THE POWER TO KNOW.

SAS/ETS® 14.2 User's Guide
Specifying Forecasting Models

This document is an individual chapter from *SAS/ETS® 14.2 User's Guide*.

The correct bibliographic citation for this manual is as follows: SAS Institute Inc. 2016. *SAS/ETS® 14.2 User's Guide*. Cary, NC: SAS Institute Inc.

SAS/ETS® 14.2 User's Guide

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

November 2016

Chapter 58
Specifying Forecasting Models

Contents

This chapter explores the tools available through the Develop Models window for investigating the properties of time series and for specifying and fitting models. The first section shows you how to diagnose time series properties in order to determine the class of models appropriate for forecasting series with such properties. Later sections show you how to specify and fit different kinds of forecasting models.

Series Diagnostics

The series diagnostics tool helps you determine the kinds of forecasting models that are appropriate for the data series so that you can limit the search for the best forecasting model. The series diagnostics address these three questions: Is a log transformation needed to stabilize the variance? Is a time trend present in the data? Is there a seasonal pattern to the data?

The automatic model fitting process, which you used in the previous chapter through the Automatic Model Fitting window, performs series diagnostics and selects trial models from a list according to the results. You can also look at the diagnostic information and make your own decisions as to the kinds of models appropriate for the series. The following example illustrates the series diagnostics features.

Select “Develop Models” from the Time Series Forecasting window. Select the library SASHELP, the data set CITIMON, and the series RCARD. This series represents domestic retail sales of passenger cars. To look at this series, select “View Series” from the Develop Models window. This opens the Time Series Viewer window, as shown in Figure 58.1.

Figure 58.1 Automobile Sales Series

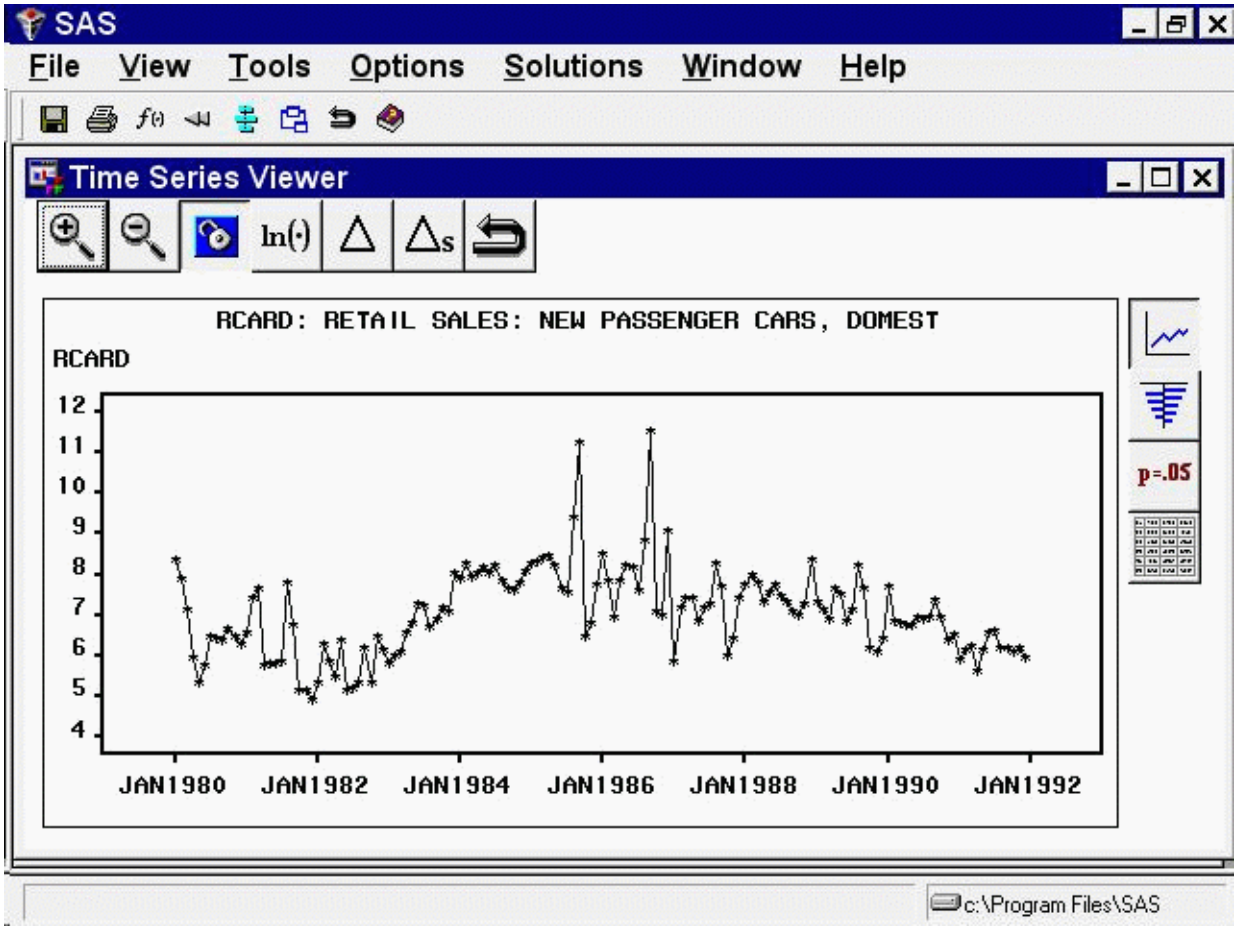

Select “Diagnose Series” from the Tools menu. You can do this from the Develop Models window or from the Time Series Viewer window. Figure 58.2 shows this from the Develop Models window.

Figure 58.2 Selecting Series Diagnostics

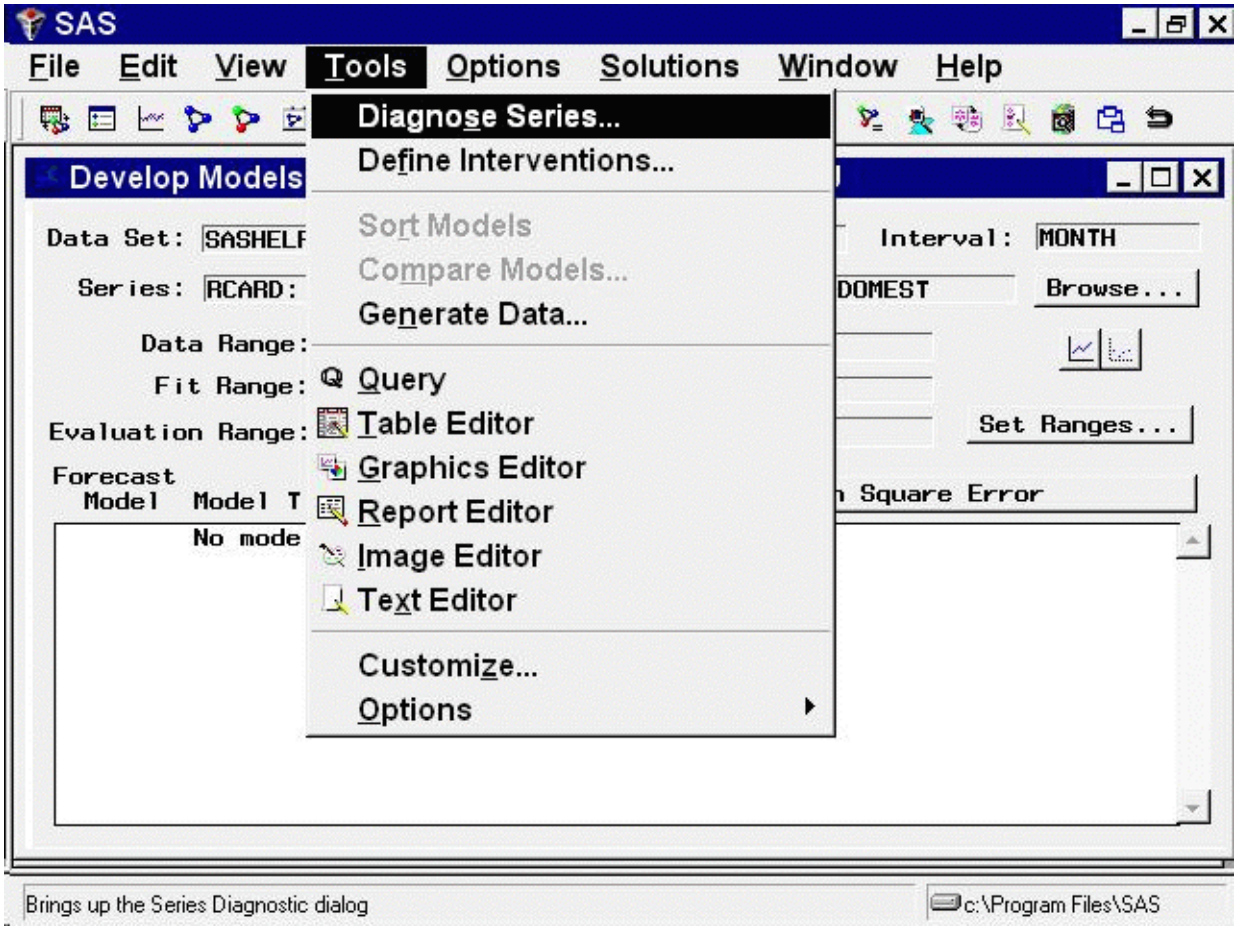

This opens the Series Diagnostics window, as shown in Figure 58.3.

Figure 58.3 Series Diagnostics Window

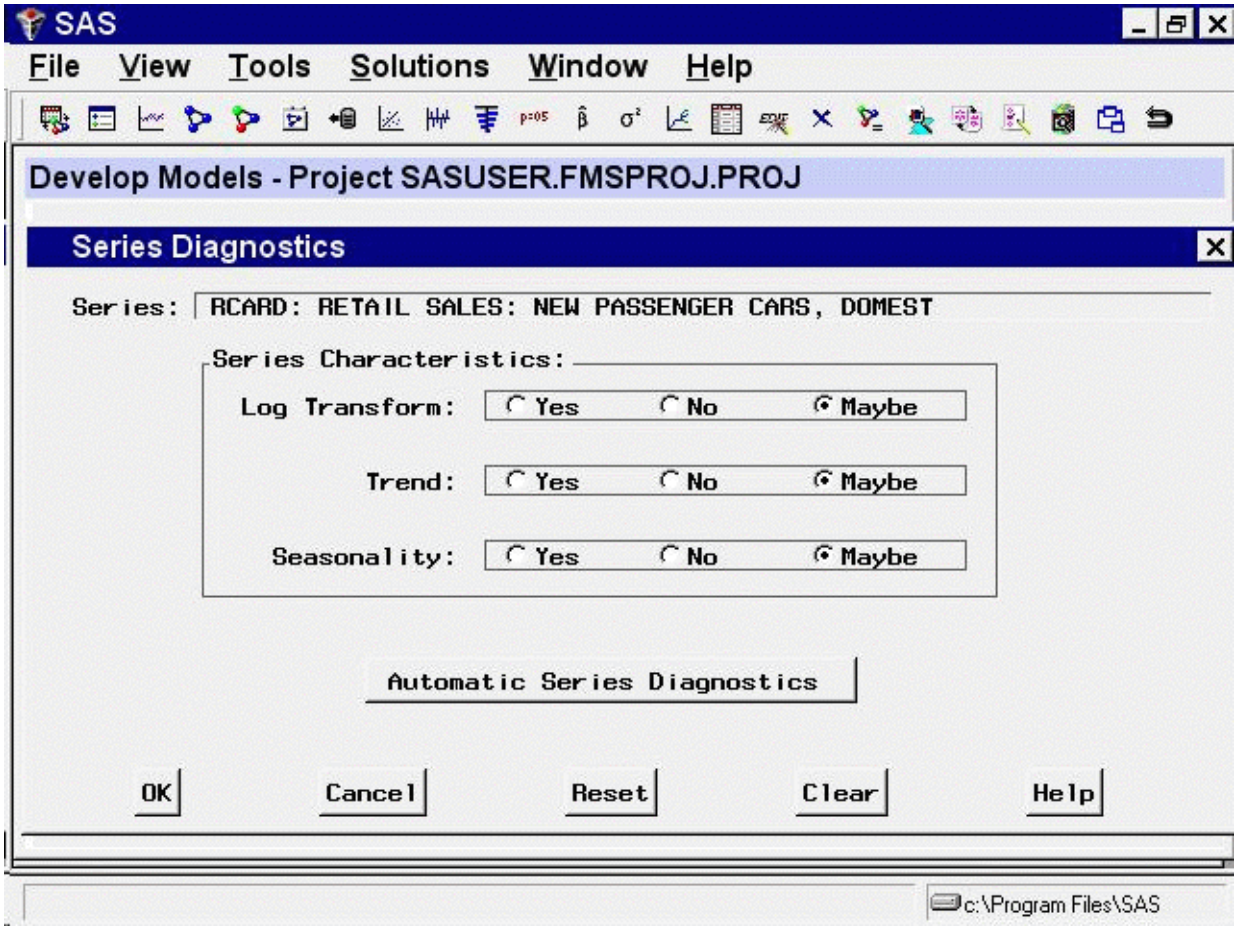

Each of the three series characteristics—need for log transformation, presence of a trend, and seasonality—has a set of options for `Yes`, `No`, and `Maybe`. `Yes` indicates that the series has the characteristic and that forecasting models fit to the series should be able to model and predict this behavior. `No` indicates that you do not need to consider forecasting models designed to predict series with this characteristic. `Maybe` indicates that models with and without the characteristic should be considered. Initially, all these values are set to `Maybe`.

To have the system diagnose the series characteristics, click the Automatic Series Diagnostics button. This runs the diagnostic routines described in Chapter 63, "Forecasting Process Details," and sets the options according to the results. In this example, `Trend` and `Seasonality` are changed from `Maybe` to `Yes`, while `Log Transform` remains set to `Maybe`.

These diagnostic criteria affect the models displayed when you use the Models to Fit window or the Automatic Model Selection model-fitting options described in the following section. You can set the criteria manually, according to your judgment, by selecting any of the options, whether you have used the Automatic Series Diagnostics button or not. For this exercise, leave them as set by the automatic diagnostics. Click the OK button to close the Series Diagnostics window.

Models to Fit Window

As you saw in the previous chapter, you can select models from a list. Invoke the Models to Fit window by clicking the middle of the table and selecting "Fit Models from List" from the menu. This can also be selected from the toolbar or the Fit Model submenu of the Edit menu. The Models to Fit window comes up, as shown in Figure 58.4.

Figure 58.4 Models to Fit Window

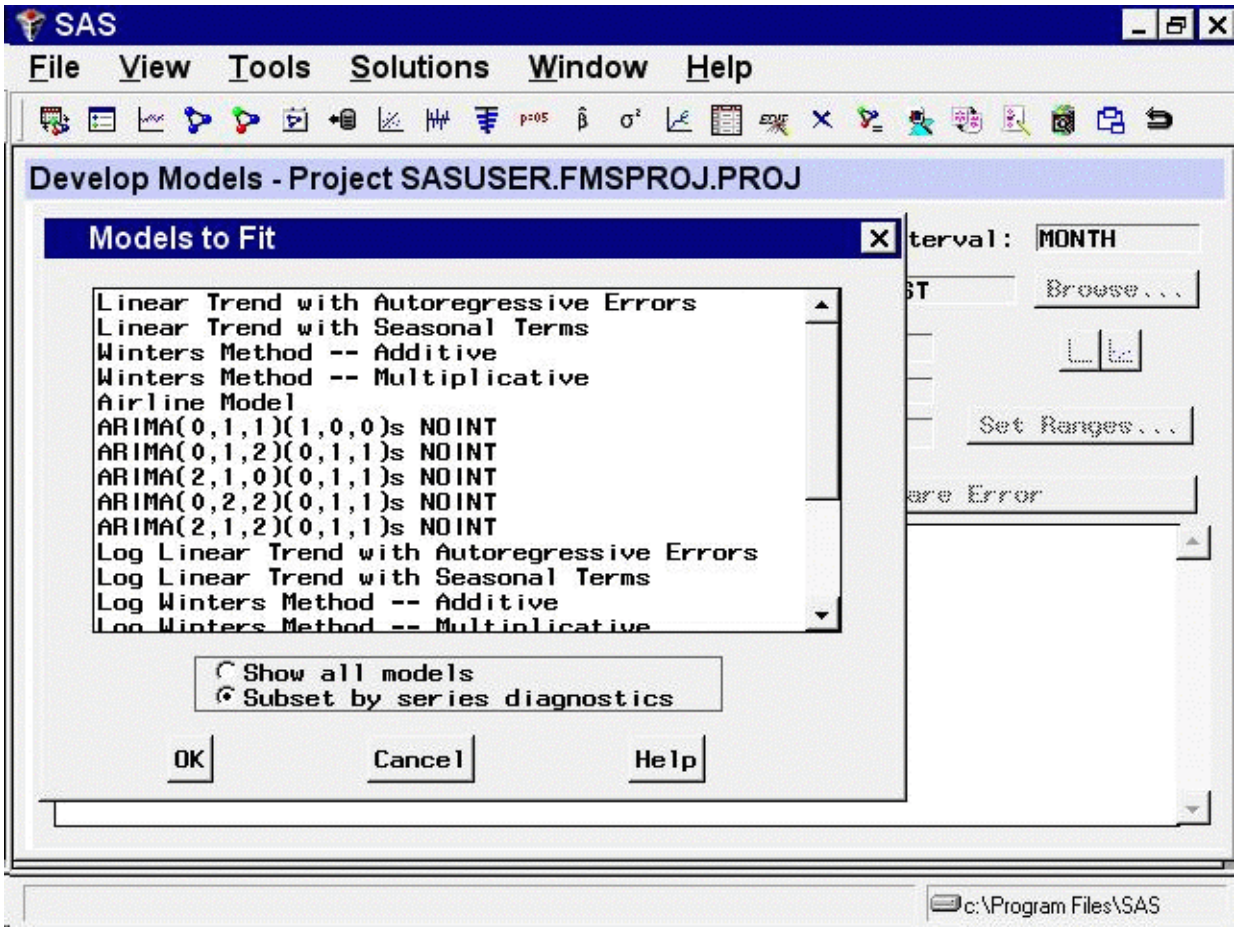

Since you have performed series diagnostics, the models shown are the subset that fits the diagnostic criteria.

