project Manager runs on Windows.

If your site licenses SAS Forecast Server, you can run the SAS Forecast Project Manager in either of the following ways:

- from a local installation of SAS Forecast Server on your computer. If you installed on a computer in the Windows operating environment, select **Start** ⇒ **Programs** ⇒ **SAS** ⇒ **SAS Forecast Studio** ⇒ **SAS Forecast Project Manager 14.1**.
- by using Java Web Start. Contact your site administrator for the URL for your site. For more information about the Java Web Start client, see the *SAS Forecast Server: Administrator's Guide*.

To start the SAS Forecast Project Manager if your site licenses SAS Forecasting for Desktop, select **Start** ⇒ **Programs** ⇒ **SAS** ⇒ **SAS Forecast Studio** ⇒ **SAS Forecast Project Manager for Desktop 14.1**.

Note: The Java Web Start client is not available if you license SAS Forecasting for Desktop. For more information, see [Chapter 2, “About SAS Forecast Studio for Desktop,” on page 7](#).

Additional Documentation

For more information about how to complete a task in SAS Forecast Project Manager, see the online Help that is available in SAS Forecast Project Manager.

Chapter 20

Working in Batch Mode

Working with the SAS Code	229
About the SAS Code	229
View the Project Code	229
Save the Code to a File	230
Running Projects in Batch Mode	230
Saving the Project Code to a Batch File	230
Descriptions of the Project Code in Each Batch File	231

Working with the SAS Code

About the SAS Code

Note: You can work in batch mode only if your site licenses SAS Forecast Server (which includes the batch interface). SAS Forecasting for Desktop does not include the batch interface. If your site licenses SAS Forecasting for Desktop, you can view the SAS code that is generated by SAS Forecast Studio for Desktop, but you cannot run it in batch mode. For more information, see [Chapter 2, “About SAS Forecast Studio for Desktop,” on page 7](#).

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 229](#).

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 230](#).

View the Project Code

To view the project code that SAS Forecast Studio or SAS Forecast Studio for Desktop generated:

1. Select **Project** ⇒ **SAS Code**. The SAS Code dialog box appears.
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. (Optional) 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.
4. (Optional) Select the **Include code to generate index files** check box to generate the index files that can provide faster access to specific observations in your data sets. This option is selected by default. If you do not select this option, SAS Forecast Studio generates some but not all of the index files. For more information about indexes, see the *SAS Language Reference: Concepts* at <http://support.sas.com/documentation/onlinedoc/base>.

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:

1. Select **Project** ⇒ **SAS Code**. The SAS Code dialog box appears.
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 appears.
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 with 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

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.

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 code files contain 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` code 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` code 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` code 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.

Chapter 21

Working with SAS Forecast Studio Tasks

Overview of the SAS Forecast Studio Tasks	237
What Are the SAS Forecast Studio Tasks?	237
What Is the SAS Add-In for Microsoft Office?	237
What Is SAS Enterprise Guide?	238
Prerequisites for Using the SAS Forecast Studio Tasks	239
The Forecast Studio Create Project Task	239
The Forecast Studio Open Project Task	240
The Forecast Studio Override Project Task	240
Additional Information	241

Overview of the SAS Forecast Studio Tasks

What Are the SAS Forecast Studio Tasks?

Note: The SAS Forecast Studio tasks are available only if your site licenses SAS Forecast Server. These tasks are not available if you are licensing SAS Forecasting for Desktop. For more information, see [Chapter 2, “About SAS Forecast Studio for Desktop,” on page 7](#).

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 a Microsoft 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 Override Project 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, a **SAS** tab is automatically integrated into Microsoft Excel, Microsoft Word, and Microsoft PowerPoint.

Figure 21.1 SAS Tab in Microsoft Excel



The SAS add-in includes approximately 80 SAS tasks that enable you to perform a variety of analyses. The Forecast Studio Create Project task, the Forecast Studio Open Project task, and the Forecast Studio Submit Overrides task are available in the Analyze Data window. In the SAS Tasks folder, click **Time Series** to view these tasks.


Figure 21.2 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 21.3 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 the SAS Add-In for Microsoft Office or SAS Enterprise Guide, your site administrator must complete these steps:

- install SAS Forecast Server.
- install the SAS Add-In for Microsoft Office or SAS Enterprise Guide.
- configure SAS Forecast 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 a Microsoft 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, Microsoft Word, or Microsoft PowerPoint.

Here are some examples of the results that you can open:

- 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 are shipped with the SAS Add-In for Microsoft Office or SAS Enterprise Guide.

The Forecast Studio Override Project Task

The Forecast Studio Override Project 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 forecast hierarchically.

- If the data in the project is forecast 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 forecast 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 an error message appears in the SAS Add-In for Microsoft Office or SAS Enterprise Guide. The conflicts cannot be reconciled in the SAS Add-In for Microsoft Office or SAS Enterprise Guide. 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. Alternatively, the SAS Forecast Studio user can close this dialog box and update the overrides in the data table.

Additional Information

The availability of these tasks depends on the capabilities that were assigned to you. (For more information about capabilities, see [“Understanding Role-Based Settings in SAS Forecast Studio” on page 6](#).) If the Forecast Studio tasks are not available to you in the SAS Add-In for Microsoft Office or SAS Enterprise Guide, contact your site administrator.

