

Chapter 27

Using Predictor Variables

Chapter Table of Contents

LINEAR TREND	1329
TIME TREND CURVES	1330
REGRESSORS	1332
ADJUSTMENTS	1334
DYNAMIC REGRESSOR	1335
INTERVENTIONS	1339
The Intervention Specification Window	1339
Specifying a Trend Change Intervention	1341
Specifying a Level Change Intervention	1343
Modeling Complex Intervention Effects	1343
Fitting the Intervention Model	1345
Limitations of Intervention Predictors	1347
SEASONAL DUMMIES	1347
REFERENCES	1349

Chapter 27

Using Predictor Variables

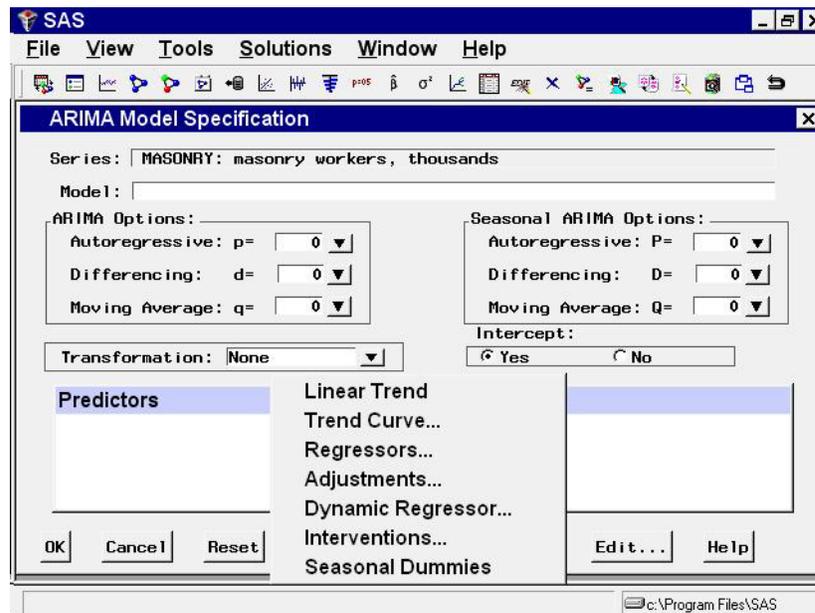
Forecasting models predict the future values of a series using two sources of information: the past values of the series and the values of other time series variables. Other variables used to predict a series are called *predictor variables*.

Predictor variables that are used to predict the dependent series may be variables in the input data set, such as regressors and adjustment variables, or they can be special variables computed by the system as functions of time, such as trend curves, intervention variables, and seasonal dummies.

You can specify seven different types of predictors in forecasting models using the ARIMA Model or Custom Model Specification windows. You cannot specify predictor variables with the Smoothing Model Specification window.

Display 27.1 shows the menu of options for adding predictors to an ARIMA model that is brought up by the Add button. The Add menu for the Custom Model Specification menu is similar.

Display 27.1. Add Predictors Menu



These types of predictors are as follows.

- | | |
|--------------|---|
| Linear Trend | adds a variable that indexes time as a predictor series. A straight line time trend is fit to the series by regression when you specify a linear trend. |
| Trend Curve | provides a menu of various functions of time that you can add to the model to fit nonlinear time trends. The Linear Trend option |

Part 3. General Information

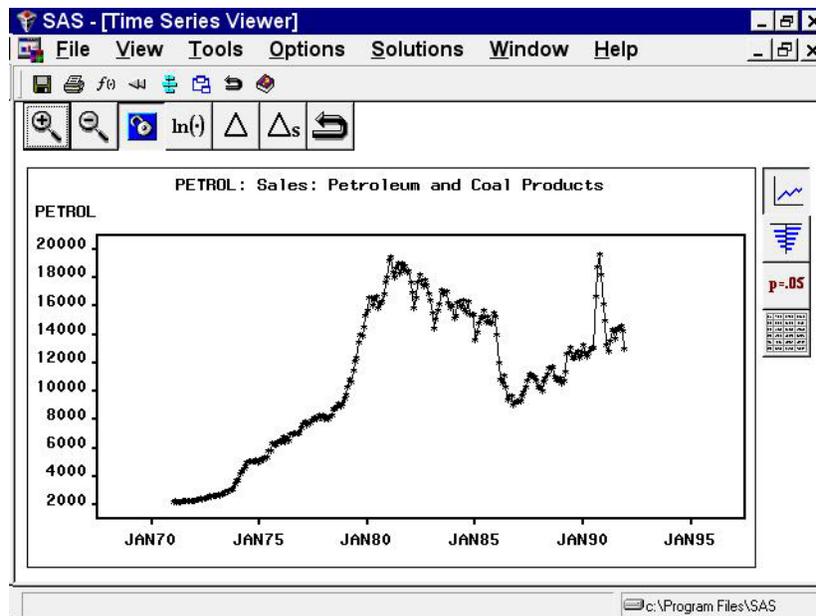
	is a special case of the Trend Curve option for which the trend curve is a straight line.
Regressors	allows you to predict the series by regressing it on other variables in the data set.
Adjustments	allows you to specify other variables in the data set that supply adjustments to the forecast.
Dynamic Regressor	allows you to select a predictor variable from the input data set and specify a complex model for the way that the predictor variable affects the dependent series.
Interventions	allows you to model the effect of special events that "intervene" to change the pattern of the dependent series. Examples of intervention effects are strikes, tax increases, and special sales promotions.
Seasonal Dummies	adds seasonal indicator or "dummy" variables as regressors to model seasonal effects.

You can add any number of predictors to a forecasting model, and you can combine predictor variables with other model options.

The following sections explain these seven kinds of predictors in greater detail and provide examples of their use. The examples illustrate these different kinds of predictors using series in the SASHELP.USECON data set.

Select the `Develop Models` button from the main window. Select the data set `SASHELP.USECON` and select the series `PETROL`. Then select the `View Series Graphically` button from the `Develop Models` window. The plot of the example series `PETROL` appears as shown in Display 27.2.