Suppose you want to consider models other than those in this subset because you are undecided about including a trend in the model. Select the Show all models option. Now the entire model selection list is shown. Scroll through the list until you find `Log Seasonal Exponential Smoothing`, as shown in Figure 58.5.

Figure 58.5 Selecting a Model from List

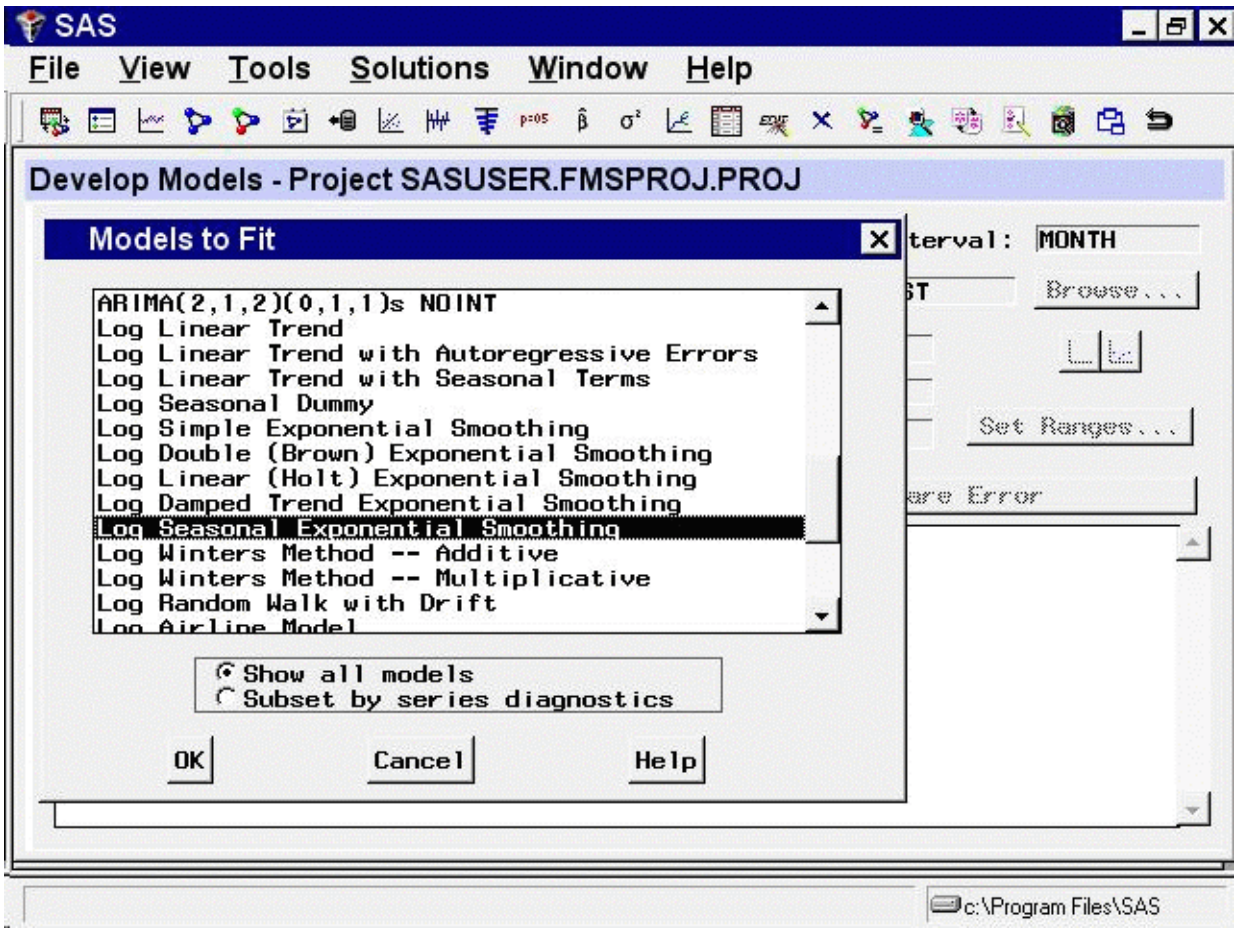

This is a nontrended model, which seems a good candidate. Select this model, and then click the OK button. The model is fit to the series and then appears in the table with the value of the selected fit criterion, as shown in Figure 58.6.

Figure 58.6 Develop Models Window Showing Model Fit

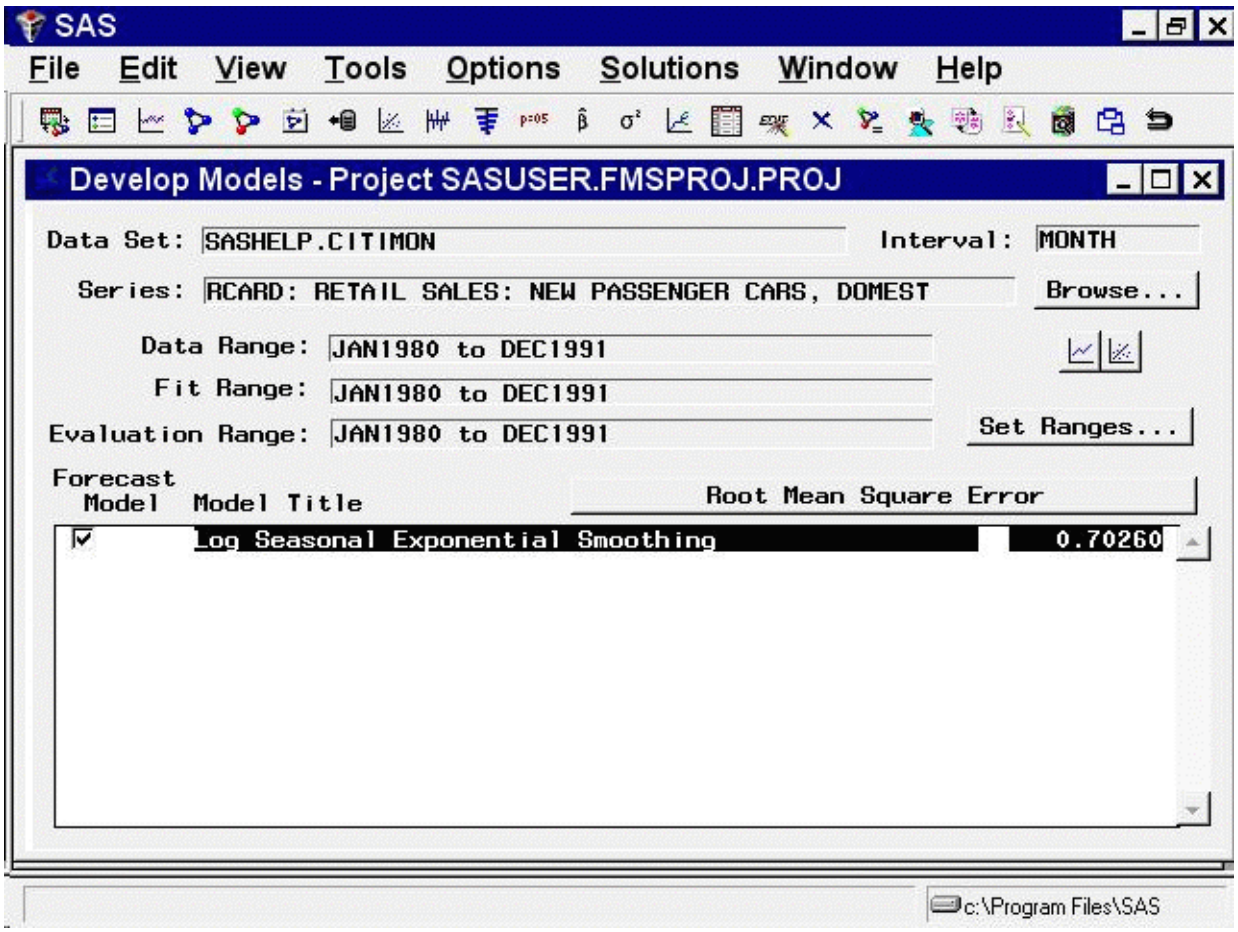

You can edit the model list that appears in the Models to Fit window by selecting “Options” and “Model Selection List” from the menu bar or by selecting the Edit Model List toolbar icon. You can then delete models you are not interested in from the default list and add models using any of the model specification methods described in this chapter. When you save your project, the edited model selection list is saved in the project file. In this way, you can use the Select from List item and the Automatic Model Selection item to select models from a customized search set.

Automatic Model Selection

Automatic model selection is equivalent to choosing Select from List, as you did in the preceding section, fitting all the models in the subset list and then deleting all except the best fitting of the models. If series diagnostics have not yet been done, they are performed automatically to determine the model subset to fit. If you set the series diagnostics for log, trend, or seasonal criteria manually using the radio buttons, these choices are honored by the automatic fitting process.

Using automatic selection, the system does not pause to warn you of model fitting errors, such as failure of the estimates to converge (you can track these using the audit trail feature).

By default, only the best fitting model is kept. However, you can control the number of automatically fit models retained in the Develop Models list, and the following example shows how to do this.

From the menu bar, select "Options" and "Automatic Fit." This opens the Automatic Model Selection Options window. Click the `Models to Keep` list arrow, and select "All models," as shown in Figure 58.7. Now click OK.

Figure 58.7 Selecting Number of Automatic Fit Models to Keep

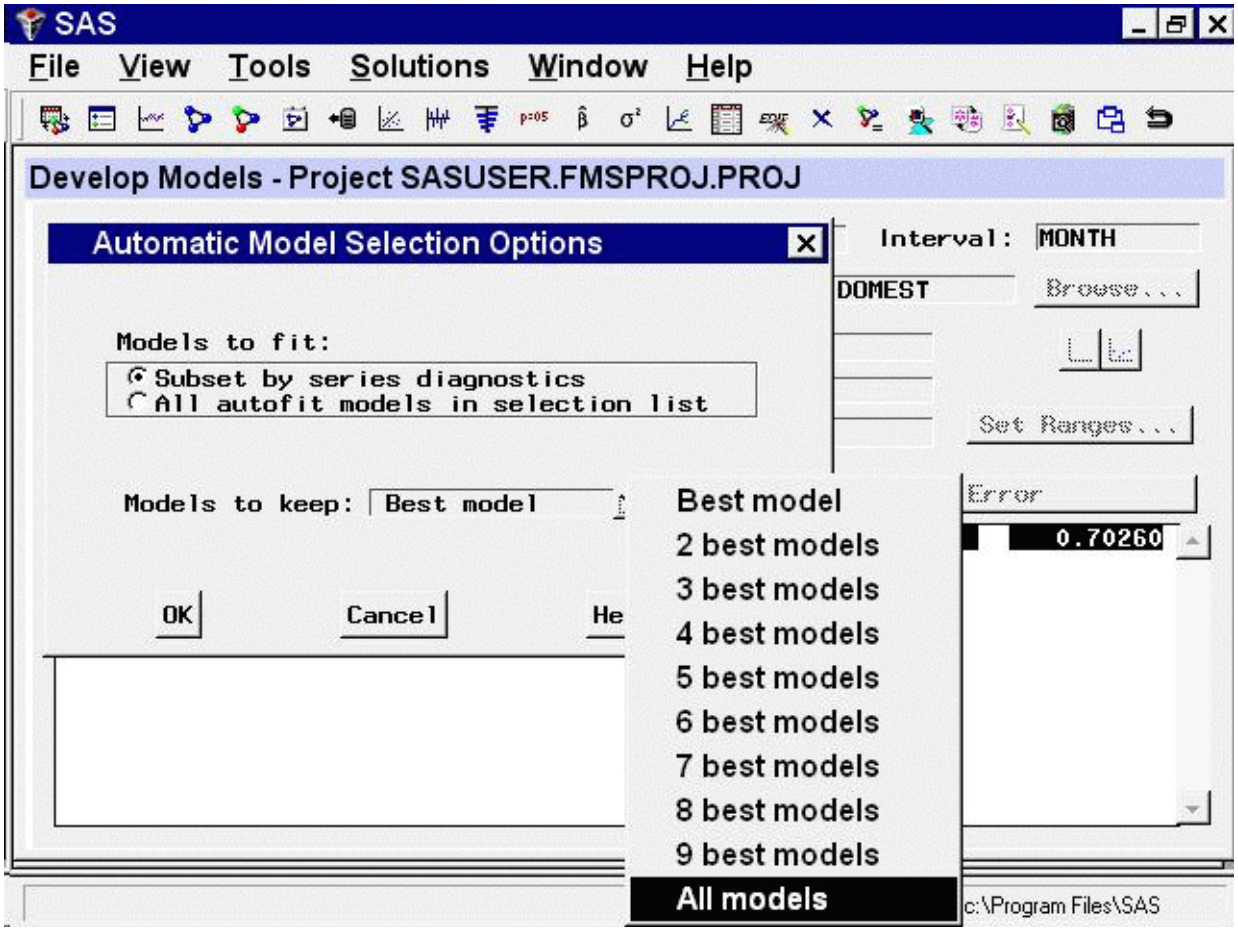

Next, select “Fit Models Automatically” by clicking the middle of the table or using the toolbar or Edit menu. The Automatic Model Selection window appears, showing the diagnostic criteria in effect and the number of models to be fit, as shown in Figure 58.8.

Figure 58.8 Automatic Model Selection Window

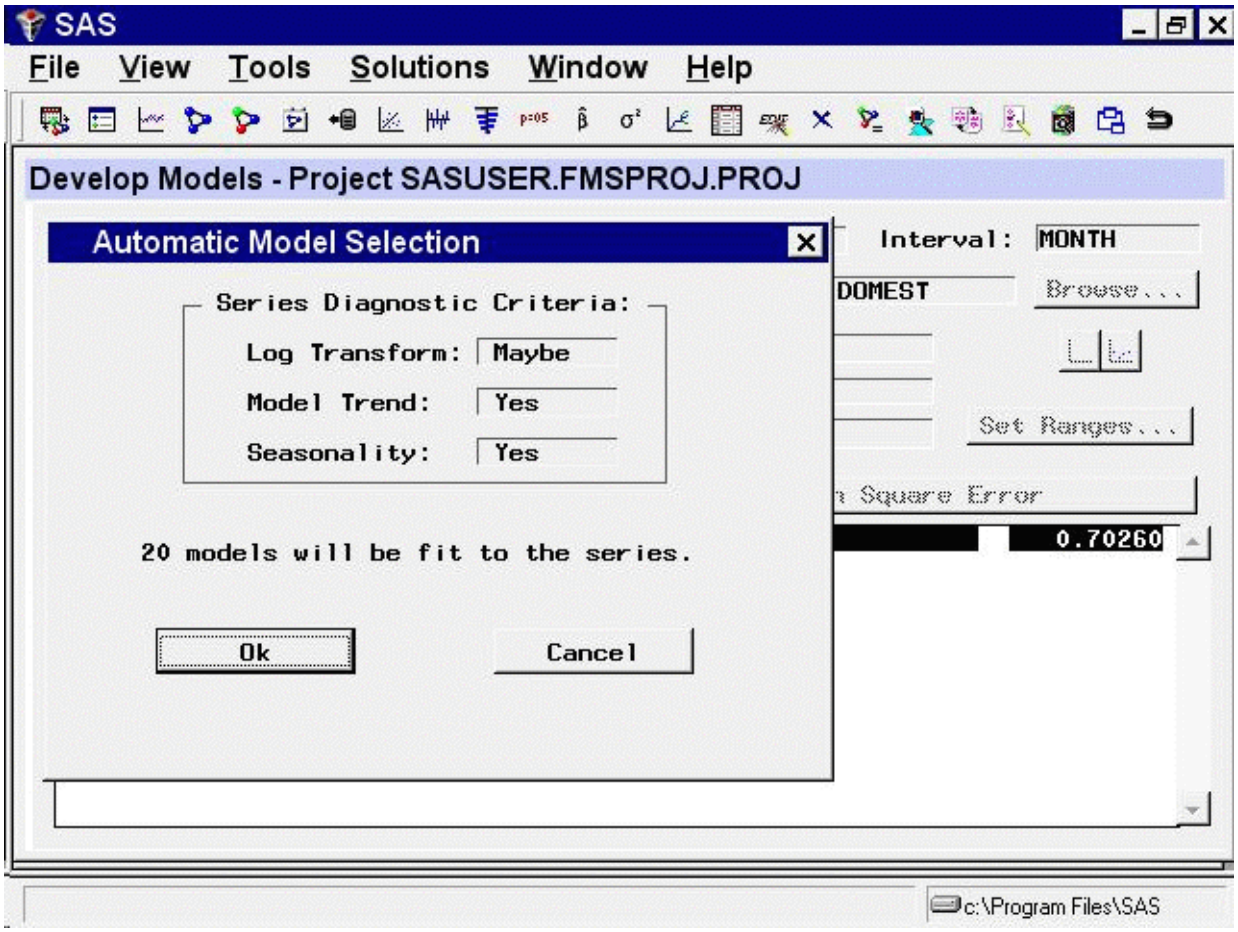

Click the OK button. After the models have been fit, all of them appear in the table, in addition to the model which you fit earlier, as shown in Figure 58.9.