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 6

Appendixes

<i>Appendix 1</i>	
Troubleshooting Tips	245
<i>Appendix 2</i>	
Reserved Names for Variables and Models	249
<i>Appendix 3</i>	
Output Data Sets in the Project Directory	255
<i>Appendix 4</i>	
Statistics of Fit	267
<i>Appendix 5</i>	
Default Model Selection Lists	273
<i>Appendix 6</i>	
Autoregressive, Differencing, and Moving Average Orders	277
<i>Appendix 7</i>	
Sample Reports in SAS Forecast Studio	281

Appendix 1

Troubleshooting Tips

Viewing the SAS Log	245
Troubleshooting the Time ID Variable	245
Troubleshooting Other Types of Variables	247
Performance	247

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 or SAS Forecast Studio for Desktop creates when generating the forecast, select **Tools** ⇒ **SAS Log**. If an error occurs when SAS Forecast Studio or SAS Forecast Studio for Desktop is generating the forecasts, 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 and SAS Forecast Studio for Desktop require a specific format for the input data. For more information about how SAS Forecast Studio and SAS Forecast Studio for Desktop create the time series data, see [Chapter 5, “Preparing an Input Data Set,”](#) on page 33.

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*.

Here are some common errors:

- SAS Forecast Studio and SAS Forecast Studio for Desktop determine the time interval that best fits the data. If SAS Forecast Studio or SAS Forecast Studio for Desktop detects multiple time intervals, SAS Forecast Studio selects the most frequently used base interval. Although SAS Forecast Studio and SAS Forecast Studio for Desktop select the default interval, you can change this interval using the **Interval** option in the New Project wizard or the Forecasting Settings dialog box.
- SAS Forecast Studio or SAS Forecast Studio for Desktop issue warnings about gaps in the data or duplicate values in the time ID variable. It is important to investigate these warnings. When possible, SAS Forecast Studio and SAS Forecast Studio for Desktop provide graphs of your time ID data to help you determine whether the gaps in the data are correct. The diagnostic graphs are shown for the first BY group with a problem. If these data issues are fixed, the warning might appear again if there is a problem in another BY group. If your data is transactional or you are confident that the data is correct, then you can ignore these warnings.
- SAS Forecast Studio or SAS Forecast Studio for Desktop is unable to determine the time interval. When SAS Forecast Studio or SAS Forecast Studio for Desktop is unable to determine the time interval for the data, you must select another variable as the time ID variable or to manually set the time interval.

Here are the reasons why SAS Forecast Studio or SAS Forecast Studio for Desktop 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 this case, you must specify the time interval to use. Then SAS Forecast Studio or SAS Forecast Studio for Desktop 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 or SAS Forecast Studio for Desktop cannot detect the time interval. For example, SAS Forecast Studio or SAS Forecast Studio for Desktop cannot detect the time interval if the data represents a multi-weekday interval, such as the **WEEKDAY3.** interval.

You must manually specify the time interval to use.

- The wrong variable has been selected as the time ID variable.

In the Assign Variables to Roles step in the New Project wizard, you must 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 and SAS Forecast Studio for Desktop support 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 and SAS Forecast Studio for Desktop do 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 and SAS Forecast Studio for Desktop assume 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 or SAS Forecast Studio for Desktop. 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 or SAS Forecast Studio for Desktop 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 the product to find the format catalog and recognize your custom format, you must specify the FMTSEARCH system option in your start-up files. For more information about how to add a custom format, see the *SAS Forecast Server: Administrator's Guide*.
- 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.

This error is detected after you finish the New Project wizard. You will not see this error when you assign a variable to the BY variable role.

Performance

If you are experiencing performance issues, here are some areas to investigate.

- If your project has many independent variables or events, reduce the number of inputs. A large number of inputs requires more resources.
- If **System-generated unobserved components models** is selected under Model Generation settings, disable the option. This class of model requires a high level of processing to estimate.
- If **Detect outliers** is selected under Diagnostic settings, disable the option. Outlier detection is an extremely I/O intensive analysis.

Appendix 2

Reserved Names for Variables and Models

Reserved Variable Names	249
Overview of Reserved Variable Names	249
Reserved Variable Names	249
Reserved Model Names	254

Reserved Variable Names

Overview of Reserved Variable Names

For each project, SAS Forecast Studio and SAS Forecast Studio for Desktop create 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 and SAS Forecast Studio for Desktop create, see [Appendix 3, “Output Data Sets in the Project Directory,” on page 255](#).

Reserved Variable Names

The following table lists alphabetically the variables that are used by SAS Forecast Studio and SAS Forecast Studio for Desktop.

Variable Name	Description
<i>_VariableName</i>	Any variable name that begins with an underscore
AADJRSQ	Amemiya’s adjusted R-square
ACTUAL	Dependent series value
ADJRSQ	Adjusted R-square

Variable Name	Description
AGGCHILDPREDICT	Aggregated prediction of child nodes
AIC	Akaike Information Criterion
AICC	Finite sample corrected Akaike Information Criterion
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

Variable Name	Description
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

Variable Name	Description
MPE	Mean percent error
MPPE	Mean predictive percent error
MRAE	Mean relative absolute error
MRE	Mean relative error
MSE	Mean square error
MSPE	Mean symmetric percent error
N	Number of nonmissing observations or number of variance products
NAME	Variable name
NMISS	Number of missing observations
NMISSA	Number of missing actuals
NMISSP	Number of missing predicted
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

Variable Name	Description
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 and SAS Forecast Studio for Desktop use the following keywords in the names of the automatically generated models:

- LEAF
- TOP
- HPF

These words are reserved for SAS Forecast Studio and SAS Forecast Studio for Desktop 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, then SAS Forecast Studio(or SAS Forecast Studio for Desktop) names the copied model LEAFCOPY n , where n is the lowest available integer value. Because this name is unique, SAS Forecast Studio and SAS Forecast Studio for Desktop allow this name.