Display 27.2. Sales of Petroleum and Coal

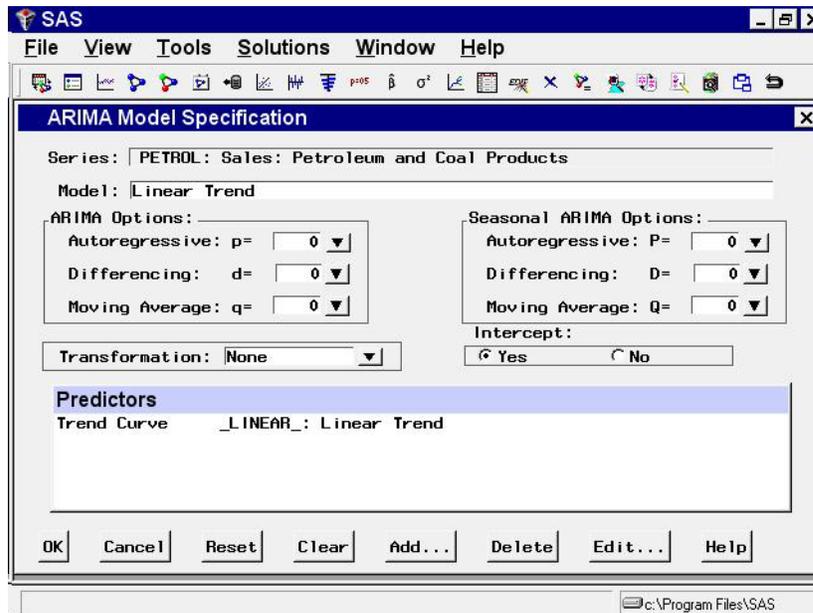


Linear Trend

From the Develop Models window, select **Fit ARIMA Model**. From the ARIMA Model Specification window, select **Add** and then select **Linear Trend** from the menu (shown in Display 27.1).

A linear trend is added to the Predictors list, as shown in Display 27.3.

Display 27.3. Linear Trend Predictor Specified



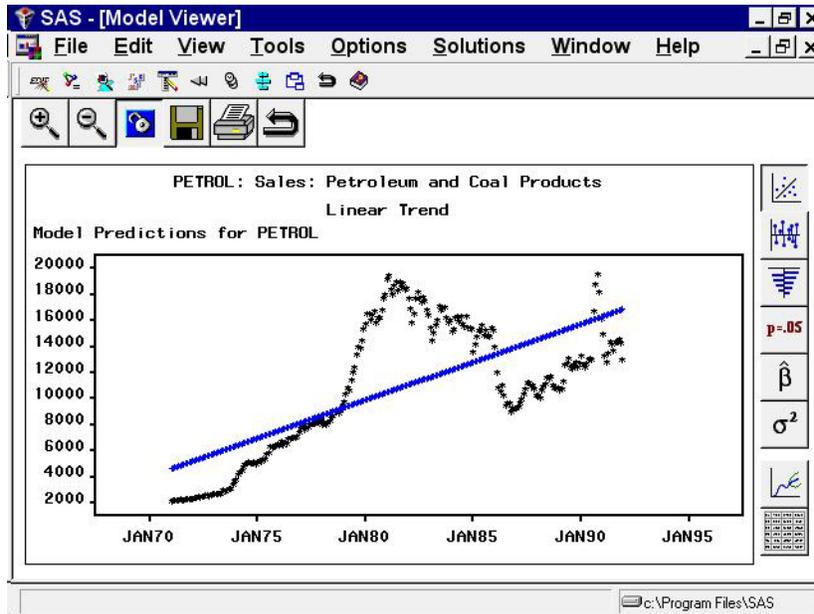
The description for the linear trend item shown in the Predictors list has the following meaning. The first part of the description, Trend Curve, describes the type of predictor. The second part, `_LINEAR_`, gives the variable name of the predictor series. In this case, the variable is a time index that the system computes. This variable will be included in the output forecast data set. The final part, Linear Trend, describes the predictor.

Notice that the model you have specified consists only of the time index regressor `_LINEAR_` and an intercept. Although this window is normally used to specify ARIMA models, in this case no ARIMA model options are specified, and the model is a simple regression on time.

Select the **OK** button. The Linear Trend model is fit and added to the model list in the Develop Models window.

Now bring up the Model Viewer using the **View Model Graphically** icon or the **Model Predictions** item under the **View** pull-down menu or toolbar. This displays a plot of the model predictions and actual series values, as shown in Display 27.4. The predicted values lie along the least squares trend line.

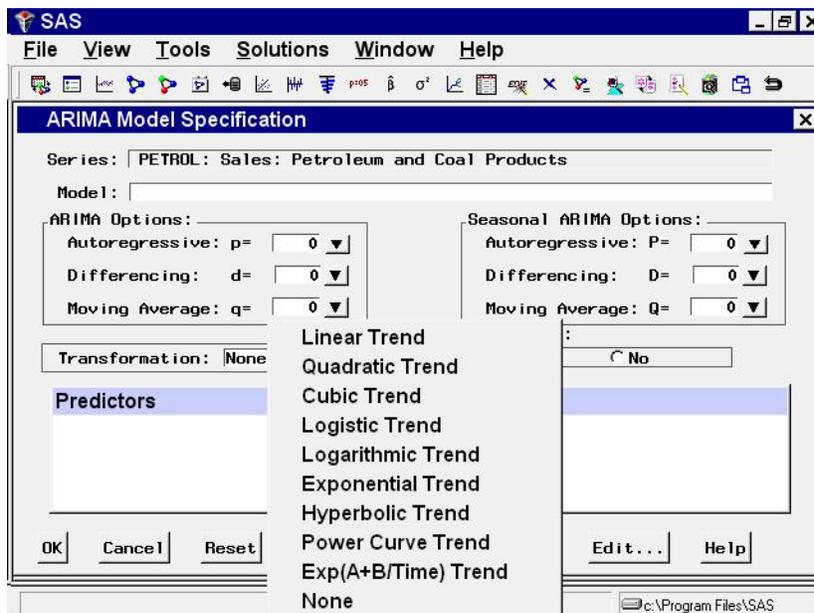
Display 27.4. Linear Trend Model



Time Trend Curves

From the Develop Models window, select **Fit ARIMA Model**. From the ARIMA Model Specification window, select **Add** and then select **Trend Curve** from the menu (shown in Display 27.1). A menu of different kinds of trend curves is displayed, as shown in Display 27.5.