Figure 58.9 Automatically Fit Models

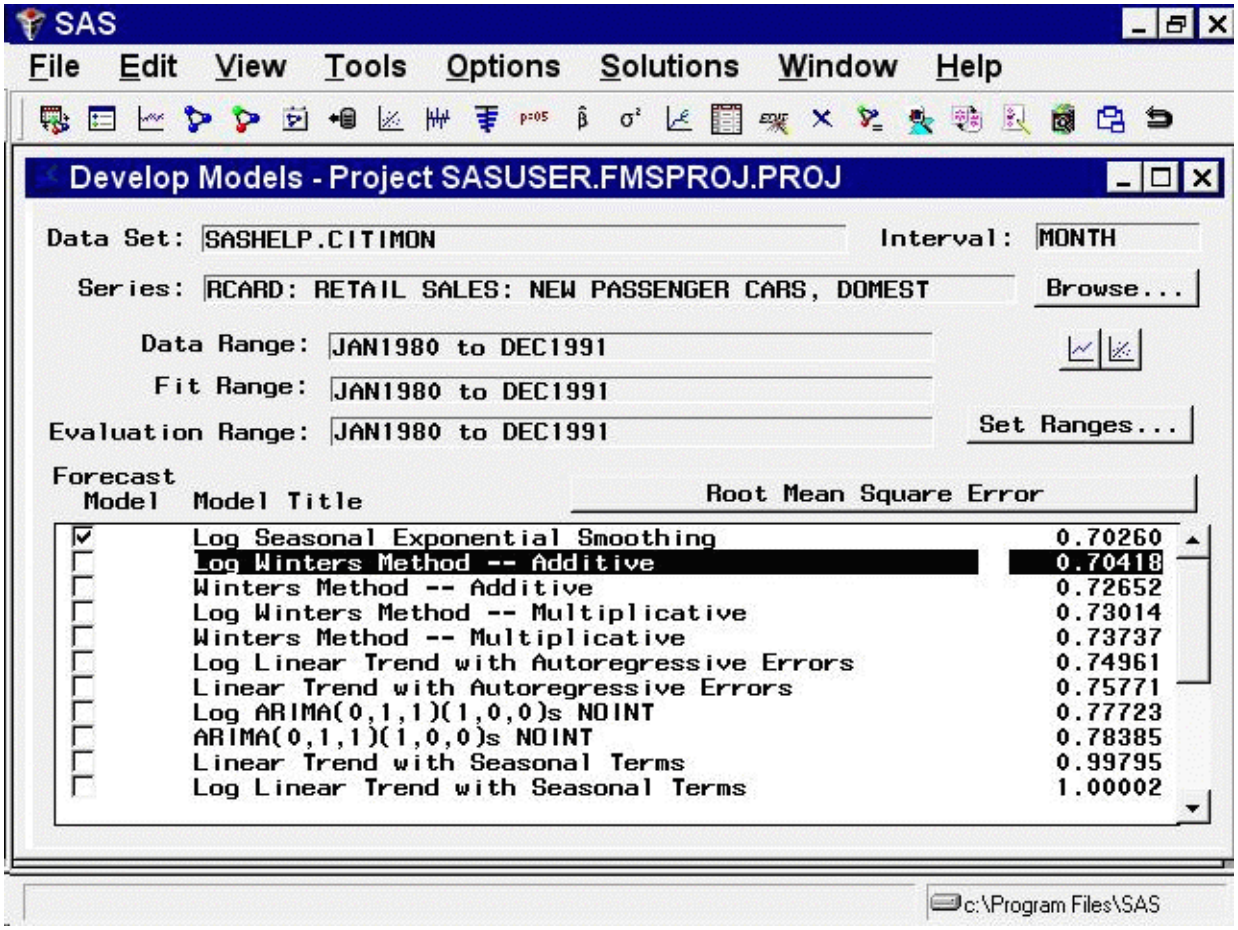

Smoothing Model Specification Window

To fit exponential smoothing and Winters models not already provided in the Models to Fit window, select "Fit Smoothing Model" from the pop-up menu or toolbar or select "Smoothing Model" from the Fit Model submenu of the Edit menu. This opens the Smoothing Model Specification window, as shown in Figure 58.10.

Figure 58.10 Smoothing Model Specification Window

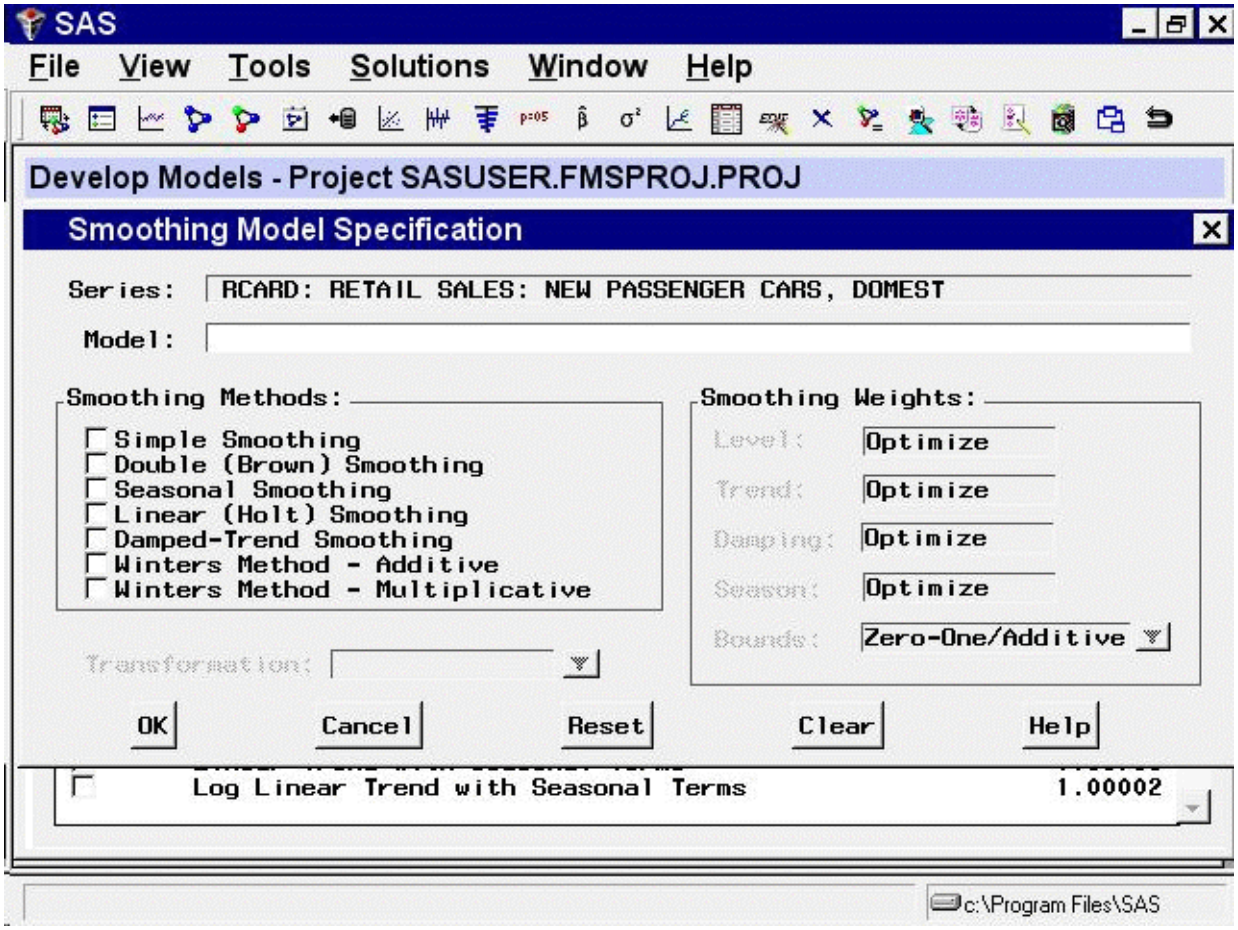

The Smoothing Model Specification window consists of several parts. At the top is the series name and a field for the label of the model you are specifying. The model label is filled in with an automatically generated label as you specify options. You can type over the automatic label with your own label for the model. To restore the automatic label, enter a blank label.

The `Smoothing Methods` box lists the different methods available. Below the `Smoothing Methods` box is the Transformation field, which is used to apply the smoothing method to transformed series values.

The `Smoothing Weights` box specifies how the smoothing weights are determined. By default, the smoothing weights are automatically set to optimize the fit of the model to the data. For more information about how the smoothing weights are fit, see Chapter 63, "Forecasting Process Details."

Under smoothing methods, select "Winters Method – Additive." Notice the smoothing weights box to the right. The third item, `Damping`, is grayed out, while the other items, `Level`, `Trend`, and `Season`, show the word `Optimize`. This tells you that these three smoothing weights are applicable to the smoothing method that you selected and that the system is currently set to optimize these weights for you.

Next, specify a transformation using the `Transformation` list. A menu of transformation choices pops up, as shown in Figure 58.11.

Figure 58.11 Transformation Options

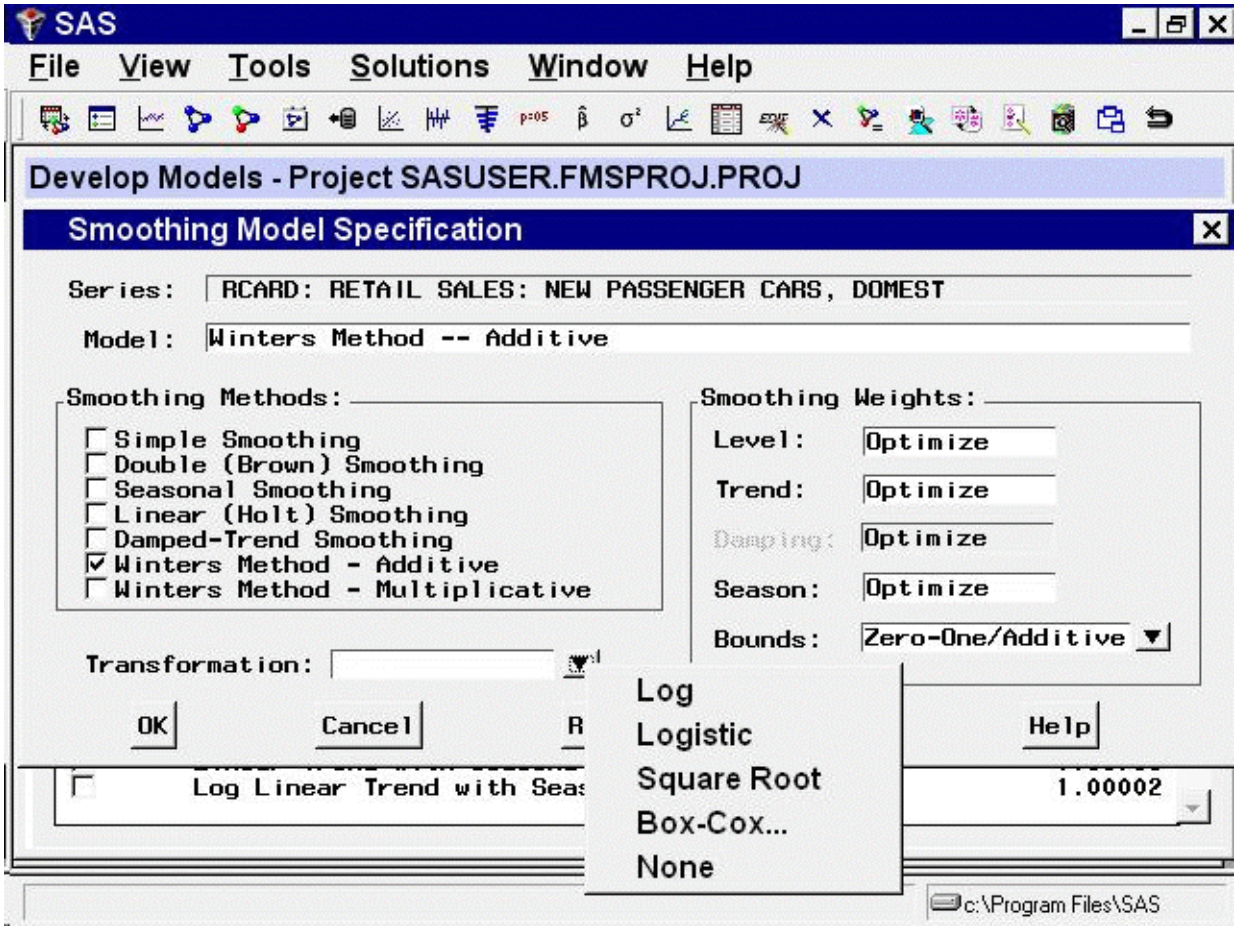

You can specify a logarithmic, logistic, square root, or Box-Cox transformation. For this example, select "Square Root" from the list. The Transformation field is now set to Square Root.

This means that the system will first take the square roots of the series values, apply the additive version of the Winters method to the square root series, and then produce the predictions for the original series by squaring the Winters method predictions (and multiplying by a variance factor if the Mean Prediction option is set in the Forecast Options window). For more information about predictions from transformed models, see Chapter 63, "Forecasting Process Details."

The Smoothing Model Specification window should now appear as shown in Figure 58.12. Click the OK button to fit the model. The model is added to the table of fitted models in the Develop Models window.

Figure 58.12 Winter's Method Applied to Square Root Series

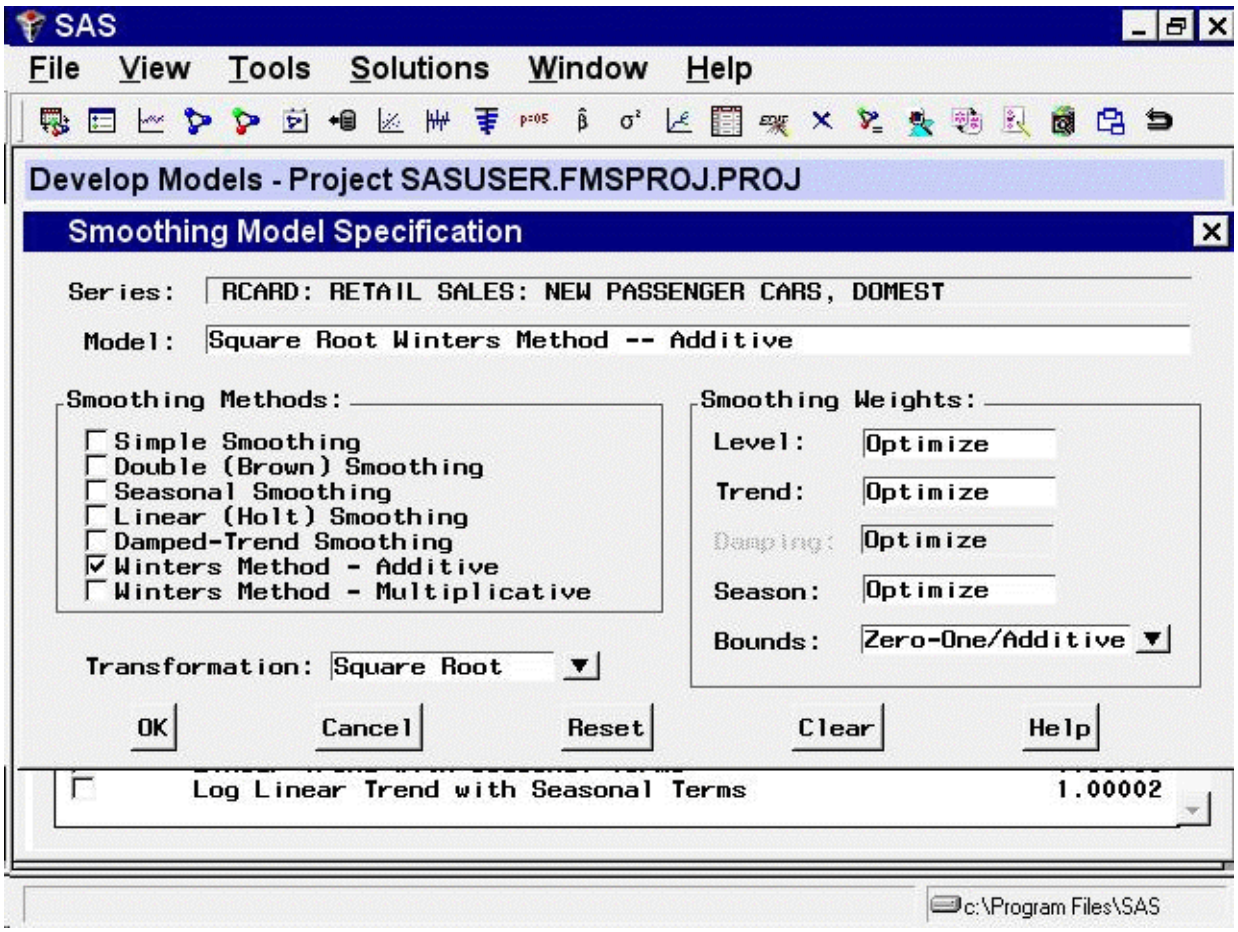

ARIMA Model Specification Window

To fit ARIMA or Box-Jenkins models not already provided in the Models to Fit window, select the ARIMA model item from the pop-up menu, toolbar, or Edit menu. This opens the ARIMA Model Specification window, as shown in Figure 58.13.

Figure 58.13 ARIMA Model Specification Window

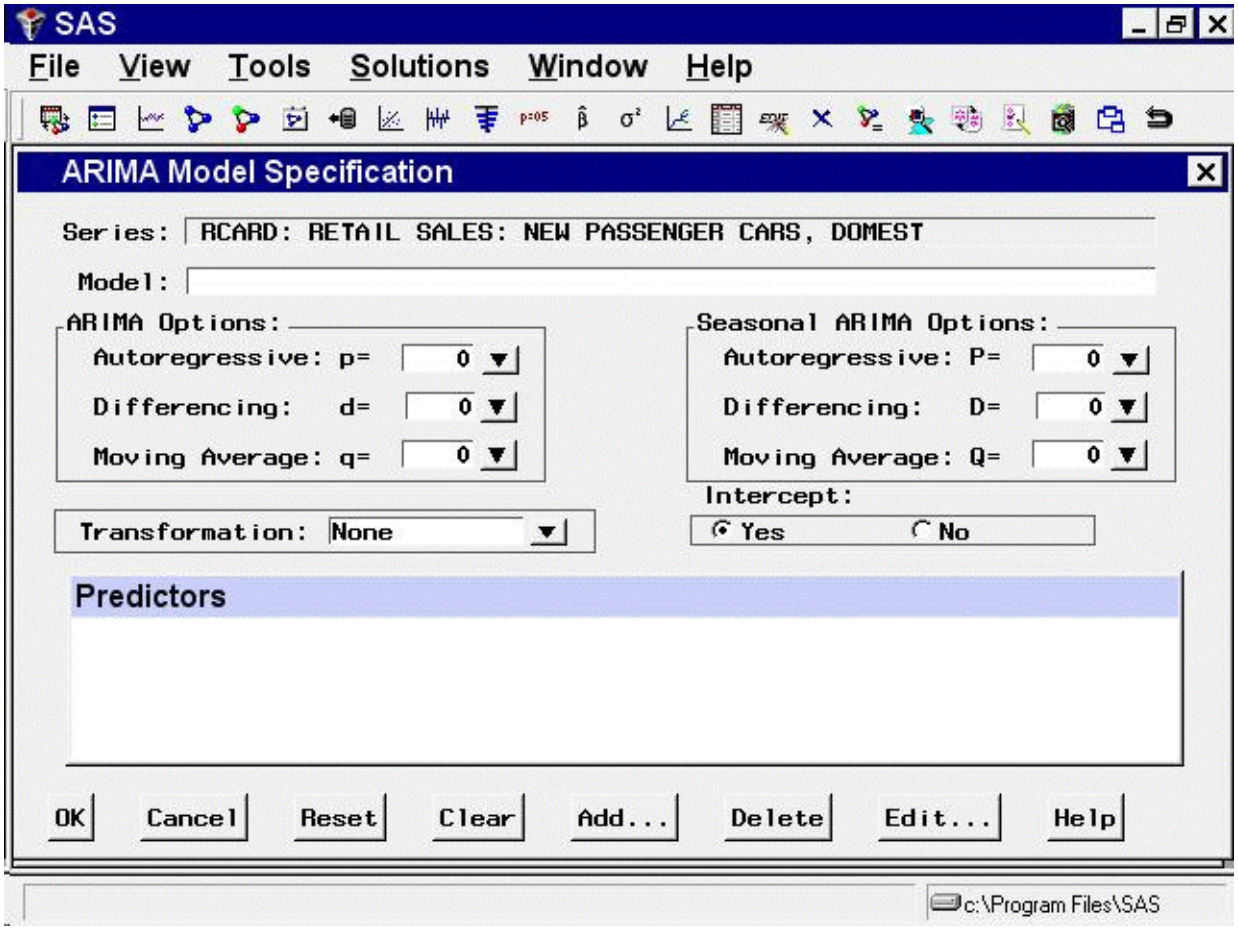

This ARIMA Model Specification window is structured according to the Box and Jenkins approach to time series modeling. You can specify the same time series models with the Custom Model Specification window and the ARIMA Model Specification window, but the windows are structured differently, and you may find one more convenient than the other.