Appendix 3

Output Data Sets in the Project Directory

Overview of Output Data Sets	255
FINALFOR= Data Set	256
OUT= Data Set	257
OUTCOMPONENT= Data Set	257
OUTINDEP= Data Set	258
OUTEST= Data Set	258
OUTFOR= Data Set	259
OUTOVRD= Data Set	260
OUTSTAT= Data Set	260
OUTSTATSELECT= Data Set	262
OUTSUM= Data Set	262
RECFOR= Data Set	263
RECSTAT= Data Set	265
RECSUM= Data Set	265
SCENARIO= Data Set	266

Overview of Output Data Sets

SAS Forecast Studio and SAS Forecast Studio for Desktop create 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 project directory for the specified environment. For example, if you are using SAS Forecast Studio, the output data sets could be saved in these locations:

Windows environments

```
SAS-configuration-directory\Lev1\AppData
\ForecastServer14.1\Projects\project-name\hierarchy
```

UNIX environments

```
SAS-configuration-directory/Lev1/AppData/
SAS/ForecastServer14.1/Projects/project-name/hierarchy
```

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. 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. These missing values appear as .F in the data.

This appendix documents the output data sets that are unique to SAS Forecast Server. If you do not see the documentation for an output data set in this appendix, see the *SAS Forecast Server Procedures: User's Guide*.

FINALFOR= Data Set

The FINALFOR= data set contains the observations that have a time ID value in the forecast horizon. If there are no overrides, the PREDICT column in the FINALFOR= data set is equivalent to the PREDICT column in the OUTFOR= or RECFOR= data sets. When overrides are specified, the values in the FINALFOR= data set reflect the effect of the overrides.

The FINALFOR= data set contains the variables that you assign as BY variables and the variables listed below:

- ACTUAL - values of the dependent series
- DATE - values of the time ID variable
- ERROR - prediction errors
- LOWER - lower confidence limits
- PREDICT - predicted values
- STD - prediction standard errors
- UPPER - upper confidence limits
- _NAME_ - variable name
- _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
400	An unlocked equality constraint has been imposed.
500	A locked equality constraint has been imposed.
600	A lower bound is active.
700	An upper bound is active.
1000	ID value out of 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 STDMETHOD= option that was provided by the user has been changed for this observation. |
| 9000 | The CLMETHOD= option that was provided by the user has been changed for this observation. |
| 10000 | The standard error hit the limits that are imposed by the STDDIFBD= option. |
| 11000 | Multiple warnings have been printed to the log for this observation. |
| 12000 | The number of missing values in the STD variable in the DISAGGDATA= data set is different from the number of missing values in the union of the PREDICT and ACTUAL variables. |
| 13000 | The solution might be suboptimal. This means that the optimizer did not find an optimal solution, but the solution that was provided satisfies all constraints. |
| 14000 | A failed forecast “.F” has been detected in a relevant input variable. |
- LOWBFOVR – the lower confidences limits before reconciling overrides
 - PREBFOVR – the predicted values before reconciling overrides
 - STDBFOVR – the predicted standard errors before reconciling overrides
 - UPPBFOVR – the upper confidences limits before reconciling overrides

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 and 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

OUTINDEP= Data Set

The OUTINDEP= data set contains the forecasting values for the independent variables. These future values of the independent variables are supplied by the HPFENGINE procedure. For more information, see [“Create a Forecast Data Set for Independent Variables” on page 76](#).

The OUTINDEP= data set contains the variables that you assign as BY variables and the following variables:

- `_NAME_` - variable name
- `_TIMEID_` - time ID values
- `_X_` - values of the input variable `_NAME_`

Note: For more information about the models used to generate the forecasts, see [“OUTEST= Data Set” on page 258](#).

OUTEST= Data Set

The OUTEST= data set contains the variables that you assign as BY variables. This data set also contains 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
- `_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.

Here are the values for the `_STATUS_` variable:

0	The forecast completed successfully.
3000	Model selection could not be completed. Forecast values are set to missing.
3001	Model selection could not be completed because no models can be fit from the initial selection list. Forecast values are set to missing.
3002	The forecast completed, but one or more input variables were omitted from the selected model. This can occur only for ARIMAX or unobserved component models (UCM).
3003	The selected model could not be forecast. The forecast reverted to the default exponential smoothing model.
3004	An arithmetic exception occurred when trying to forecast the selected model. The forecast is generated by using the default exponential smoothing model for the variable that produces the arithmetic exception.
3005	An arithmetic exception occurred when trying to forecast the selected model. The forecast is generated by using a zero-drift random walk model for the BY variable that produces the exception.
3006	An arithmetic exception occurred when trying to forecast the selected model. The forecast of missing values is generated for the BY variable that produces the exception.
3007	The mean value of the forecast is generated because minimum number of observations is the selection criterion.
3008	There were insufficient nonmissing observations in the variable to forecast. A missing value is produced for the forecast.
3009	There were insufficient nonzero observations in the variable to forecast. A missing value is produced for the forecast.

OUTFOR= Data Set

The OUTFOR= data set contains the variables that you assign as BY variables and 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 or SAS Forecast Studio for Desktop. 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 specified overrides for your project. If you did not specify any overrides, then SAS Forecast Studio or SAS Forecast Studio for Desktop does not create this data set.

OUTSTAT= Data Set

The OUTSTAT= data set contains the variables that you assign as BY variables and the variables listed below. This data set contains the statistics for the current model for the most recent data update. The OUTSTAT data set contains the statistics-of-fit that appear in the Forecasting View.

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 squared 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 statistics that were used to select the current model. This data set contains the statistics-of-fit that appear in the Modeling View.