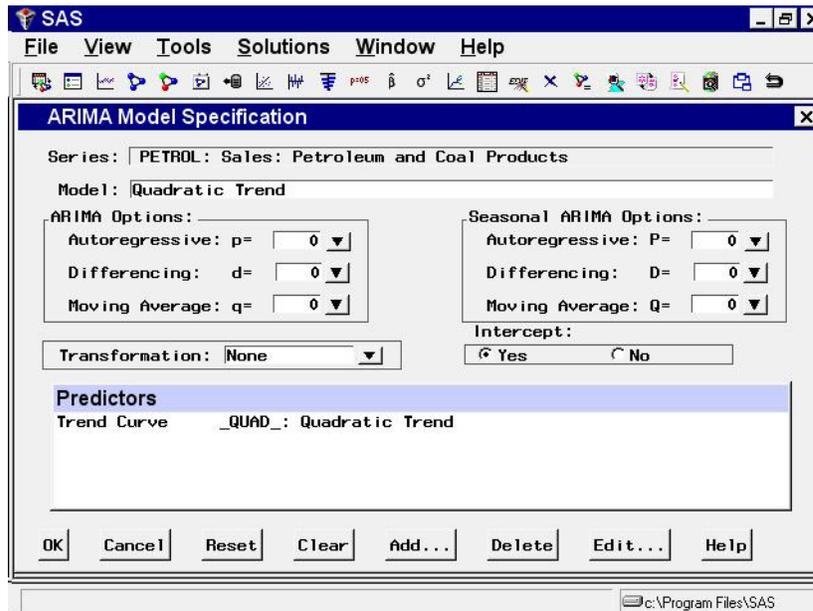
Display 27.5. Time Trend Curves Menu



These trend curves work in a similar way as the Linear Trend option (which is a special case of a trend curve and one of the choices on the menu), but with the Trend Curve menu you have a choice of various nonlinear time trends.

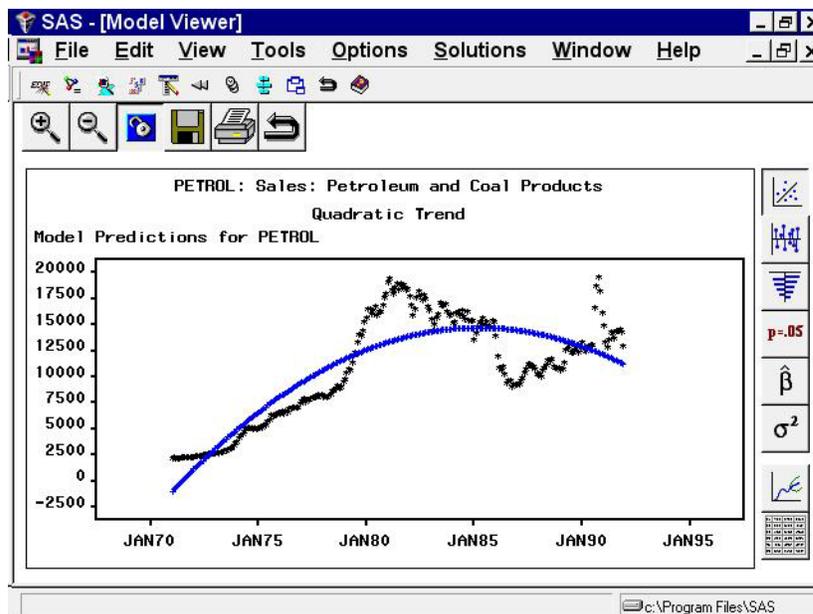
Select Quadratic Trend. This adds a quadratic time trend to the Predictors list, as shown in Display 27.6.

Display 27.6. Quadratic Trend Specified



Now select the OK button. The quadratic trend model is fit and added to the list of models in the Develop Models window. The Model Viewer displays a plot of the quadratic trend model, as shown in Display 27.7.

Display 27.7. Quadratic Trend Model



This curve does not fit the PETROL series very well, but the View Model plot illustrates how time trend models work. You may want to experiment with different trend models to see what the different trend curves look like.

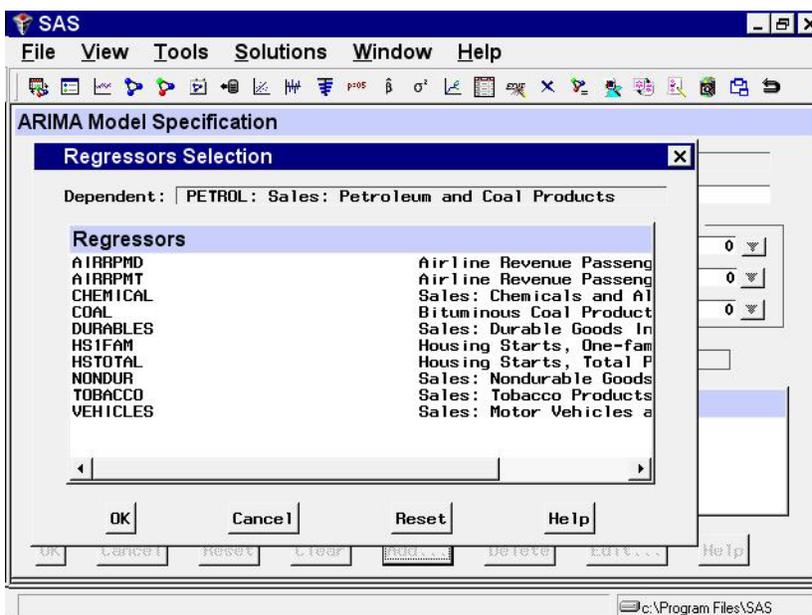
Some of the trend curves require transforming the dependent series. When you specify one of these curves, a notice is displayed reminding you that a transformation is needed, and the Transformation field is automatically filled in. Therefore, you cannot control the Transformation specification when some kinds of trend curves are specified.