At the top of the ARIMA Model Specification window is the name and label of the series and the label of the model you are specifying. The model label is filled in with an automatically generated label as you specify options. You can type over the automatic label with your own label for the model. To restore the automatic label, enter a blank label.

Using the ARIMA Model Specification window, you can specify autoregressive (p), differencing (d), and moving average (q) orders for both simple and seasonal factors. You can specify transformations with the `Transformation` list. You can also specify whether an intercept is included in the ARIMA model.

In addition to specifying seasonal and nonseasonal ARIMA processes, you can also specify predictor variables and other terms as inputs to the model. ARIMA models with inputs are sometimes called ARIMAX models or Box-Tiao models. Another term for this kind of model is *dynamic regression*.

In the lower part of the ARIMA Model Specification window is the list of *predictors* to the model (initially empty). You can specify predictors by using the Add button. This opens a menu of different kinds of independent effects, as shown in Figure 58.14.

Figure 58.14 Add Predictors Menu

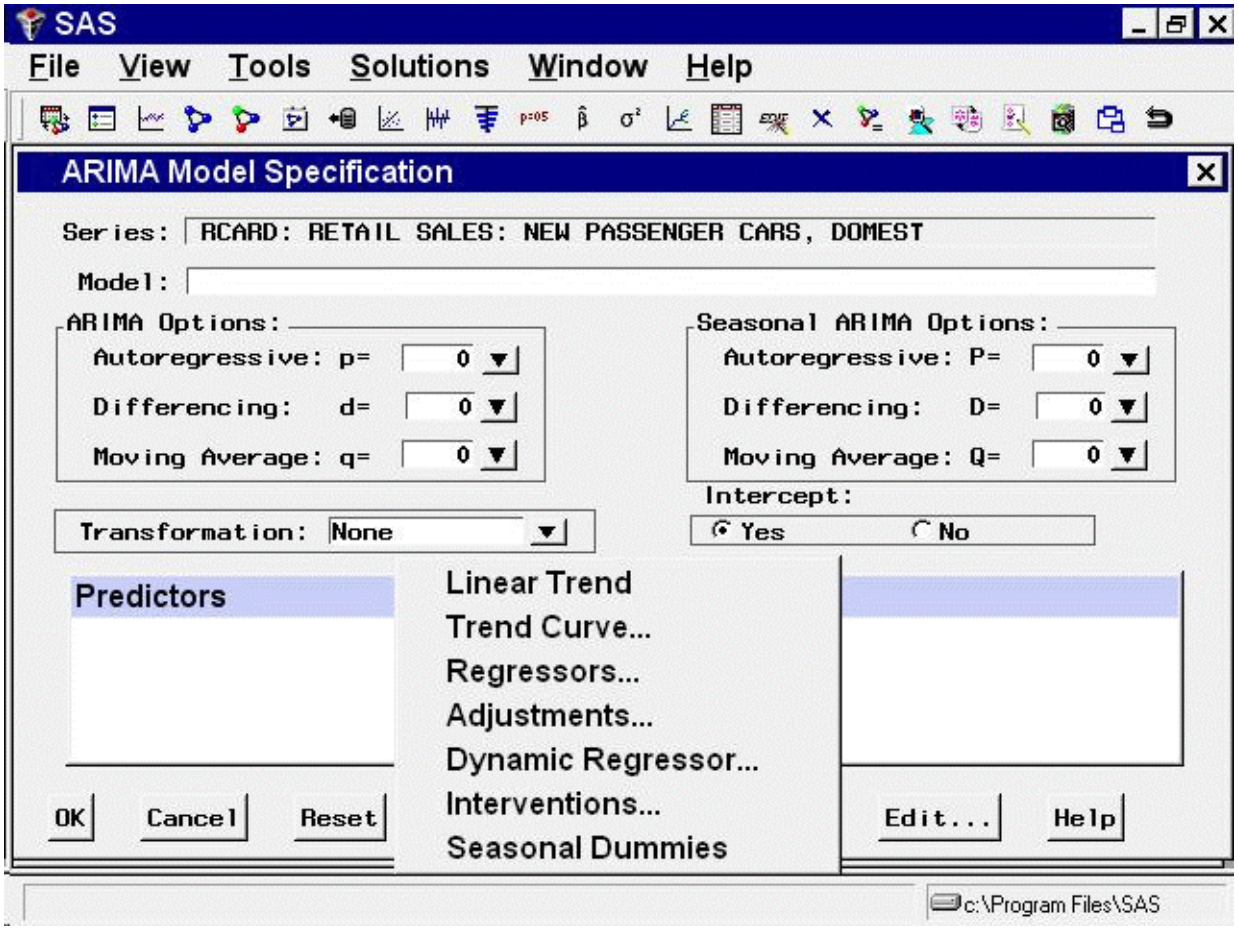

The kinds of predictor effects allowed include time trends, regressors, adjustments, dynamic regression (transfer functions), intervention effects, and seasonal dummy variables. How to use different kinds of predictors is explained in Chapter 60, "Using Predictor Variables."

As an example, in the `ARIMA Options` box, set the order of differencing `d` to `1` and the moving average order `q` to `2`. You can either type in these values or click the arrows and select the values from pop-up lists.

These selections specify an ARIMA(0,1,2) or IMA(1,2) model. (For more information about the notation used for ARIMA models, see Chapter 7, "The ARIMA Procedure.") Notice that the model label at the top is now `IMA(1,2) NOINT`, meaning that the data are differenced once and a second-order moving-average term is included with no intercept.

In the `Seasonal ARIMA Options` box, set the seasonal moving-average order `Q` to `1`. This adds a first-order moving-average term at the seasonal (12 month) lag. Finally, select "Log" in the Transformation combo box.

The model label is now `Log ARIMA(0,1,2)(0,0,1)s NOINT`, and the window appears as shown in Figure 58.15.

Figure 58.15 Log ARIMA(0,1,2)(0,0,1)s Specified

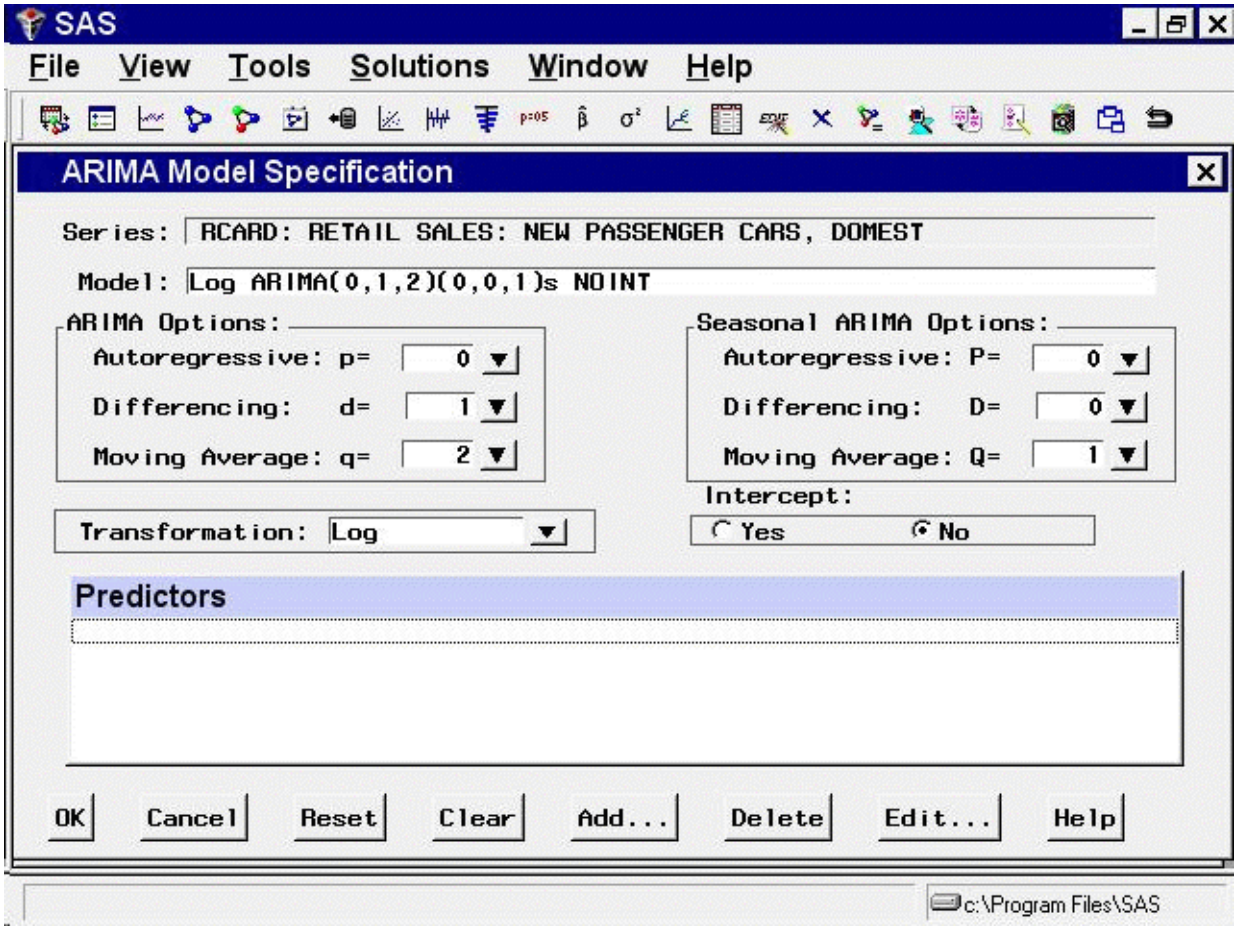

Click the OK button to fit the model. The model is fit and added to the Develop Models table.

Factored ARIMA Model Specification Window

To fit a factored ARIMA model, select the Factored ARIMA model item from the pop-up menu, toolbar, or Edit menu. This brings up the Factored ARIMA Model Specification window, shown in Figure 58.16.

Figure 58.16 Factored ARIMA Model Specification Window

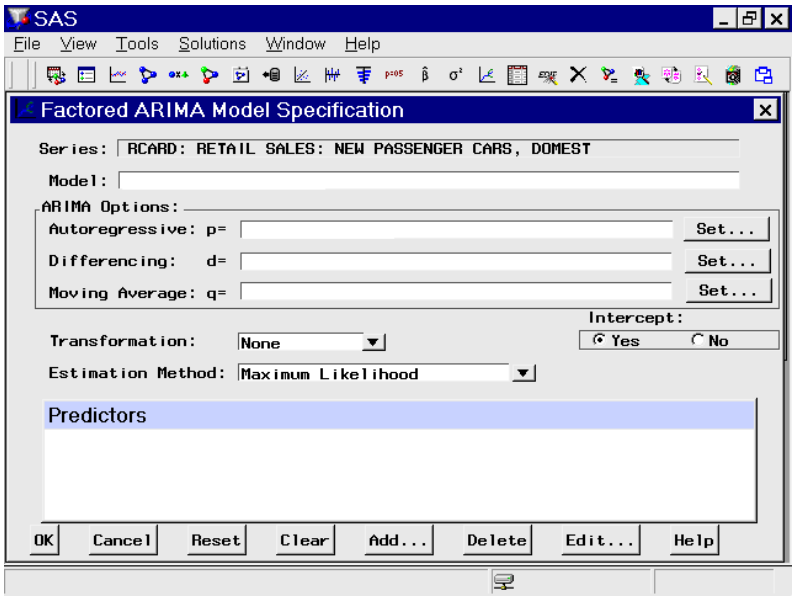

The Factored ARIMA Model Specification window is similar to the ARIMA Model Specification window and has the same features, but it uses a more general specification of the autoregressive (p), differencing (d), and moving-average (q) terms. To specify these terms, click the corresponding Set button, as shown in Figure 58.16. For example, to specify autoregressive terms, click the first Set button. This opens the AR Polynomial Specification Window, shown in Figure 58.17.

Figure 58.17 AR Polynomial Specification Window

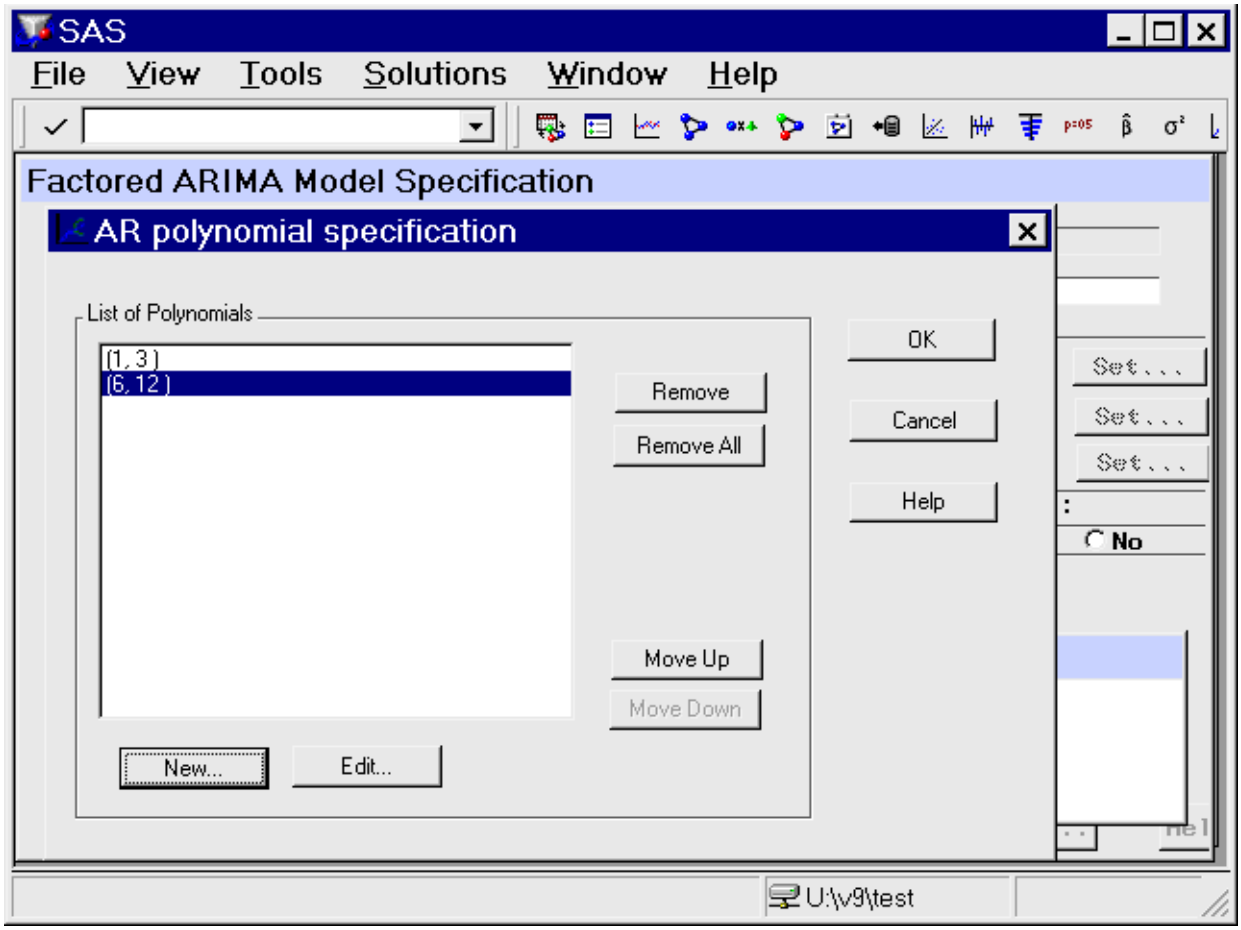

To add AR polynomial terms, click the New button. This opens the Polynomial Specification Window, shown in Figure 58.18. Specify the first lag you want to include by using the `Lag` spin box, then click the Add button. Repeat this process, adding each lag you want to include in the current list. All lags must be specified. For example, if you add only lag 3, the model contains only lag 3, not 1 through 3.

As an example, add lags 1 and 3, then click the OK button. The AR Polynomial Specification Window now shows (1,3) in the list of polynomials. Now select "New" again. Add lags 6 and 12 and click "OK". Now the AR Polynomial Specification Window shows (1,3) and (6,12) as shown in Figure 58.17. Select "OK" to close this window. The Factored ARIMA Model Specification Window now shows the factored model `p=(1,3)(6,12)`. Use the same technique to specify the q terms, or moving-average part of the model. There is no limit to the number of lags or the number of factors you can include in the model.