This 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

If the selection of the model was used with an older data or if there is a holdout sample, the values in the OUTSTATSELECT= data set could differ from the values in the OUTSTAT= data set.

OUTSUM= Data Set

The OUTSUM= data set contains the variables that you assign as BY variables. This data set also records the summary statistics for each forecast variable and the multistep forecasts as determined by the value of the LEAD= option.

The following variables that are related to summary statistics 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.

Here are the values for the _STATUS_ variable:

- 0 The forecast completed successfully.
- 3000 Model selection could not be completed. Forecast values are set to missing.

- 3001 Model selection could not be completed because no models can be fit from the initial selection list. Forecast values are set to missing.
- 3002 The forecast completed, but one or more input variables were omitted from the selected model. This can occur only for ARIMAX or unobserved component models (UCM).
- 3003 The selected model could not be forecast. The forecast reverted to the default exponential smoothing model.
- 3004 An arithmetic exception occurred when trying to forecast the selected model. The forecast is generated by using the default exponential smoothing model for the variable that produces the arithmetic exception.
- 3005 An arithmetic exception occurred when trying to forecast the selected model. The forecast is generated by using a zero-drift random walk model for the BY variable that produces the exception.
- 3006 An arithmetic exception occurred when trying to forecast the selected model. The forecast of missing values is generated for the BY variable that produces the exception.
- 3007 The mean value of the forecast is generated because minimum number of observations is the selection criterion.
- 3008 There were insufficient nonmissing observations in the variable to forecast. A missing value is produced for the forecast.
- 3009 There were insufficient nonzero observations in the variable to forecast. A missing value is produced for the forecast.

- 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 `_LEAD $n$ _` variable is related to the multistep forecast (n ranges from one to the value of the horizon).

If the forecast step fails for a particular variable, the variables related to forecasting are set to missing.

RECFOR= Data Set

The RECFOR= data set contains the values of the reconciled forecasts. 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
400	An unlocked equality constraint has been imposed.
500	A locked equality constraint has been imposed.
600	A lower bound is active.
700	An upper bound is active.
1000	ID value out of 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 STDMETHOD= option that was provided by the user has been changed for this observation.
9000	The CLMETHOD= option that was provided by the user has been changed for this observation.
10000	The standard error hit the limits that are imposed by the STDDIFBD= option.
11000	Multiple warnings have been printed to the log for this observation.
12000	The number of missing values in the STD variable in the DISAGGDATA= data set is different from the number of missing values in the union of the PREDICT and ACTUAL variables.
13000	The solution might be suboptimal. This means that the optimizer did not find an optimal solution, but the solution that was provided satisfies all constraints.
14000	A failed forecast “.F” has been detected in a relevant input variable.

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 or SAS Forecast Studio for Desktop 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 or SAS Forecast Studio for Desktop 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 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 also contains 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 or SAS Forecast Studio for Desktop 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 or SAS Forecast Studio for Desktop.
- 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 4

Statistics of Fit

What Are Statistics of Fit?	267
Understanding the Formulas and Definitions	268
Descriptions of Statistics of Fit	268
Additional Information	271
Descriptive Statistics	271

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.

Understanding the Formulas and Definitions

To understand the formulas and definitions for the statistics of fit, review the following information:

- n is the number of nonmissing observations.
 - k is the number of fitted parameters in the model.
 - $APE = 1100 * (y_t - \hat{y}_t)/y_t$ is the absolute percent error.
 - $ASPE = 1100 * (y_t - \hat{y}_t)/0.5(y_t + \hat{y}_t)$ is the absolute symmetric percent error.
 - $APPE = 1100 * (y_t - \hat{y}_t)/\hat{y}_t$ is the absolute predictive percent error.
 - $RAE = |(y_t - \hat{y}_t)/(y_t - y_{t-1})|$ is relative absolute error.
-

Descriptions of Statistics of Fit

The definitions and formulas for the statistics of fit that are available in SAS Forecast Studio and or SAS Forecast Studio for Desktop are described below. 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.

Adjusted R-squared (ADJRSQ)

The adjusted R^2 statistic, $1 - (\frac{n-1}{n-k})(1 - R^2)$.

Akaike Information Criterion (AIC)

Akaike's information criterion, $n \ln(\frac{SSE}{n}) + 2k$.

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-squared (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, $100\max(\frac{(y_t - \hat{y}_t)}{y_t})$. 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, $\frac{1}{n}\sum_{t=1}^n \frac{(y_t - \hat{y}_t)}{y_t}$. 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, $100\min(\frac{(y_t - \hat{y}_t)}{y_t})$. 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, $R^2 = 1 - \frac{SSE}{SST}$. 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), $1 - (\frac{n-1}{n})\frac{SSE}{RWSSE}$, 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, $n * \ln(\frac{SSE}{n}) + k * \ln(n)$.