See the section "Time Trend Curves" in Chapter 30, "Forecasting Process Details," for more information about the different trend curves.

Regressors

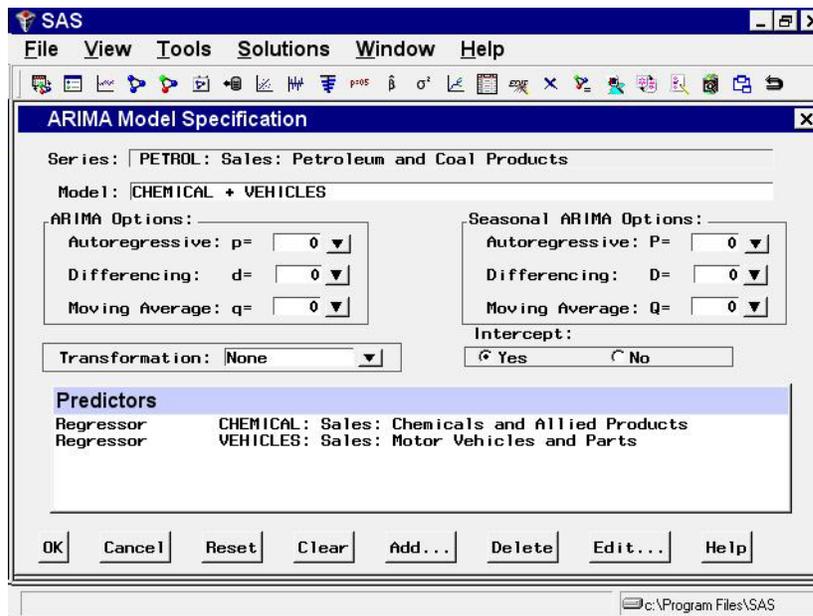
From the Develop Models window, select **Fit ARIMA Model**. From the ARIMA Model Specification window, select **Add** and then select **Regressors** from the menu (shown in Display 27.1). This displays the **Regressors Selection** window, as shown in Display 27.8. This window allows you to select any number of other series in the input data set as regressors to predict the dependent series.

Display 27.8. Regressors Selection Window



For this example, select **CHEMICAL, Sales: Chemicals and Allied Products**, and **VEHICLES, Sales: Motor Vehicles and Parts**. (Note: You do not need to use the CTRL key when selecting more than one regressor.) Then select the **OK** button. The two variables you selected are added to the Predictors list as regressor type predictors, as shown in Display 27.9.

Display 27.9. Regressors Selected

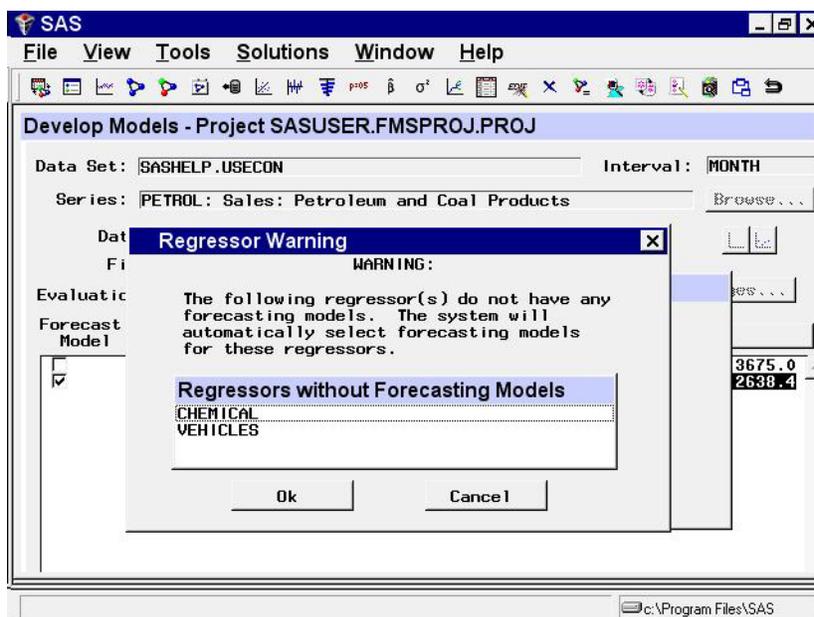


You must have forecasts of the future values of the regressor variables in order to use them as predictors. To do this, you can specify a forecasting model for each regressor, have the system automatically select forecasting models for the regressors, or supply predicted future values for the regressors in the input data set.

Even if you have supplied future values for a regressor variable, the system requires a forecasting model for the regressor. Future values that you supply in the input data set will take precedence over predicted values from the regressor's forecasting model when the system computes the forecasts for the dependent series.

Select the OK button. The system starts to fit the regression model but then stops and displays a warning that the regressors that you selected do not have forecasting models, as shown in Display 27.10.

Display 27.10. Regressors Needing Models Warning



If you want the system to create forecasting models automatically for the regressor variables using the automatic model selection process, select the OK button. If not, you can select the Cancel button to abort fitting the regression model.

For this example, select the OK button. The system now performs the automatic model selection process for CHEMICAL and VEHICLES. The selected forecasting models for CHEMICAL and VEHICLES are added to the model lists for those series. If you switch the current time series in the Develop Models window to CHEMICAL or VEHICLES, you will see the model that the system selected for that series.

Once forecasting models have been fit for all regressors, the system proceeds to fit the regression model for PETROL. The fitted regression model is added to the model list displayed in the Develop Models window.

Adjustments

An *adjustment* predictor is a variable in the input data set that is used to adjust the forecast values produced by the forecasting model. Unlike a regressor, an adjustment variable does not have a regression coefficient. No model fitting is performed for adjustments. Nonmissing values of the adjustment series are simply added to the model prediction for the corresponding period. Missing adjustment values are ignored. If you supply adjustment values for observations within the period of fit, the adjustment values are subtracted from the actual values, and the model is fit to these adjusted values.

To add adjustments, select Add and then select Adjustments from the pop-up menu (shown in Display 27.1). This displays the Adjustments Selection window.