Figure 58.18 Polynomial Specification Window

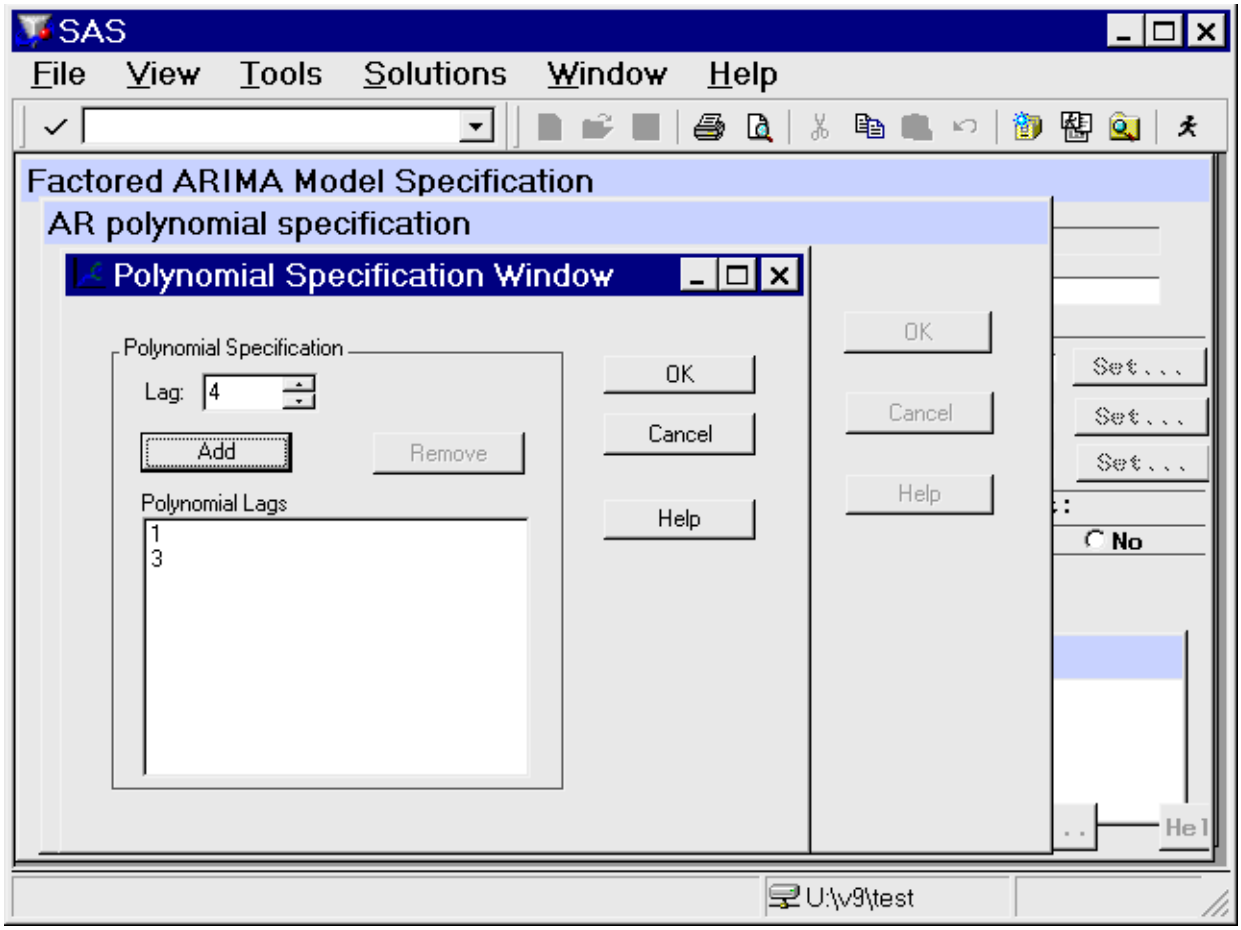

To specify differencing lags, click the middle Set button to open the Differencing Specification window. Specify lags using the spin box and add them to the list with the Add button. When you click "OK" to close the window, the differencing lags appear after `d=` in the Factored ARIMA Specification Window, within a single pair of parentheses.

You can use the Factored ARIMA Model Specification Window to specify any model that you can specify with the ARIMA Model and Custom Model windows, but the notation is more similar to that of the ARIMA procedure (see Chapter 7, "The ARIMA Procedure"). Consider as an example the classic Airline model fit to the International Airline Travel series, `SASHELP.AIR`. This is a factored model with one moving-average term at lag one and one moving-average term at the seasonal lag, with first-order differencing at the simple and seasonal lags. Using the ARIMA Model Specification Window, you specify the value 1 for the q and d terms and also for the Q and D terms, which represent the seasonal lags. For monthly data, the seasonal lags represent lag 12, since a yearly seasonal cycle is assumed.

By contrast, the Factored ARIMA Model Specification Window makes no assumptions about seasonal cycles. The Airline model is written as `IMA d=(1,12) q=(1)(12) NOINT`. To specify the differencing terms, add the values 1 and 12 in the Differencing Specification Window and click OK. Then select "New" in the MA Polynomial Specification Window, add the value 1, and select OK. To add the factored term, select "New" again, add the value 12, and click OK. Remember to select "No" in the Intercept radio box, since it is not selected by default. Click OK to close the Factored ARIMA Model Specification Window and fit the model.

You can show that the results are the same as they are when you specify the model by using the ARIMA Model Specification Window and when you select Airline Model from the default model list. If you are familiar with the ARIMA Procedure (Chapter 7, "The ARIMA Procedure"), you might want to turn on the `Show Source Statements` option before fitting the model, then examine the procedure source statements in the log window after fitting the model.

The strength of the Factored ARIMA Specification approach lies in its ability to construct unusual ARIMA models, such as the following:

Subset models
: These are models of order n, where fewer than n lags are specified. For example, an AR order 3 model might include lags 1 and 3 but not lag 2.

Unusual seasonal cycles
: For example, a monthly series might cycle two or four times per year instead of just once.

Multiple cycles
: For example, a daily sales series might peak on a certain day each week and also once a year at the Christmas season. Given sufficient data, you can fit a three-factor model, such as `IMA d=(1) q=(1)(7)(365)`.

Models with high order lags take longer to fit and often fail to converge. To save time, select the Conditional Least Squares or Unconditional Least Squares estimation method (see Figure 58.16). Once you have narrowed down the list of candidate models, change to the Maximum Likelihood estimation method.

Custom Model Specification Window

To fit a custom time series model not already provided in the Models to Fit window, select the Custom Model item from the pop-up menu, toolbar, or Edit menu. This opens the Custom Model Specification window, as shown in Figure 58.19.

Figure 58.19 Custom Model Specification Window

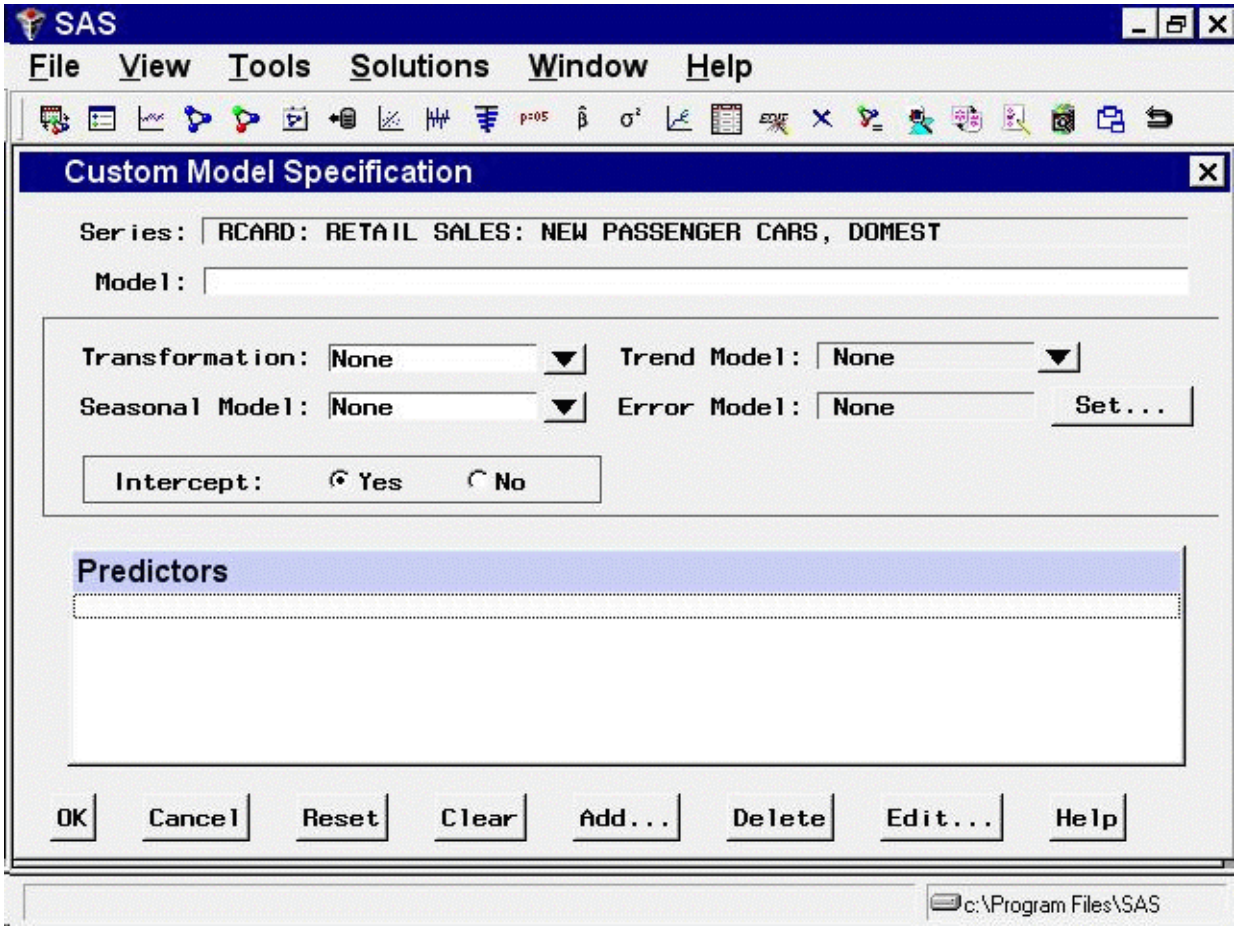

You can specify the same time series models with the Custom Model Specification window and the ARIMA Model Specification window, but the windows are structured differently, and you might find one more convenient than the other.

At the top of the Custom Model Specification window is the name and label of the series and the label of the model you are specifying. The model label is filled in with an automatically generated label as you specify options. You can type over the automatic label with your own label for the model. To restore the automatic label, enter a blank label.

The middle part of the Custom Model Specification window consists of four fields: `Transformation`, `Trend Model`, `Seasonal Model`, and `Error Model`. These fields allow you to specify the model in four parts. Each part specifies how a different aspect of the pattern of the time series is modeled and predicted.

The `Predictors` list at the bottom of the Custom Model Specification window allows you to include different kinds of predictor variables in the forecasting model. The Predictors feature for the Custom Model Specification window is like the Predictors feature for the ARIMA Model Specification window, except that time trend predictors are provided through the Trend Model field and seasonal dummy variable predictors are provided through the Seasonal Model field.

To illustrate how to use the Custom Model Specification window, the following example specifies the same model you fit by using the ARIMA Model Specification window.

First, specify the data transformation to use. Select "Log" using the Transformation combo box.

Second, specify how to model the trend in the series. Select `First Difference` in the `Trend Model` combo box, as shown in Figure 58.20.

Figure 58.20 Trend Model Options

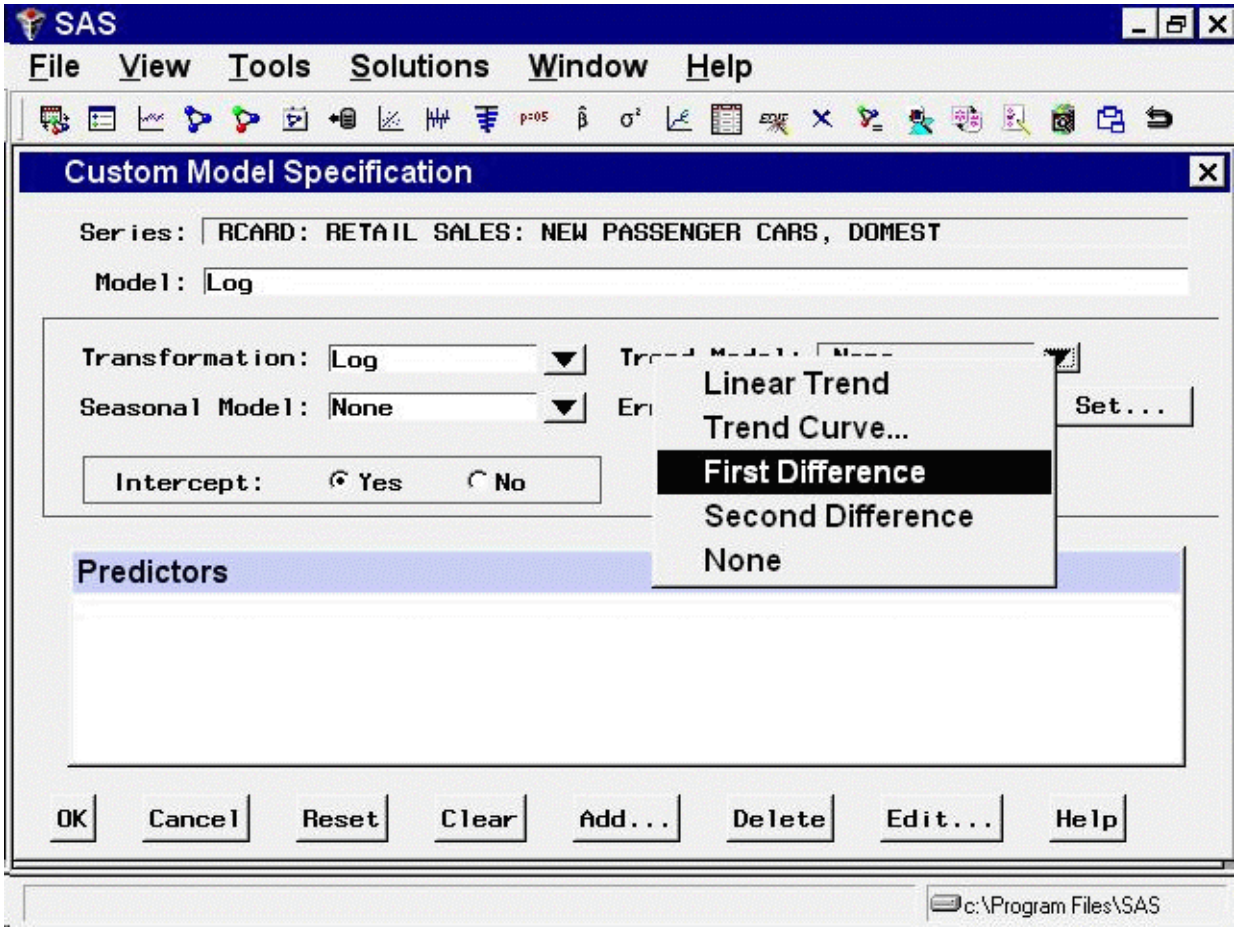

Next, specify how to model the seasonal pattern in the series. Select "Seasonal ARIMA" in the Seasonal Model combo box, as shown in Figure 58.21.

Figure 58.21 Seasonal Model Options

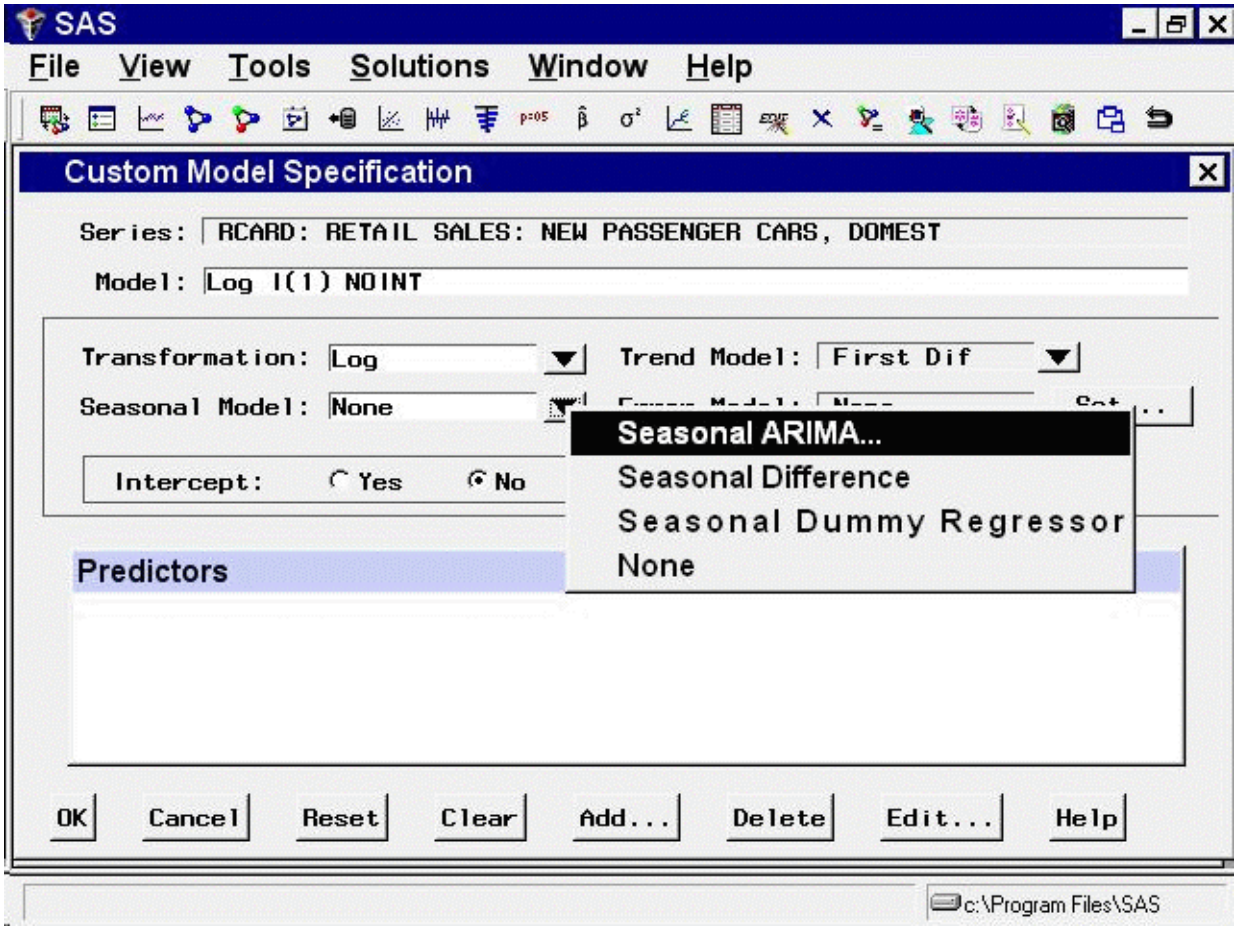

This opens the Seasonal ARIMA Model Options window, as shown in Figure 58.22.