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: $\sum_{t=1}^n (y_t - \bar{y})^2$, where $\bar{y}$ 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 Forecast Server Procedures: 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 value

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 5

Default Model Selection Lists

Overview of the Default Model Selection Lists	273
Best Smoothing Model	273
Best Nonseasonal Smoothing Model	274
Best Seasonal Smoothing Model	274
Time Series Forecasting Models	274

Overview of the Default Model Selection Lists

Although you can create your own model, SAS Forecast Studio and SAS Forecast Studio for Desktop are shipped 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.HPFDFTL.BESTN catalog. The following models are in this selection list:

- Simple Exponential Smoothing
- Double (Brown) Exponential Smoothing
- Linear (Holt) Exponential Smoothing
- Damped Trend Exponential Smoothing

Best Seasonal Smoothing Model

The model selection list for the Best Seasonal Smoothing Model is saved in the SASHELP.HPFDFTL.BESTS catalog. The following models are in this selection list:

- Additive Seasonal Exponential Smoothing
- Winters Additive Method
- Winters Multiplicative Method

Time Series Forecasting Models

The model selection list for the time series forecasting models is saved in the SASHELP.HPFDFTL.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 6

Autoregressive, Differencing, and Moving Average Orders

Which Models Use Autoregressive, Differencing, and Moving Average Orders?	277
Autoregressive Orders	277
Differencing Orders	278
Moving Average Orders	279
Example: Specifying the Orders for the Airline Model	280

Which 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 and SAS Forecast Studio for Desktop use the syntax from PROC HPFARIMASPEC to specify this order. For more information about the underlying code, see the HPFARIMASPEC procedure in the *SAS Forecast Server Procedures: 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 $p = (1, \dots, \text{integer})$. This syntax creates a model of $P = (1, \dots, \text{integer})$.
- For an autoregressive model with parameters at specified lags, use $p = (\text{integer})$. This syntax creates a model of $P = (\text{integer})$.
- For an autoregressive model with a seasonal component, use $p = (\text{integer}, \dots, \text{integer})$ ($\text{integer} \dots \text{integer}$)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 $p = (integer,...integer)(integer,...integer)(integer,...integer)$. This syntax creates a model of $P = (integer,...integer)(integer,...integer)(integer,...integer)$.

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. 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 $q = \text{integer}$. This syntax creates a model of $Q = (1, \dots, \text{integer})$.
- For a moving average model with parameters at specified lags, use $q = (\text{integer})$. This syntax creates a model of $Q = (\text{integer})$.
- For a moving average model with a seasonal component, use $q = (\text{integer}, \dots, \text{integer}) (\text{integer}, \dots, s)$, 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 $Q = (\text{integer}, \dots, \text{integer}) (\text{integer}, \dots, \text{integer}) s$.
- For a moving average model with multiple factors, use $q = (\text{integer}, \dots, \text{integer}) (\text{integer}, \dots, \text{integer})$. This syntax creates a model of $Q = (\text{integer}, \dots, \text{integer}) (\text{integer}, \dots, \text{integer}) s$.

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

Model Description	Value of q	Result in Details Pane
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 and SAS Forecast Studio for Desktop, 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	$d = (1\ 12)$	$q = (1)(1)12$	$D = (1\ 12)\ Q = (1)(12)$
Alternative way to specify monthly seasonality	$d = (1\ 12)$	$q = (1)(12)$	$D = (1\ 12)\ Q = (1)(12)$

Appendix 7

Sample Reports in SAS Forecast Studio

About the Sample Reports	281
Getting Started Reports	281
Time Series Reports	283
Statistical Model Reports	291
Joint Reports	296
Import and Export Reports	296
Quality Reports	297
Reconciliation Reports	299
Transactional Statistics Reports	301
Overrides Reports	302
Final Forecast Reports	302
Reporting Variables Reports	303

About the Sample Reports

SAS Forecast Server is shipped with several sample reports. These reports are organized by category. Your SAS Forecast Server administrator must configure SAS Forecast Studio to use these sample reports. For more information, see the *SAS Forecast Server: Administrator's Guide*.

Note: Reports are available only if your site licenses SAS Forecast Server. These reports are not available if your site licenses SAS Forecasting for Desktop. For more information, see [Chapter 2, “About SAS Forecast Studio for Desktop,”](#) on page 7.

Getting Started Reports

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.

Table A7.1 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.

Table A7.2 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 <i>SAS Forecast Server: Administrator's Guide</i> .
Using Libraries	Shows how libraries that are not pre-assigned can be used in reports. This sample includes a demonstration of the SAS Forecast Server macro variables.
Using Package Files	Shows how you include extra files in the results package. This sample demonstrates how to use SAS Forecast Server macros to manually include files in the results package.
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

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 A7.3 *Exploration*

Name of Report	Description
Branch Series Plot	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.
Level Descriptive Statistics Table	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
Node Series Plot	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.

Name of Report	Description
Path Cross Series Plot	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. <i>Note:</i> To run this project, you must have at least one independent variable in the project.
Node Cross Series Plot	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. <i>Note:</i> To run this project, you must have at least one independent variable in the project.
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.

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 A7.4 *Statistical Analysis*

Name of Report	Description
Level Seasonal Cycles Plot	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
Node X12-ARIMA Analysis Report	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
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

Name of Report	Description
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
Node Autocorrelation Analysis Report	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

Name of Report	Description
Node Cross Correlation Analysis Report	Creates a cross-correlation analysis report for the selected node of the hierarchy. When you run this report, you can customize the following parameters: • the title of the report • the independent variable to include in the report • the cross-correlation statistics • the number of lags • the number of parameters • the series and the cross series transformation • the simple differencing order • the seasonal differencing order • how to interpret missing values • how to interpret zero values

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
Node Time Series Analysis Report (cont.)	• 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.

Name of Report	Description
Node Decomposition Analysis Report	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
Node Stepwise Regression Analysis Report	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

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

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.

Table A7.5 *Evaluation*

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.

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.

Table A7.6 Forecast

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 the same path of the hierarchy. When you run this report, you can customize the title of the report.
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.

Name of Report	Description
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 which 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.

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.