The Adjustments Selection window functions the same as the Regressor Selection window (which is shown in Display 27.8). You can select any number of adjustment variables as predictors.

Unlike regressors, adjustments do not require forecasting models for the adjustment variables. If a variable that is used as an adjustment does have a forecasting model fit to it, the adjustment variable's forecasting model is ignored when the variable is used as an adjustment.

You can use forecast adjustments to account for expected future events that have no precedent in the past and so cannot be modeled by regression. For example, suppose you are trying to forecast the sales of a product, and you know that a special promotional campaign for the product is planned during part of the period you want to forecast. If such sales promotion programs have been frequent in the past, then you can record the past and expected future level of promotional efforts in a variable in the data set and use that variable as a regressor in the forecasting model.

However, if this is the first sales promotion of its kind for this product, you have no way to estimate the effect of the promotion from past data. In this case, the best you can do is to make an educated guess at the effect the promotion will have and add that guess to what your forecasting model would predict in the absence of the special sales campaign.

Adjustments are also useful for making judgmental alterations to forecasts. For example, suppose you have produced forecast sales data for the next 12 months. Your supervisor believes that the forecasts are too optimistic near the end and asks you to prepare a forecast graph in which the numbers that you have forecast are reduced by 1000 in the last three months. You can accomplish this task by editing the input data set so that it contains observations for the actual data range of sales plus 12 additional observations for the forecast period, and a new variable called, for example, ADJUSTMENT. The variable ADJUSTMENT contains the value 1000 for the last three observations, and is missing for all other observations. You fit the same model previously selected for forecasting using the ARIMA Model Specification or Custom Model Specification window, but with an adjustment added using the variable ADJUSTMENT. Now when you graph the forecasts using the Model Viewer, the last three periods of the forecast are reduced by 1000. The confidence limits are unchanged, which helps draw attention to the fact that the adjustments to the forecast deviate from what would be expected statistically.

Dynamic Regressor

Selecting `Dynamic Regressor` from the `Add Predictors` menu (shown in Display 27.1) allows you to specify a complex time series model of the way that a predictor variable influences the series that you are forecasting.

When you specify a predictor variable as a simple regressor, only the current period value of the predictor effects the forecast for the period. By specifying the predictor with the `Dynamic Regression` option, you can use past values of the predictor series, and you can model effects that take place gradually.

Part 3. General Information

Dynamic regression models are an advanced feature that you are unlikely to find useful unless you have studied the theory of statistical time series analysis. You may want to skip this section if you are not trained in time series modeling.

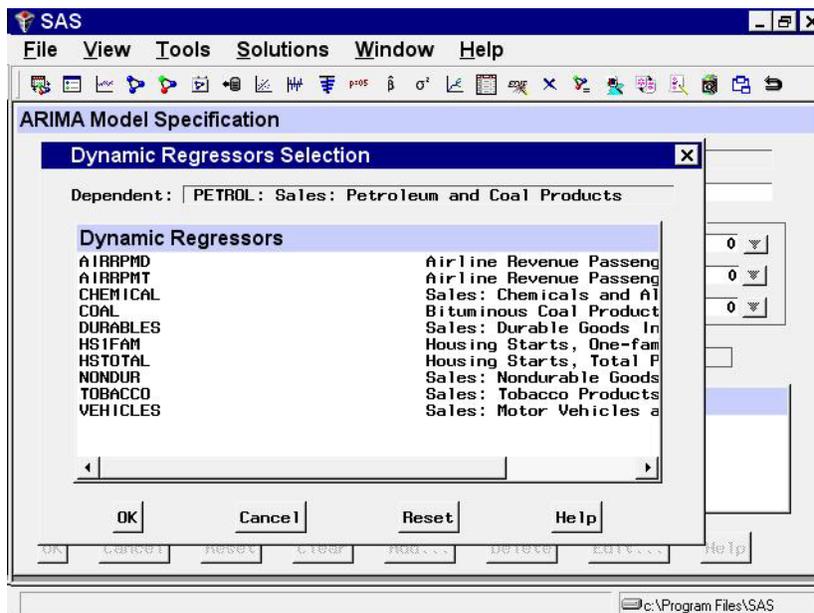
The term *dynamic regression* was introduced by Pankratz (1991) and refers to what Box and Jenkins (1976) named *transfer function models*. In dynamic regression, you have a time series model, similar to an ARIMA model, that predicts how changes in the predictor series affect the dependent series over time.

The dynamic regression model relates the predictor variable to the expected value of the dependent series in the same way that an ARIMA model relates the fluctuations of the dependent series about its conditional mean to the random error term (which is also called the innovation series). Refer to "*Forecasting with Dynamic Regression Models*" (Pankratz, 1991) and "*Time Series Analysis: Forecasting and Control*" (Box and Jenkins, 1976 for more information on dynamic regression or transfer function models. See also Chapter 7, "The ARIMA Procedure."

From the Develop Models window, select **Fit ARIMA Model**. From the ARIMA Model Specification window, select **Add** and then select **Linear Trend** from the menu (shown in Display 27.1).

Now select **Add** and select **Dynamic Regressor**. This displays the Dynamic Regressors Selection window, as shown in Display 27.11.

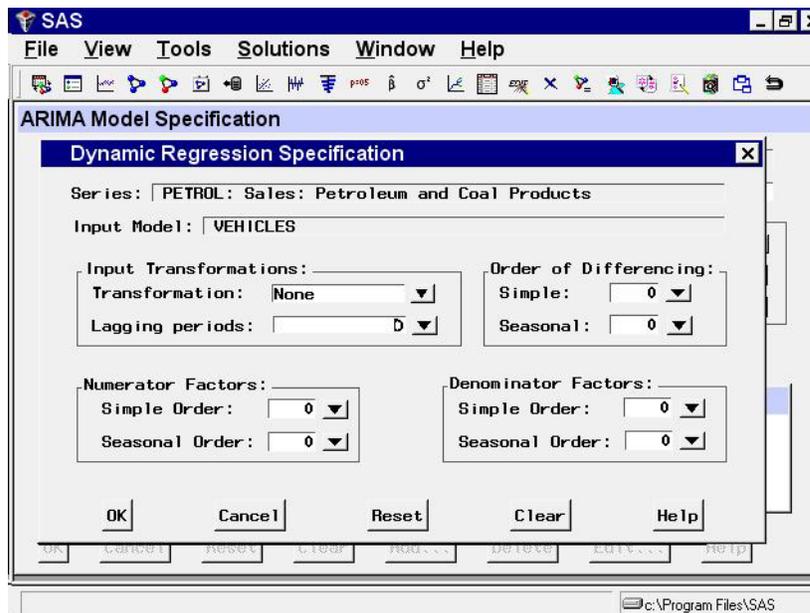
Display 27.11. Dynamic Regressors Selection Window



You can select only one predictor series when specifying a dynamic regression model. For this example, select **VEHICLES**, Sales: Motor Vehicles and Parts. Then select the **OK** button.