Figure 58.22 Seasonal ARIMA Model Options

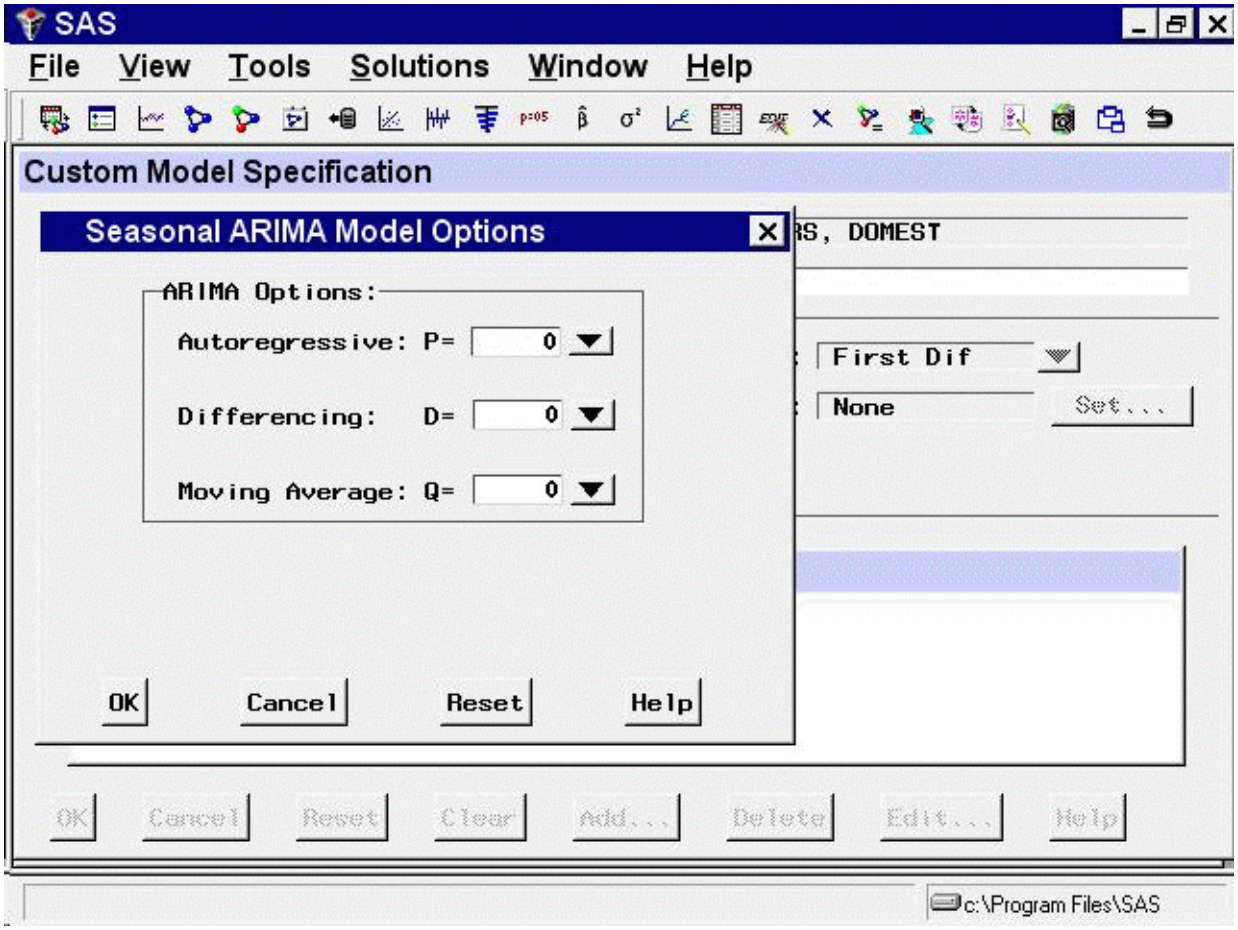

Specify a first-order seasonal moving-average term by typing `1` or by selecting "1" from the Moving Average: Q= combo box pop-up menu, and then click the OK button.

Finally, specify how to model the autocorrelation pattern in the model prediction errors. Click the Set button to the right of the Error Model field. This opens the Error Model Options window, as shown in Figure 58.23. This window allows you to specify an ARMA error process. Set the Moving Average order `q` to `2`, and then click the OK button.

Figure 58.23 Error Model Options

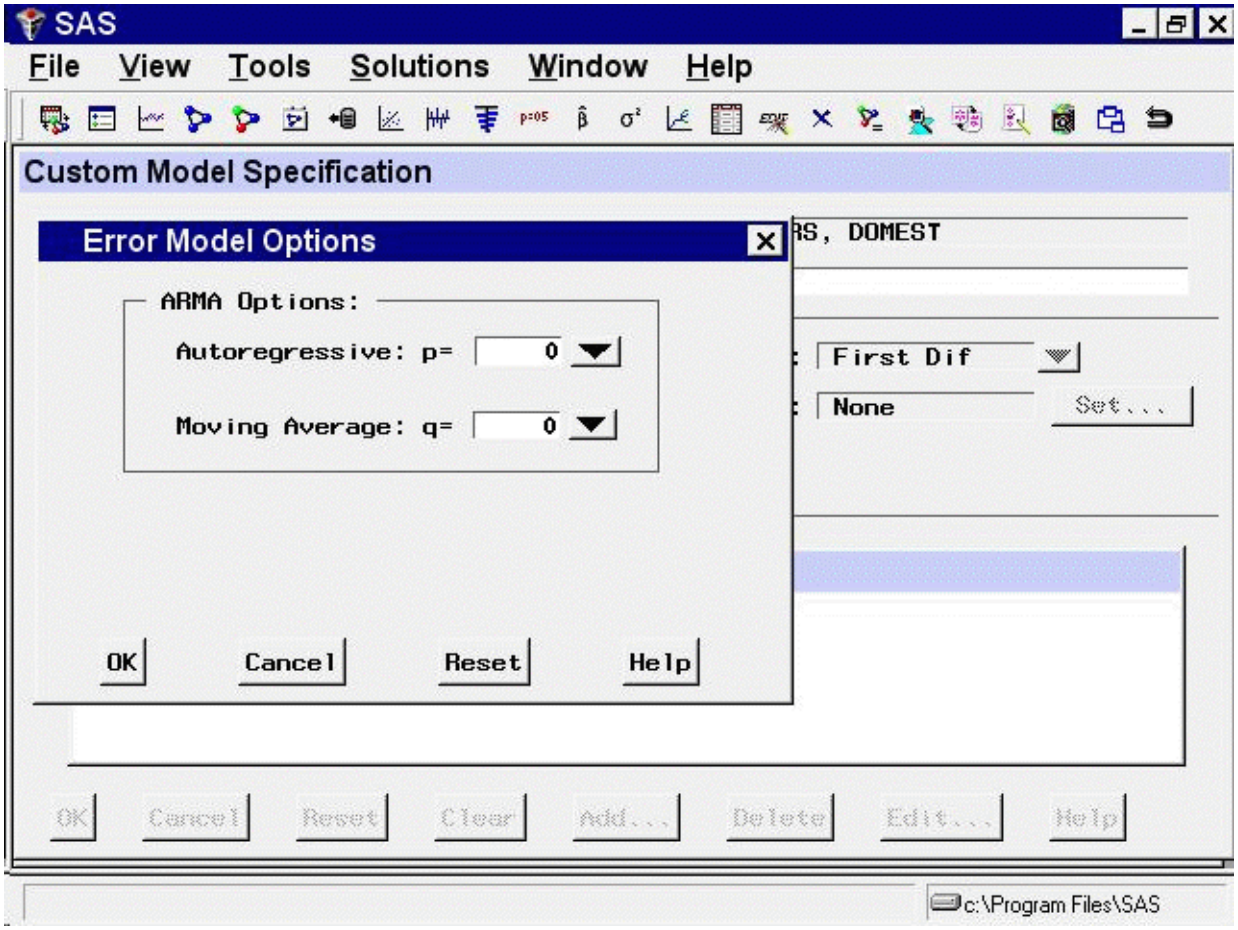

The Custom Model Specification window should now appear as shown in Figure 58.24. The model label at the top of the Custom Model Specification window should now read `Log ARIMA(0,1,2)(0,0,1)s NOINT`, just as it did when you used the ARIMA Model Specification window.

Figure 58.24 Log ARIMA(0,1,2)(0,0,1)s Specified

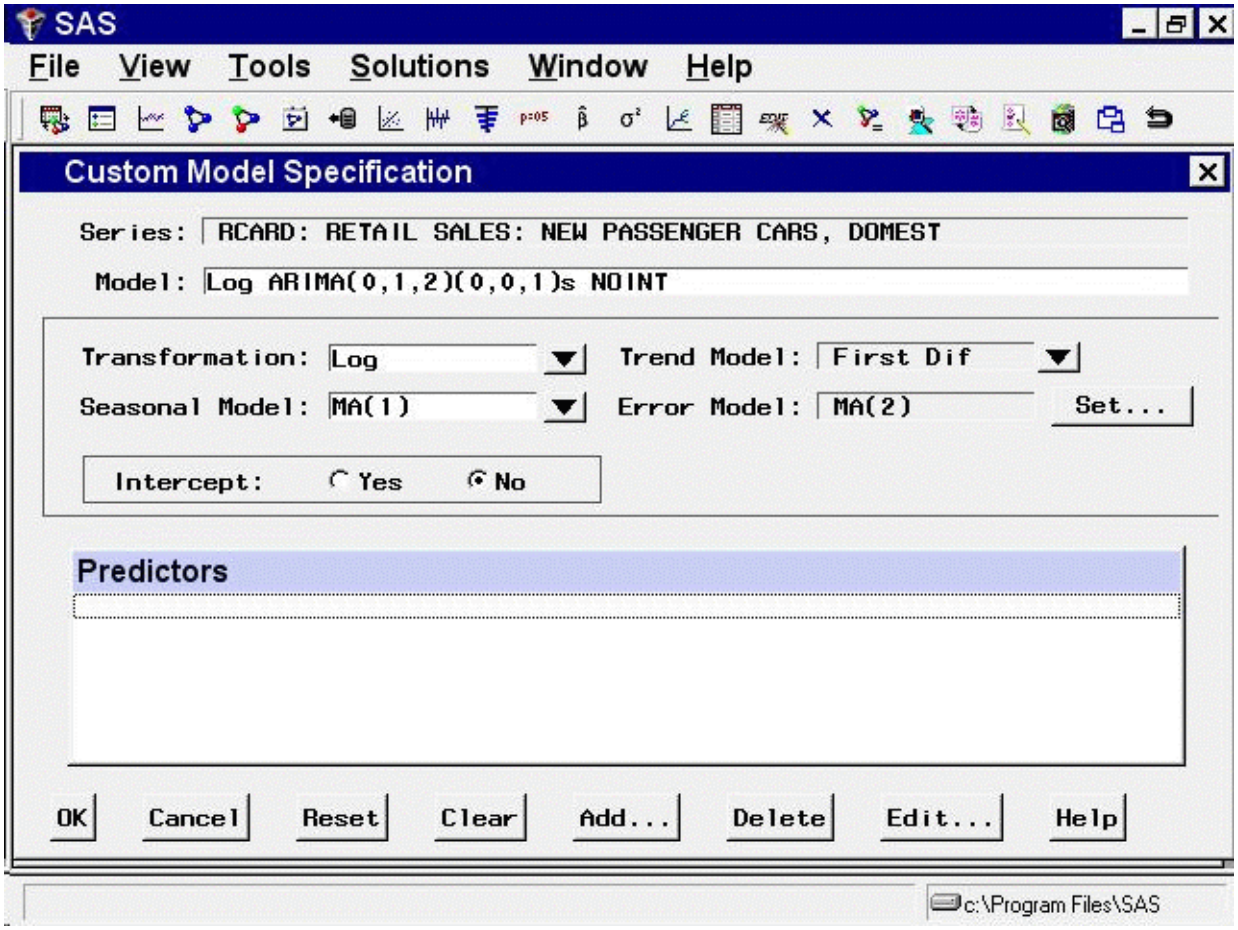

Now that you have seen how the Custom Model Specification window works, click “Cancel” to exit the window without fitting the model. This should return you to the Develop Models window.

Editing the Model Selection List

Now that you know how to specify new models that are not included in the system default model selection list, you can edit the model selection list to add models that you expect to use in the future or to delete models that you do not expect to use. When you save the forecasting project to a SAS catalog, the edited model selection list is saved with the project file, and the list is restored when you load the project.

There are two reasons why you would add a model to the model selection list. First, by adding the model to the list, you can fit the model to different time series by selecting it through the `Fit Models from List` action. You do not need to specify the model again every time you use it.

Second, once the model is added to the model selection list, it is available to the automatic model selection process. The model is then considered automatically whenever you use the automatic model selection feature for any series.

To edit the model selection list, select "Model Selection List" from the Options menu as shown in Figure 58.25, or select the Edit Model List toolbar icon.

Figure 58.25 Model Selection List Option

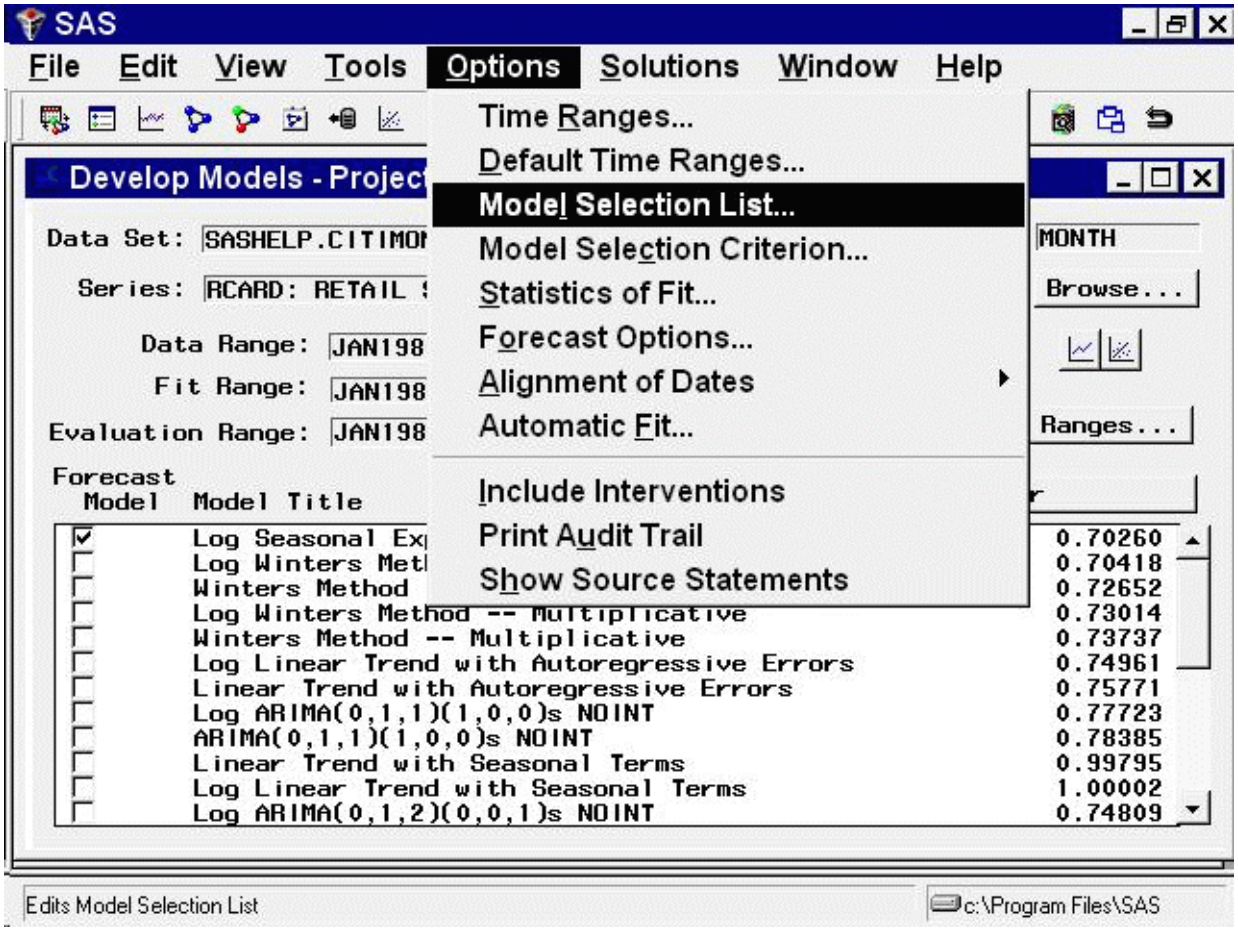

This selection brings up the Model Selection List editor window, as shown in Figure 58.26. This window consists of the model selection list and an "Auto Fit" column, which controls for each model whether the model is included in the list of models used by the automatic model selection process.

Figure 58.26 Model Selection List Window

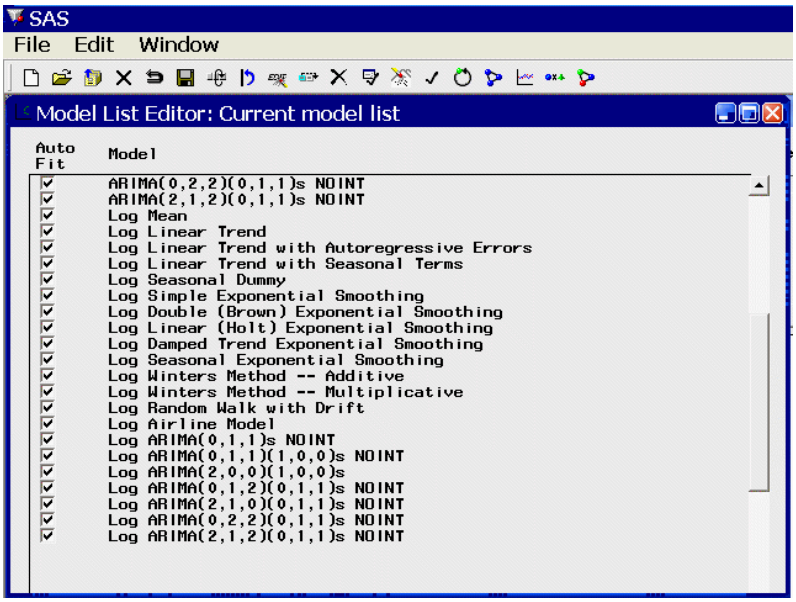

To add a model to the list, select "Add Model" from the Edit menu and then select "Smoothing Model," "ARIMA Model," "Factored ARIMA Model," or "Custom Model" from the submenu. Alternatively, click the corresponding icon on the toolbar.