Table A7.7 *Parameter Analysis*

Name of Report	Description
Level Parameter Estimates Univariate Plot	Creates a report that contains a univariate plot of the parameter estimates. These estimates 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

Name of Report	Description
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, <i>t</i>-values, and <i>p</i>-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, <i>t</i>-values, and <i>p</i>-values of the parameter estimates. When you run this report, you can customize the title of the report.
Level Event Parameter Estimates Report	Creates a report that lists the value of the parameter estimate, the standard error of the parameter estimate, the <i>T</i>-values of the parameter estimate, and the <i>p</i>-values of the parameter estimate for the selected event. When you run this report, you can customize its title. You can also specify whether to include a univariate plot in the results (when appropriate).
Level Input Parameter Estimates Report	Creates a report that lists the value of the parameter estimate, the standard error of the parameter estimate, the <i>T</i>-values of the parameter estimate, and the <i>p</i>-values of the parameter estimate for the selected independent variable. When you run this report, you can customize its title. You can also specify whether to include a univariate plot in the results (when appropriate).

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.

Table A7.8 *Component*

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.

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.

Table A7.9 *Summary*

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.

Table A7.10 Joint Reports

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.

Table A7.11 Import and Export

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.

Name of Report	Description
Export as Data Set	Exports the current project to a SAS data set. You have the option of registering this data set in the metadata. When you run this report, you must specify the following information: • whether the generated data is included in the report results • whether to register the generated data set in the metadata • the library in which to store the exported data • the name of the data set for the exported data
Export as Information Map	Exports the data sets in the project to a SAS Information Map. When you run this report, you must specify the metadata library and data set for the exported data.
Export as OLAP Cube	Exports the forecast 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 worksheet to a SAS data set. When you run this report, you must specify the location of this worksheet.

Quality Reports

The following table gives a brief explanation of the quality reports and lists the parameters that you can customize at run time.

Table A7.12 Quality Reports

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.

Name of Report	Description
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

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.

Table A7.13 *Forecast*

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.

Name of Report	Description
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.

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.

Table A7.14 Evaluation

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.

Name of Report	Description
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, and 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.
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.

Overrides 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.

Table A7.15 *Override*

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 events.

Final Forecast Reports

These reports contain information about 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.

Table A7.16 *Final Forecast*

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 the 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.

Table A7.17 *Reporting Variables*

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.

Recommended Reading

Here is the recommended reading list for this title:

- the online Help for SAS Forecast Studio
- *SAS Forecast Server: Administrator's Guide*
- *SAS/ETS User's Guide*

SAS also publishes the following books, which are companions to the *SAS/ETS User's Guide* :

- *SAS/ETS Software: Applications Guide 1, Version 6, First Edition* discusses features of SAS/ETS software for time series modeling and forecasting, financial reporting, and loan analysis.
- *SAS/ETS Software: Applications Guide 2, Version 6, First Edition* discusses features of SAS/ETS software for econometric modeling and simulation.
- *SAS Forecast Server Procedures: User's Guide*

The recommended reading list from SAS Press includes the following titles:

- *Demand-Driven Forecasting: A Structured Approach to Forecasting*
- *The Business Forecasting Deal: Exposing Myths, Eliminating Bad Practices, Providing Practical Solutions*
- *Applied Data Mining for Forecasting Using SAS*

For a complete list of SAS publications, go to sas.com/store/books. If you have questions about which titles you need, please contact a SAS Representative:

SAS Books
SAS Campus Drive
Cary, NC 27513-2414
Phone: 1-800-727-0025
Fax: 1-919-677-4444
Email: sasbook@sas.com
Web address: sas.com/store/books

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

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.

aggregation

the process of combining more than one time series to form a single series within the same time interval. For example, data can be combined 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.

autocorrelation

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

autocorrelation function plot

See 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.

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.

dummy variable

See indicator variable.

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 intervals into the future, beyond a base date, for which analyses and predictions are made.

IACF plot

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

indicator variable

a numeric variable with a value of either 1 or 0 that is used to indicate whether or not unusual events occur. The variable takes the value of 1 during the event and 0 otherwise.

inverse autocorrelation

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

inverse autocorrelation function plot

See IACF plot.

ISO year

a year as specified by ISO 8601, and represented by a minimum of four digits for the year (YYYY), followed by digits that represent the week number and the weekday number. For example, 2006-W52-7 (or 2006W527) is the Sunday of the 52nd week of 2006.

leap week

the 53rd week that can occur when an ISO year is measured in weeks. Normally, a year contains 52 weeks of seven days each, for a total of 364 days. However, some years can have an extra one or two days, for a total of 365 or 366 days.

Approximately every six years, seven extra days have accumulated, and these extra days are called a leap week.

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.

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.

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.

PACF plot

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

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

See PACF plot.

performance statistic

a statistic of fit that is used to determine how well a forecasting process is predicting the future.

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

a statistical technique designed to even out seasonal swings in time series data in order to analyze nonseasonal trends.

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 $\hat{y} = f(x)$ performs by comparing the actual data to the predictions.

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.

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.

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.

variable label

up to 256 characters of descriptive text that can be printed in the output by certain procedures instead of, or in addition to, the variable name.