This displays the Dynamic Regression Specification window, as shown in Display 27.12.

Display 27.12. Dynamic Regression Specification Window



This window consists of four parts. The `Input Transformations` fields allow you to transform or lag the predictor variable. For example, you might use the lagged logarithm of the variable as the predictor series.

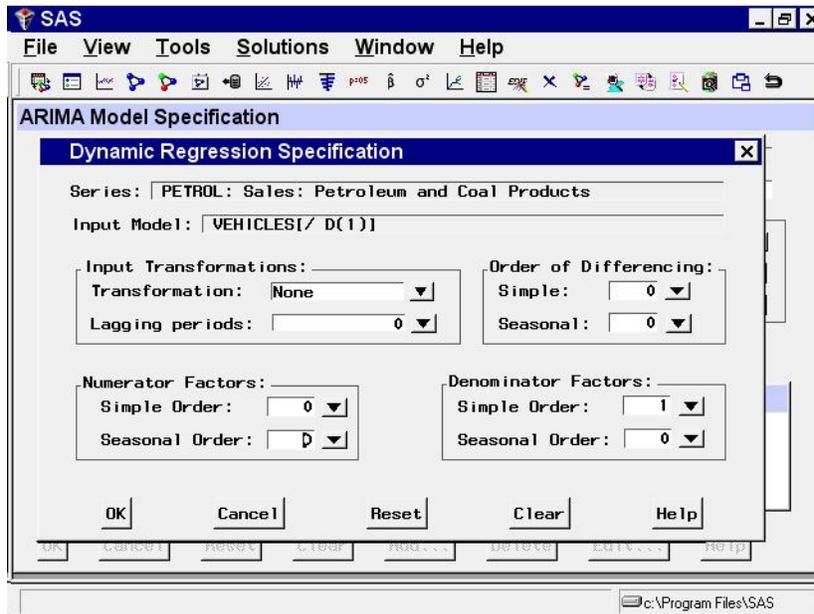
The `Order of Differencing` fields allow you to specify simple and seasonal differencing of the predictor series. For example, you might use changes in the predictor variable instead of the variable itself as the predictor series.

The `Numerator Factors` and `Denominator Factors` fields allow you to specify the orders of simple and seasonal numerator and denominator factors of the transfer function.

Simple regression is a special case of dynamic regression in which the dynamic regression model consists of only a single regression coefficient for the current value of the predictor series. If you select the `OK` button without specifying any options in the `Dynamic Regression Specification` window, a simple regressor will be added to the model.

For this example, use the `Simple Order` combo box for `Denominator Factors` and set its value to 1. The window now appears as shown in Display 27.13.

Display 27.13. Distributed Lag Regression Specified

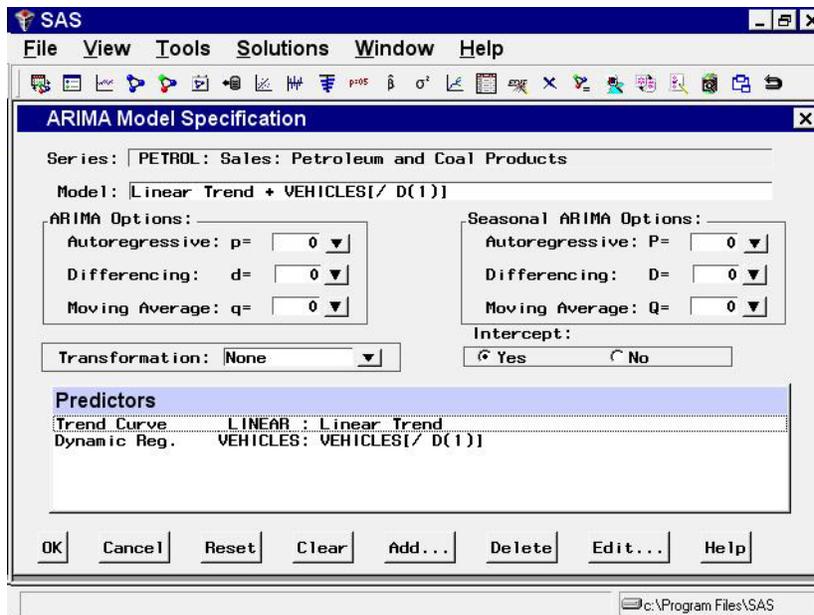


This model is equivalent to regression on an exponentially weighted infinite distributed lag of VEHICLES (in the same way an MA(1) model is equivalent to single exponential smoothing).

Select the OK button to add the dynamic regressor to the model predictors list.

In the ARIMA Model Specification window, the Predictors list should now contain two items, a linear trend and a dynamic regressor for VEHICLES, as shown in Display 27.14.

Display 27.14. Dynamic Regression Model



This model is a multiple regression of PETROL on a time trend variable and an infinite distributed lag of VEHICLES. Select the OK button to fit the model.

As with simple regressors, if VEHICLES does not already have a forecasting model, an automatic model selection process is performed to find a forecasting model for VEHICLES before the dynamic regression model for PETROL is fit.

Interventions

An *intervention* is a special indicator variable, computed automatically by the system, that identifies time periods affected by unusual events that influence or intervene in the normal path of the time series you are forecasting. When you add an intervention predictor, the indicator variable of the intervention is used as a regressor, and the impact of the intervention event is estimated by regression analysis.