As an example, select "Smoothing Model." This brings up the Smoothing Model Specification window. Note that the series name is "-Null-." This means that you are not specifying a model to be fit to a particular series, but are specifying a model to be added to the selection list for later reference.

Specify a smoothing model. For example, select "Simple Smoothing" and then select the Square Root transformation. The window appears as shown in Figure 58.27.

Figure 58.27 Adding a Model Specification

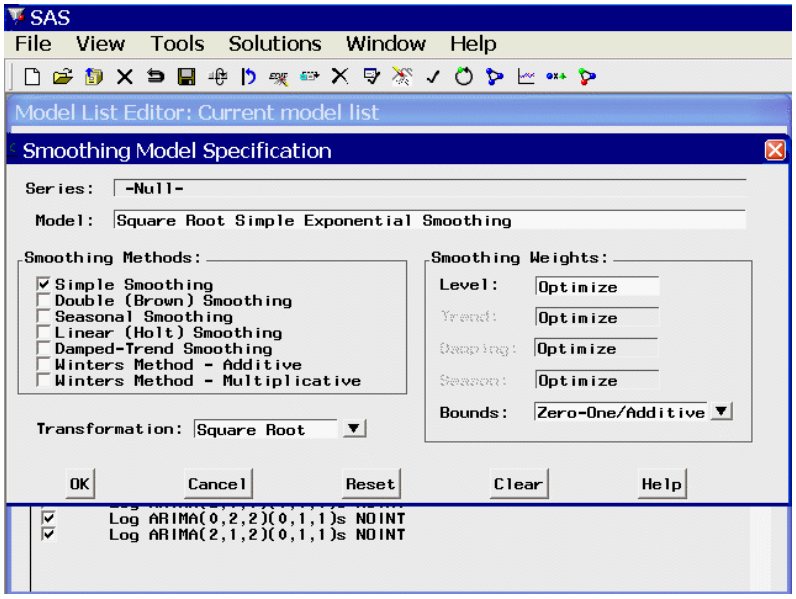

Click the OK button to add the model to the end of the model selection list and return you to the Model Selection List window, as shown in Figure 58.28. You can now select the Fit Models from List model-fitting option to use the edited selection list.

Figure 58.28 Model Added to Selection List

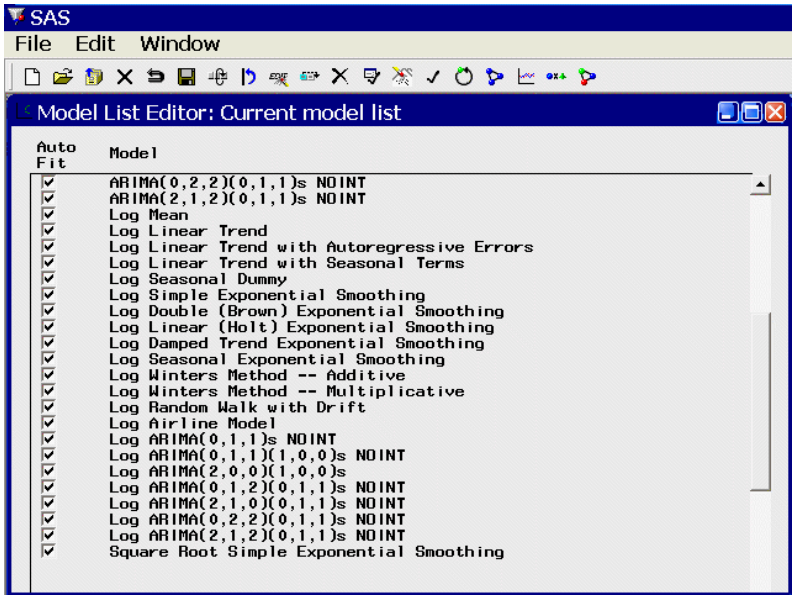

If you want to delete one or more models from the list, select the model labels to highlight them in the list. Click a second time to clear a selected model. Then select "Delete" from the Edit pull-down menu, or the corresponding toolbar icon. As an example, delete the `Square Root Simple Exponential Smoothing` model that you just added.

The Model Selection List editor window gives you a lot of flexibility for managing multiple model lists, as explained in the section "Model Selection List Editor Window" on page 3986. For example, you can create your own model lists from scratch or modify or combine previously saved model lists and those provided with the software, and you can save them and designate one as the default for future projects.

Now select "Close" from the File menu (or the Close icon) to close the Model Selection List editor window.

Forecast Combination Model Specification Window

Once you have fit several forecasting models to a series, you face the question of which model to use to produce the final forecasts. One possible answer is to combine or average the forecasts from several models. Combining the predictions from several different forecasting methods is a popular approach to forecasting.

The way that you produce forecast combinations with the Time Series Forecasting System is to use the Forecast Combination Model Specification window to specify a new forecasting model that performs the averaging of forecasts from the models you want to combine. This new model is added to the list of fitted models just like other models. You can then use the Model Viewer window features and Model Fit Comparison window features to examine the fit of the combined model.

To specify a forecast combination model, select "Combine Forecasts" from the pop-up menu or toolbar, or select "Edit" and "Fit Model" from the menu bar. This brings up the Forecast Combination Model Specification window, as shown in Figure 58.29.

Figure 58.29 Forecast Combination Window

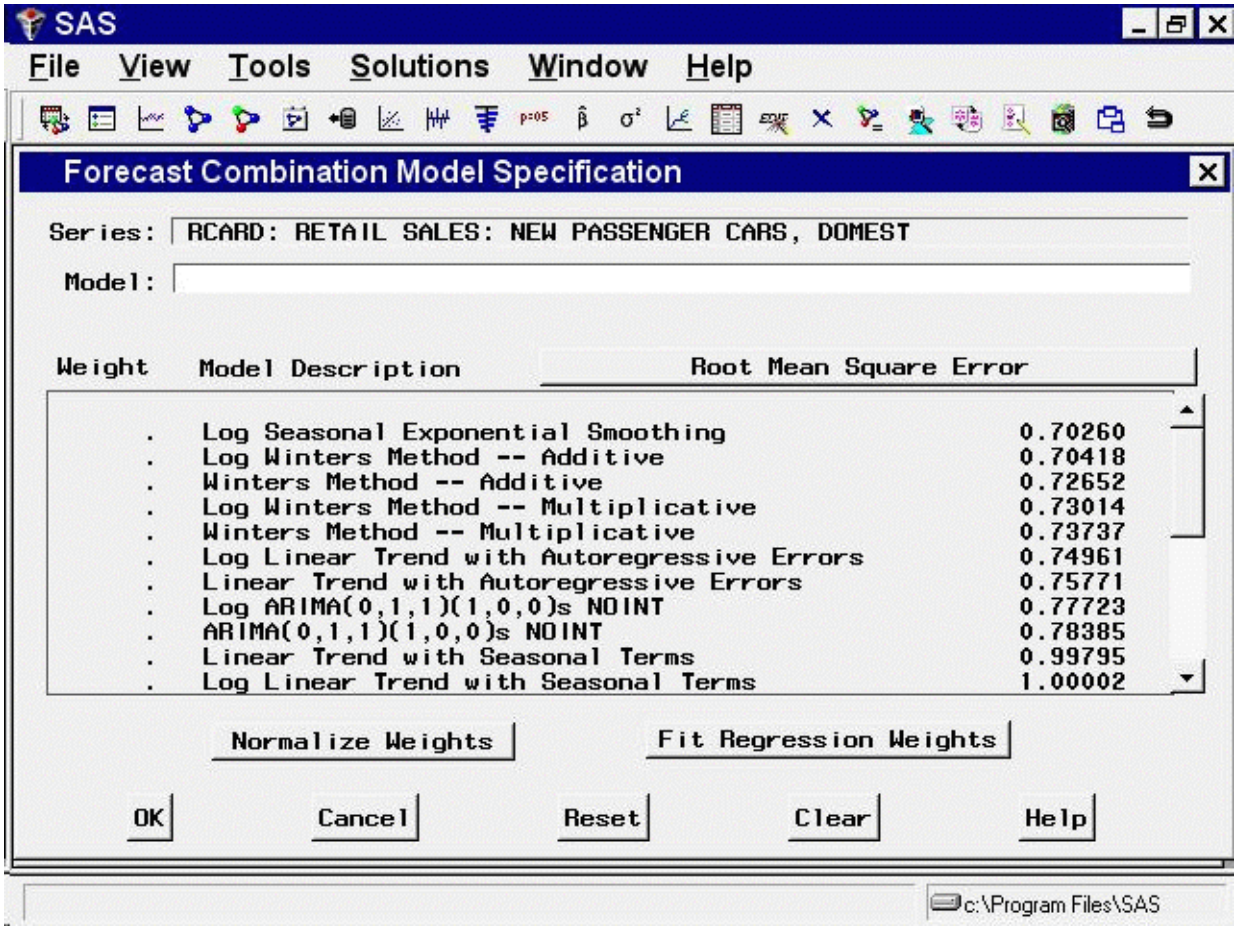

At the top of the Forecast Combination window is the name and label of the series and the label of the model you are specifying. The model label is filled in with an automatically generated label as you specify options. You can type over the automatic label with your own label for the model. To restore the automatic label, enter a blank label.

The middle part of the Forecast Combination window consists of the list of models that you have fit to the series. This table shows the label and goodness-of-fit measure for each model and the combining weight assigned to the model.

The `Weight` column controls how much weight is given to each model in the combined forecasts. A missing weight means that the model is not used. Initially, all the models have missing weight values.

You can enter the weight values you want to use in the `Weight` column. Alternatively, you can select models from the Model Description column, and weight values for the models you select are set automatically. To remove a model from the combination, select it again. This resets its weight value to missing.

At the bottom of the Forecast Combination window are two buttons: `Normalize Weights` and `Fit Regression Weights`. The Normalize Weights button adjusts the nonmissing weight values so that they sum to one. The Fit Regression Weights button uses linear regression to compute the weight values that produce the combination of model predictions with the best fit to the series.

If no models are selected, the Fit Regression Weights button fits weights for all the models in the list. You can compute regression weights for only some of the models by first selecting the models you want to combine and then selecting Fit Regression Weights. In this case, only the nonmissing Weight values are replaced with regression weights.

As an example of how to combine forecasting models, select all the models in the list. After you have finished selecting the models, all the models in the list should now have equal weight values, which implies a simple average of the forecasts.

Now click the Fit Regression Weights button. The system performs a linear regression of the series on the predictions from the models with nonmissing weight values and replaces the weight values with the estimated regression coefficients. These are the combining weights that produce the smallest mean square prediction error within the sample.

The Forecast Combination window should now appear as shown in Figure 58.30. (Note that some of the regression weight values are negative.)

Figure 58.30 Combining Models

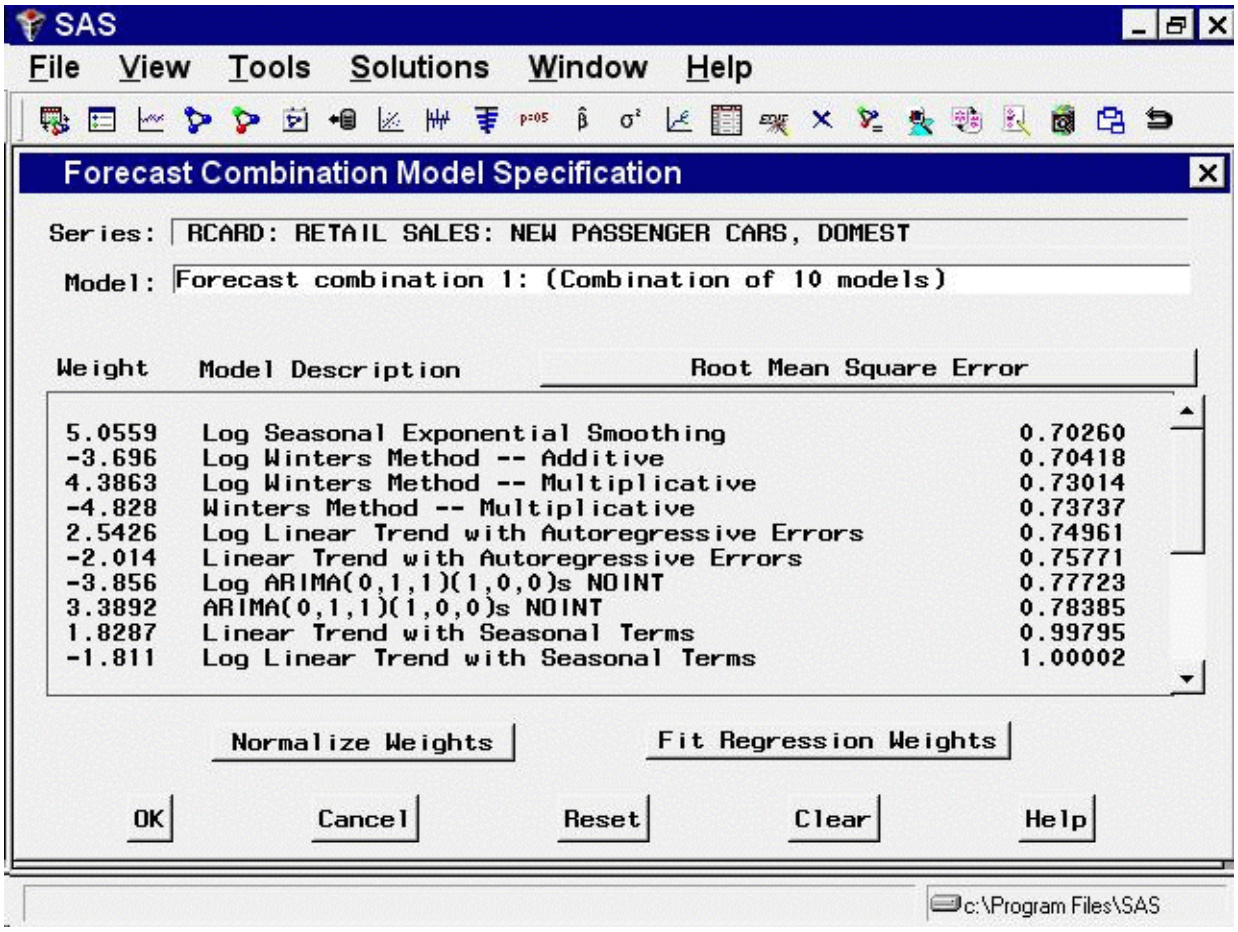

Click the OK button to fit the combined model. Now the Develop Models window shows this model to be the best fitting according to the root mean square error, as shown in Figure 58.31.

Figure 58.31 Develop Models Window Showing All Models Fit

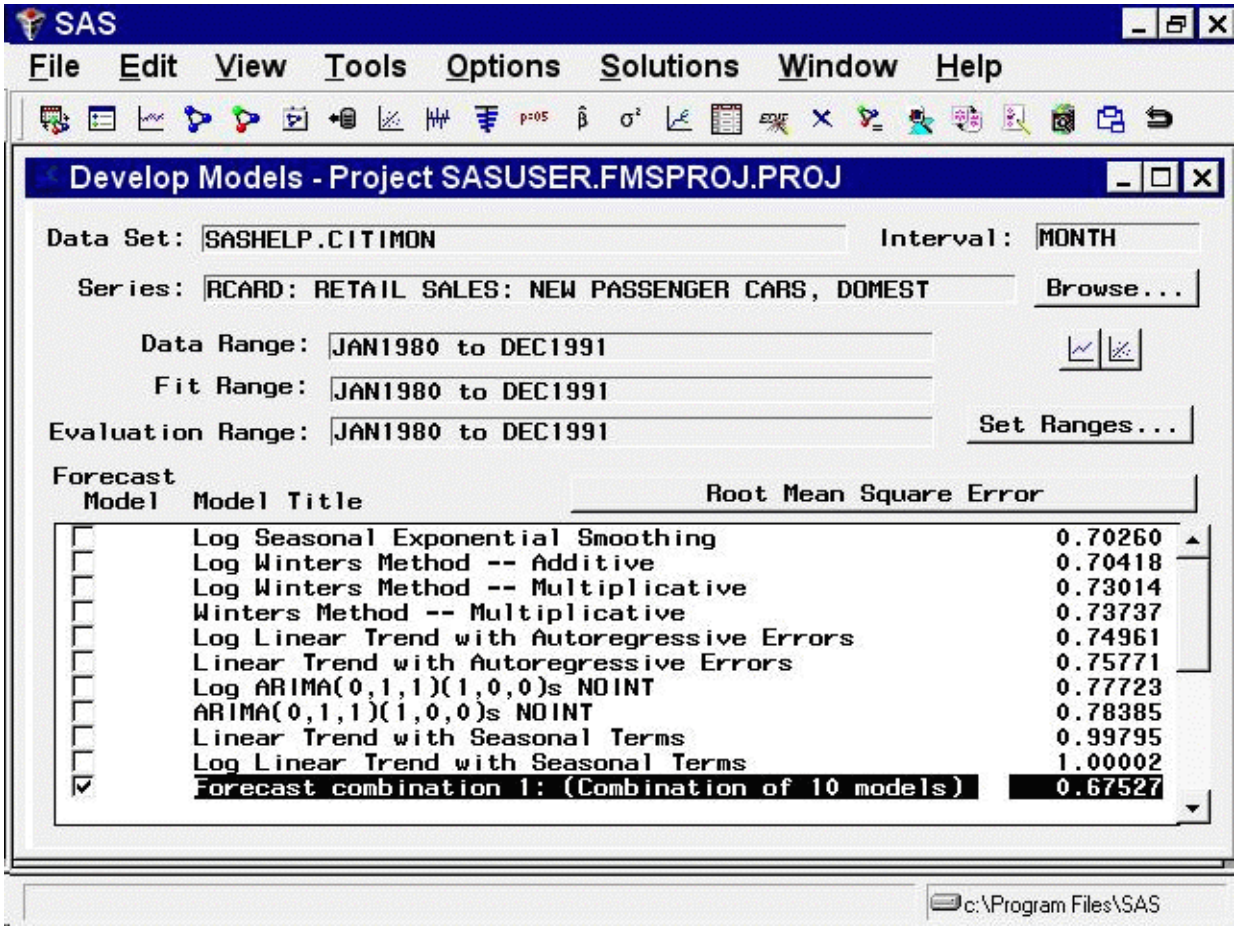

Notice that the combined model has a smaller root mean square error than any one of the models included in the combination. The confidence limits for forecast combinations are produced by taking a weighted average of the mean square prediction errors for the component forecasts, ignoring the covariance between the prediction errors.