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

A

- accumulation
 - about [57](#)
 - methods [58, 64](#)
- active series [23, 77](#)
- adding
 - adjustment variables [52](#)
 - filters [80](#)
 - notes [125](#)
 - overrides [120](#)
- adjustment variables
 - about [51](#)
 - adding [52](#)
 - examples [52](#)
- aggregation
 - about [57](#)
 - methods [58, 64](#)
- analyzing
 - scenarios [28](#)
 - time series [26](#)
- ARIMA models
 - about [190](#)
 - autoregressive orders [277](#)
 - differencing orders [278](#)
 - events [111](#)
 - moving average orders [279](#)
- autoregressive orders
 - about [277](#)
 - syntax [277](#)

B

- Base SAS [5, 10](#)
- baseline model [174](#)
- batch mode [230](#)
- best nonseasonal smoothing model [274](#)
- best seasonal smoothing model [274](#)
- best smoothing model [273](#)
- bottom-up reconciliation method [56](#)
- BY variables
 - about [49](#)
 - specifying [63](#)
 - troubleshooting [247](#)

C

- combining events [105](#)
- comparing
 - models [179](#)
 - scenarios [215](#)
- component series data set [64](#)
- confidence limit [75](#)
- copying
 - events [109](#)
 - models [183](#)
 - scenarios [216](#)
- creating
 - ARIMA models [190](#)
 - curve fitting models [208](#)
 - events [101](#)
 - exponential smoothing models [201](#)
 - intermittent demand models [203](#)
 - models [181](#)
 - moving average models [207](#)
 - multiple regression models [205](#)
 - projects [62](#)
 - random walk models [210](#)
 - reports [92](#)
 - scenarios [214](#)
 - Subset (factored) ARIMA models [194](#)
 - unobserved components models [197](#)
- curve fitting models [208](#)
- custom models
 - creating from default models [181](#)
 - model repository [170](#)
- custom time intervals [45](#)
- customizing
 - Forecasting View [29](#)
 - table view [29](#)

D

- data sets
 - forecast [76](#)
 - preparing [64](#)
 - selecting for project [63](#)
 - troubleshooting formats [247](#)
- data table
 - adding overrides [120](#)

- copying 24
- Forecasting View 23
- data view 67
- default models
 - model repository 170
- dependent variables 49
- descriptive statistics 271
- differencing orders
 - about 277
 - example 280
 - syntax 278

E

- editing
 - models 182
- event occurrence 101
- event recurrence 101
- event types 100
- events
 - about 99
 - annual promotion 102
 - ARIMA models 111
 - combining 105
 - copying 109
 - creating 64, 101
 - data table 115
 - editing 109
 - effects 115
 - exporting 110
 - forecast plot 115
 - holiday 104
 - importing 109
 - including in model generation 110
 - interpretation by SAS 110
 - multiple regression models 114
 - status 111
 - Subset (factored) ARIMA models 112
 - unobserved components models 113
- exception rules 79
- exponential smoothing models 201
- exporting
 - events 110
 - models 181

F

- Failed Forecasts panel 21
- filters
 - about 79
 - adding 80
 - examples 81
 - multiple expressions 83
 - simple 81
 - table view 20
 - viewing results 80

- FINALFOR= data set 256
- forecast data sets 76
- forecast model 171, 174
- forecast plots 22
- Forecast Summary dialog box 76
- forecasting settings 64
- Forecasting View
 - components 22
 - customizing 29
- forecasts
 - final forecasts 23
 - forecasts options 75
 - processes for generating 73
 - reforecasting data 78
 - start date 64
 - summary 76
 - viewing 22
- format of input data 35
- Foundation SAS 5, 10
- future values 38

G

- graphs
 - copying 24

H

- hierarchy
 - about 55
 - creating 63
 - example 60
 - reconcile 63
 - reconciling 131, 136
- hierarchy view 19
- historical data 23
- holdout sample 175
- horizon 64, 75

I

- importing
 - events 109
 - models 180
- inactive series 77
- independent variables 50
 - forecast data sets 76
- input data set
 - future values 38
 - required format 35
- input variables
 - editing from input table 214
 - input calculator 215
 - resetting values 216
- intermittency test 74
- intermittent demand models 203

L

linking to projects 70
 locked overrides 118, 125
 log 245

M

middle-out reconciliation method 56
 migrating
 projects 65
 missing values
 filters 85
 interpretation of 36, 64
 model generation
 including events 110
 model repositories 170, 183
 model selection lists
 about 171
 default lists 273
 editing a model 182
 exporting 181
 generating 173
 rejected models 173
 removing models 187
 model selection options 74
 model specifications
 ARIMA 190
 curve fitting 208
 exponential smoothing 201
 intermittent demand 203
 moving average 207
 multiple regression 205
 random walk 210
 Subset (factored) ARIMA 194
 unobserved components 197
 Modeling View 24
 specifying the number of lags 29
 models
 adding to a series 180
 assessing performance 177
 comparing 179
 copying 183
 creating 181
 deleting 187
 editing 182, 183
 exporting 181
 external 170
 fitting 171
 forecast 174
 forecast model 171
 importing 180
 options for project 173
 plots 178
 rejected 173
 selecting forecast model 174
 selection criterion 175

stored 170

tables 179

moving average models 207

moving average orders

 about 277

 example 280

 syntax 279

multiple regression models

 about 205

 events 114

N

name of project 68

negative forecasts 75

New Project wizard

 additional options 64

 assign data to roles 63

 assign variables to roles 64

 generating forecasts 65

 naming the project 62

 preparing data 64

 saving project code 65

 selecting a data set 63

 specifying the time ID 63

notes 125

O

occurrence 101

options

 diagnostics 74

 forecasts 75, 177

 models 173

out-of-sample region 175, 177

OUT= data set 257

OUTCOMPONENT= data set 64, 257

OUTEST= data set 258

OUTFOR= data set 259

outliers

 about 193

 detecting 74, 193

 removing from model 194

OUTOVRD= data set 260

output data sets

 FINALFOR= 256

 OUT= 257

 OUTCOMPONENT= 257

 OUTEST= 258

 OUTFOR= 259

 OUTOVRD= 260

 OUTSTAT= 260

 OUTSTATSELECT= 262

 OUTSUM= 262

 RECFOR= 263

 RECSTAT= 265

- RECSUM= 265
- SCENARIO 266
- OUTSTAT= data set 260
- OUTSTATSELECT= data set 262
- OUTSUM= data set 262
- Override Calculator 122
- override conflicts
 - about 126
 - resolving 128
- overrides
 - about 117
 - adding 120
 - basing on existing value 122
 - branches in the hierarchy 124
 - considerations when specifying 118
 - data table 23
 - filters 124
 - locking 125
 - Override Calculator 122
 - processing 118
 - removing 126
 - splitting value proportionally 123
 - types 118
 - unlocking 125
- overview panel
 - failed forecasts 21
 - hierarchy view 19
 - table view 20, 29