To add an intervention to the Predictors list, you must use the Intervention Specification window to specify the time or times that the intervening event took place and to specify the type of intervention. You can add interventions either through the Interventions item of the Add action or by selecting Tools from the menu bar and then selecting Define Interventions.

Intervention specifications are associated with the series. You can specify any number of interventions for each series, and once you define interventions you can select them for inclusion in forecasting models. If you select the Include Interventions option in the Options pull-down menu, any interventions that you have previously specified for a series are automatically added as predictors to forecasting models for the series.

From the Develop Models window, invoke the series viewer by selecting the View Series Graphically icon or Series under the View pull-down menu. This displays the Time Series Viewer, as was shown in Display 27.2.

Note that the trend in the PETROL series shows several clear changes in direction. The upward trend in the first part of the series reverses in 1981. There is a sharp drop in the series towards the end of 1985, after which the trend is again upwardly sloped. Finally, in 1991 the series takes a sharp upward excursion but quickly returns to the trend line.

You may have no idea what events caused these changes in the trend of the series, but you can use these patterns to illustrate the use of intervention predictors. To do this, you fit a linear trend model to the series, but modify that trend line by adding intervention effects to model the changes in trend you observe in the series plot.

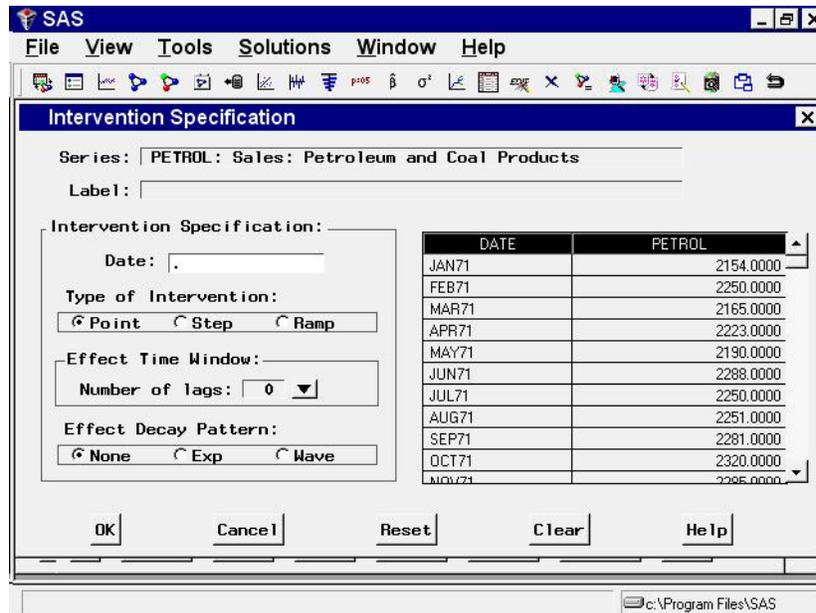
The Intervention Specification Window

From the Develop Models window, select Fit ARIMA model. From the ARIMA Model Specification window, select Add and then select Linear Trend from the menu (shown in Display 27.1).

Select Add again and then select Interventions. If you have any interventions already defined for the series, this selection displays the Interventions for

Series window. However, since you have not previously defined any interventions, this list is empty. Therefore, the system assumes that you want to add an intervention and displays the Intervention Specification window instead, as shown in Display 27.15.

Display 27.15. Interventions Specification Window



The top of the Intervention Specification window shows the current series and the label for the new intervention (initially blank). At the right side of the window is a scrollable table showing the values of the series. This table helps you locate the dates of the events you want to model.

At the left of the window is an area titled Intervention Specification that contains the options for defining the intervention predictor. The Date field specifies the time that the intervention occurs. You can type a date value in the Date field, or you can set the Date value by selecting a row from the table of series values at the right side of the window.

The area titled Type of Intervention controls the kind of indicator variable constructed to model the intervention effect. You can specify the following kinds of interventions:

- Point** is used to indicate an event that occurs in a single time period. An example of a point event is a strike that shuts down production for part of a time period. The value of the intervention's indicator variable is zero except for the date specified.
- Step** is used to indicate a continuing event that changes the level of the series. An example of a step event is a change in the law, such as a tax rate increase. The value of the intervention's indicator variable is zero before the date specified and 1 thereafter.

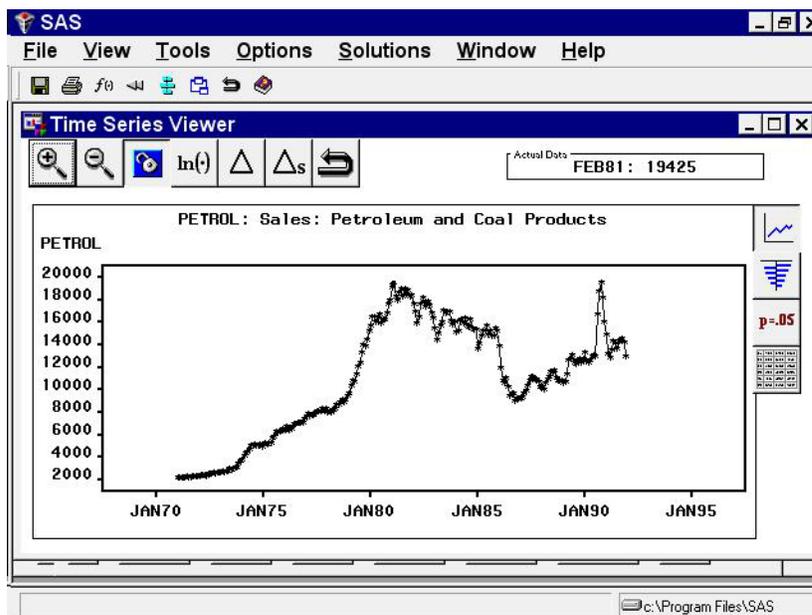
Ramp is used to indicate a continuing event that changes the trend of the series. The value of the intervention's indicator variable is zero before the date specified, and it increases linearly with time thereafter.