Incorporating Forecasts from Other Sources

You might have forecasts from other sources that you want to include in the forecasting process. Examples of other forecasts you might want to use are "best guess" forecasts based on personal judgments, forecasts produced by government agencies or commercial forecasting services, planning scenarios, and reference or "base line" projections. Because such forecasts are produced externally to the Time Series Forecasting System, they are referred to as external forecasts.

You can include external forecasts in combination models to produce compromise forecasts that split the difference between the external forecast and forecasting models that you fit. You can use external forecasts to compare them to the forecasts from models that are fit by the system.

To include external forecasts in the Time Series Forecasting process, you must first supply the external forecast as a variable in the input data set. You then specify a special kind of forecasting "model" whose predictions are identical to the external forecast recorded in the data set.

As an example, suppose you have 12 months of sales data and five months of sales forecasts based on a consensus opinion of the sales staff. The following statements create a SAS data set containing made-up numbers for this situation.

```
data widgets;
   input date monyy5. sales staff;
   format date monyy5.;
   label sales = "Widget Sales"
         staff = "Sales Staff Consensus Forecast";
   datalines;
jun94  142.1    .
jul94  139.6    .
aug94  145.0    .
sep94  150.2    .
oct94  151.1    .
nov94  154.3    .
dec94  158.7    .
jan95  155.9    .
feb95  159.2    .
mar95  160.8    .
apr95  162.0    .
may95  163.3    .
jun95     .  166.
jul95     .  168.
aug95     .  170.
sep95     .  171.
oct95     .  177.
run;
```

Submit the preceding statements in the SAS Program Editor window. From the Time Series Forecasting window, select "Develop Models." In the Series Selection window, select the data set WORK.WIDGETS and the variable SALES. The Develop Models window should now appear as shown in Figure 58.32.

Figure 58.32 Develop Models Window

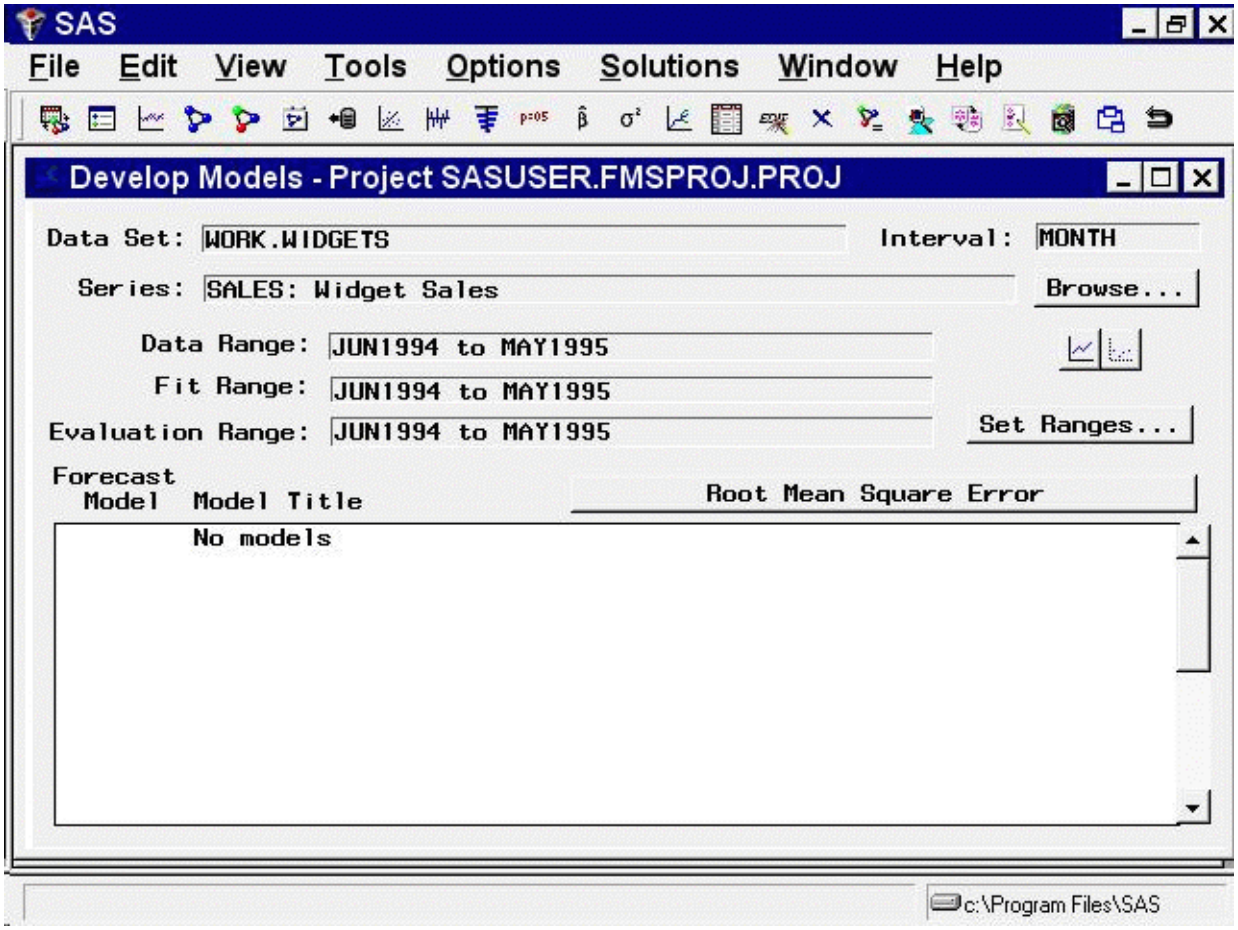

Now select "Edit," "Fit Model," and "External Forecasts" from the menu bar of the Develop Models window, as shown in Figure 58.33, or the `Use External Forecasts` toolbar icon.

Figure 58.33 Adding a Model for an External Forecast Series

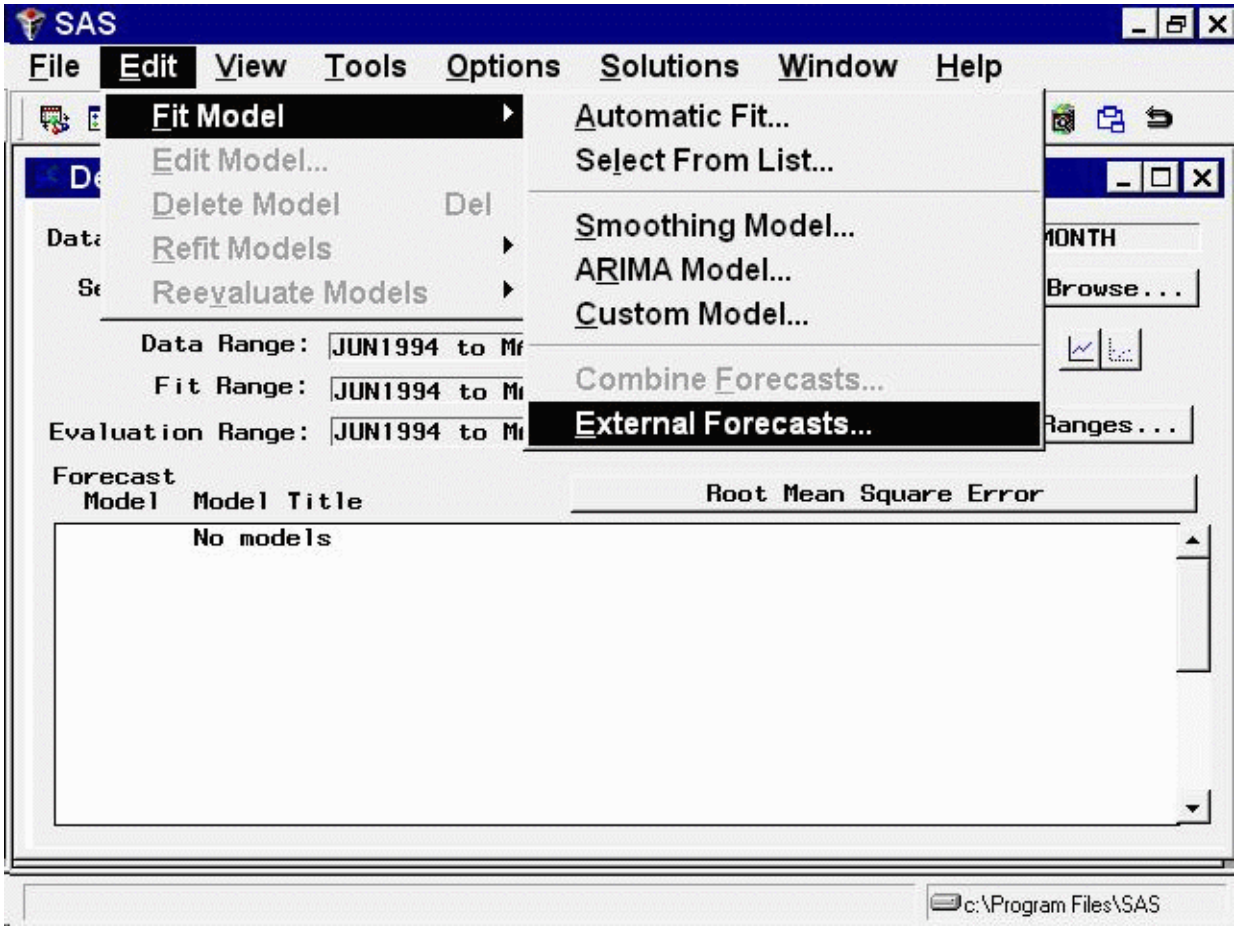

This selection opens the External Forecast Model Specification window. Select the STAFF variable as shown in Figure 58.34.

Figure 58.34 External Forecast Series Selected

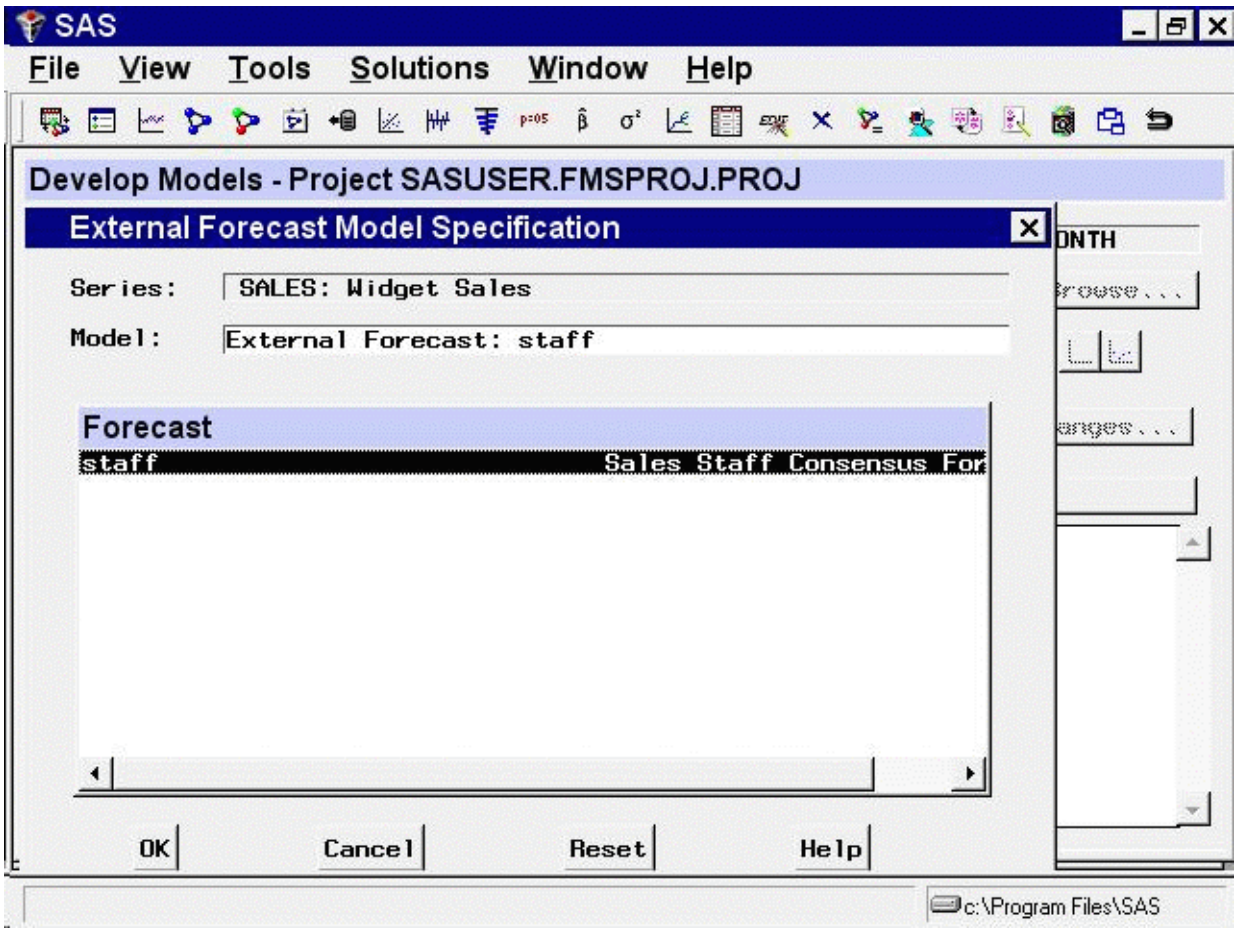

Select the OK button. The external forecast model is now "fit" and added to the Develop Models list, as shown in Figure 58.35.

Figure 58.35 Model for External Forecast

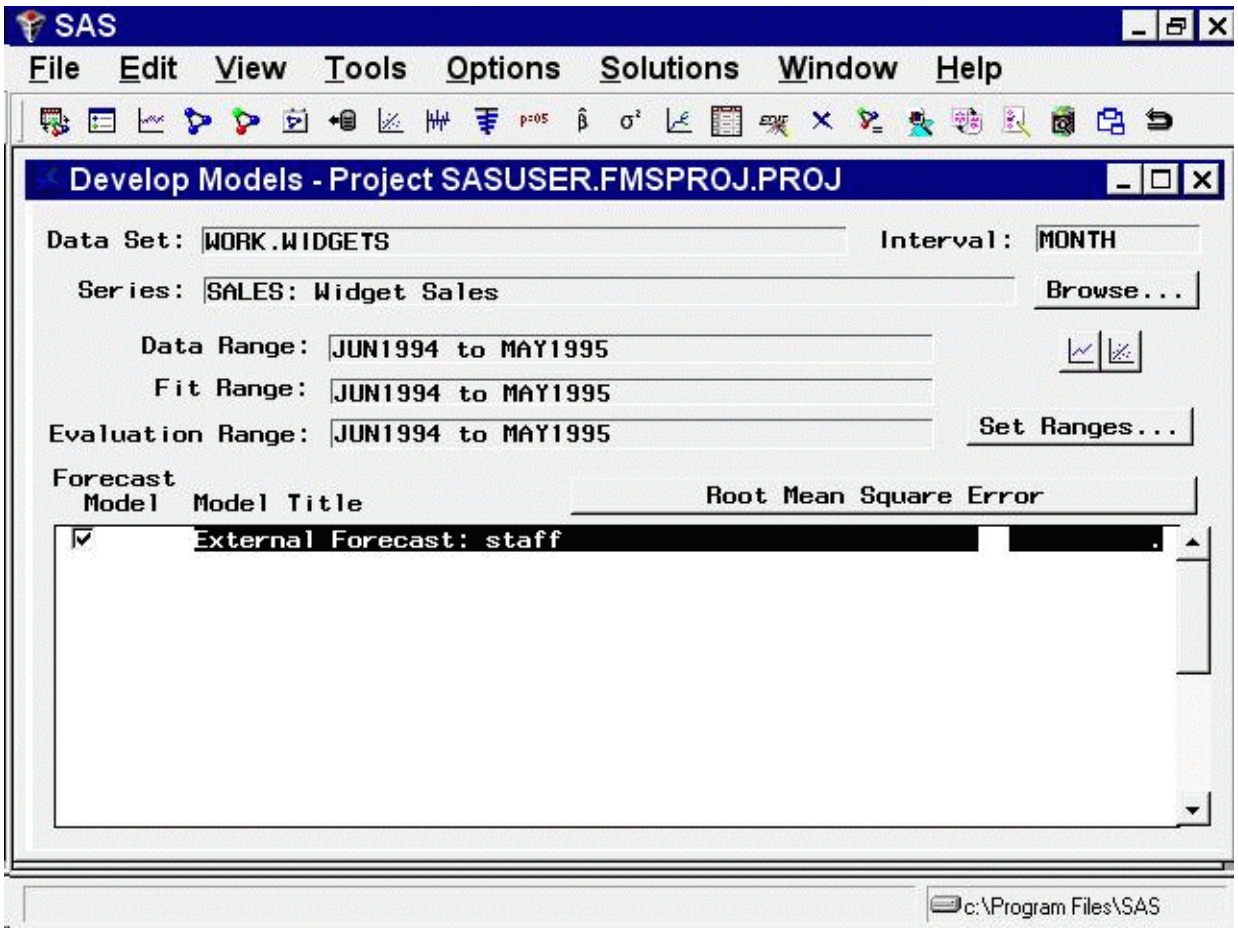

You can now use this model for comparison with the predictions from other forecasting models that you fit, or you can include it in a forecast combination model.

Note that no fitting is actually performed for an external forecast model. The predictions of the external forecast model are simply the values of the external forecast series read from the input data set. The goodness-of-fit statistics for such models will depend on the values that the external forecast series contains for observations within the period of fit. In this case, no STAFF values are given for past periods, and therefore the fit statistics for the model are missing.

Index