P

- plots
 - forecast 22
 - models 178
 - time series 163
- project code
 - batch file 230
 - saving 229, 230
- project data
 - updating 67
- project directory 61
- project name 62
- project properties
 - data source 68
 - name 68
 - storage location 68
- projects
 - about 61
 - creating 62, 66
 - forecast summary 76
 - linking to 70
 - name 62
 - new 62
 - opening 65, 66
 - reforecasting 78
 - running in batch 230

- saving 69
- updating version 65

R

- random walk models 210
- RECFOR= data set 263
- reconciled forecast 23
- reconciliation failures 132
- reconciliation methods 56
- reconciliation options
 - how to set 132
 - understanding 129
- reconciling hierarchy
 - examples 136
 - how to 131
 - reconciliation options 132
- reconciling hierarchy examples
 - change in reconciliation direction 140, 148
 - implied limits have precedence 152
 - no override conflicts 137, 156
 - unlocked overrides 146
 - unreconciled node 143, 159
- RECSTAT= data set 265
- RECSUM= data set 265
- recurrence 101
- reforecasting
 - project 78
 - series 78
- removing overrides 126
- reporting variables 50
- reports
 - about 91
 - creating 92
 - differences between reconciled and unreconciled forecasts 296
 - example 93
 - exporting 296
 - final forecasts 302
 - getting started 281
 - importing 296
 - joint 296
 - overrides 302
 - quality 297
 - reconciliation 299
 - reporting variables 303
 - running 92
 - samples 92, 281
 - statistical model 291
 - time series 283
 - transactional statistics 301
 - viewing 92
- reserved names
 - models 254
 - variables 249

reserved variable names 249
 rolling simulations 177
 running stored processes 92

S

SAS Add-In for Microsoft Office 237
 SAS code
 saving 230
 viewing 229
 SAS Enterprise Guide 238
 SAS Forecast Server
 overview 4
 relation to other products 5, 9
 SAS Forecast Server procedures 5, 9
 SAS Forecast Studio
 about 6
 benefits 4
 Java Web Start client 14
 local installation 14
 log on 14
 SAS Forecast Studio for Desktop
 log on 14
 SAS Forecast Studio tasks
 about 237
 creating a project 239
 opening a project 240
 prerequisites 239
 submitting overrides 240
 SAS Forecasting for Desktop
 overview 8
 SAS log 245
 SAS tasks 237
 SAS/ETS 5, 10
 SAS/GRAPH 6, 10
 saving projects 69
 scenario analysis
 about 213
 example 216
 Scenario Analysis View 28, 213
 SCENARIO data set 266
 scenarios
 analyzing 28
 comparing 215
 copying 216
 creating 214
 running 215
 seasonality test 74
 selection criterion 74
 series
 active 77
 inactive 77
 reforecast 78
 series properties
 statistics 271
 viewing 69

Series View 26
 specifying the number of lags 29
 simulations
 rolling 177
 statistical forecast 23
 statistics of fit
 definition 267
 descriptions 268
 storage location of project 68
 stored processes
 about 91
 parameters 91
 running 92
 Subset (factored) ARIMA models
 about 194
 autoregressive orders 277
 differencing orders 278
 events 112
 moving average orders 279
 summary of the forecast run 76

T

table view
 about 20
 customizing 29
 tables
 models 179
 time series 163
 time ID
 properties 63
 troubleshooting 245
 variables 40
 time intervals
 about 42
 custom 45
 supported intervals 40
 time series
 analyzing 26
 data requirements 34
 decomposition 164, 166
 overview of data 34
 plots 163
 process for creating data 34
 tables 163
 transforming 164
 time series forecasting model 274
 top-down reconciliation method 56
 transforming time series 164

U

unlocked overrides 118, 125
 unobserved components models
 about 197
 events 113

Unobserved Components models
 autoregressive orders [277](#)
 differencing orders [278](#)
 moving average orders [279](#)
 unreconciled nodes
 example [133](#)
 reconciliation failure report [132](#)
 viewing [132](#)


V

variable names, reserved [249](#)
 variables
 adjustment [51](#), [64](#)
 classification [49](#), [63](#)
 dependent [49](#)
 independent [50](#), [64](#)
 reporting [50](#)
 reserved names [249](#)

roles [39](#), [64](#)
 time ID [40](#), [63](#)
 troubleshooting [247](#)
 types [39](#)
 viewing
 forecasts [22](#)
 model plots [178](#)
 models [24](#)
 project code [229](#)
 reports [92](#)
 results from a filter [80](#)
 SAS code [229](#)
 time series plots [163](#)
 time series tables [163](#)

W

workspace [18](#)



Gain Greater Insight into Your SAS® Software with SAS Books.

Discover all that you need on your journey to knowledge and empowerment.

 support.sas.com/bookstore
for additional books and resources.


THE POWER TO KNOW.

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration. Other brand and product names are trademarks of their respective companies. © 2013 SAS Institute Inc. All rights reserved. S107969US.0613