The areas titled **Effect Time Window** and **Effect Decay Pattern** specify how to model the effect that the intervention has on the dependent series. These options are not used for simple interventions, they will be discussed later in this chapter.

Specifying a Trend Change Intervention

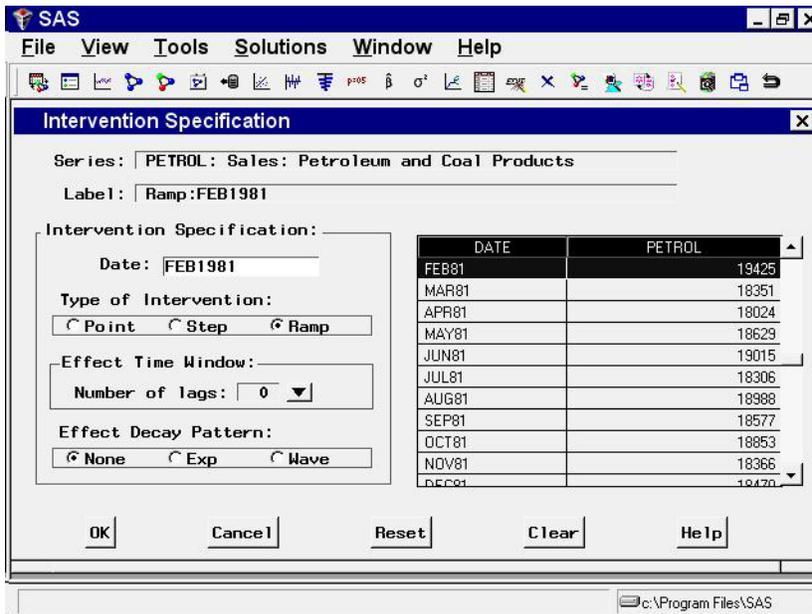
In the Time Series Viewer window position the mouse over the highest point in 1981 and select the point. This displays the data value, 19425, and date, February 1981, of that point in the upper-right corner of the Time Series Viewer, as shown in Display 27.16.

Display 27.16. Identifying the Turning Point



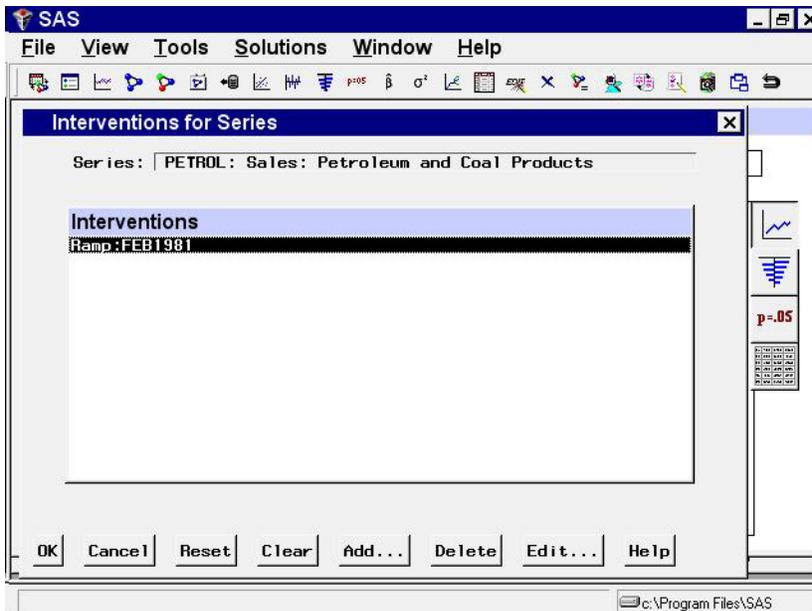
Now that you know the date that the trend reversal occurred, enter that date in the **Date** field of the Intervention Specification window. Select **Ramp** as the type of intervention. The window should now appear as shown in Display 27.17.

Display 27.17. Ramp Intervention Specified



Select the OK button. This adds the intervention to the list of interventions for the PETROL series, and returns you to the Interventions for Series window, as shown in Display 27.18.

Display 27.18. Interventions for Series Window



This window allows you to select interventions for inclusion in the forecasting model. Since you need to define other interventions, select the Add button. This returns you to the Intervention Specification window (shown in Display 27.15).

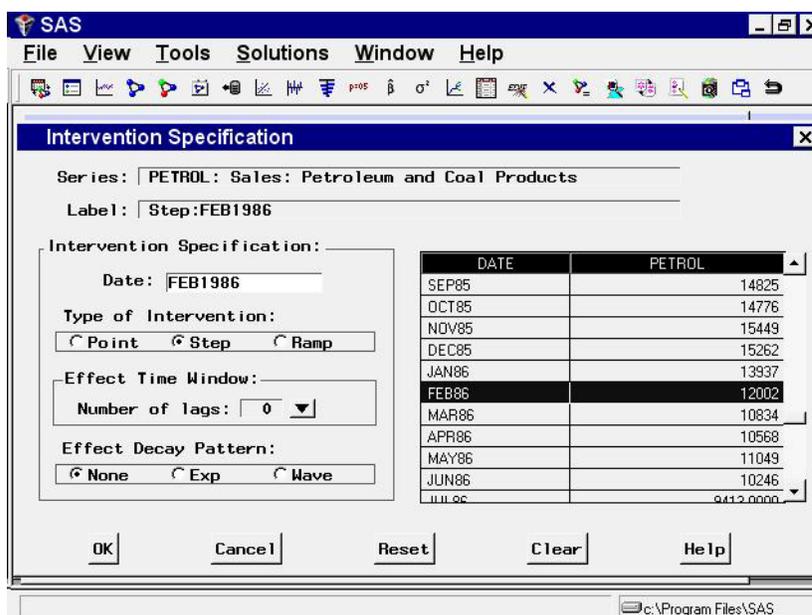
Specifying a Level Change Intervention

Now add an intervention to account for the drop in the series in late 1985. You can locate the date of this event by selecting points in the Time Series Viewer plot or by scrolling through the data values table in the Interventions Specification window. Use the latter method so that you can see how this works.

Scrolling through the table, you see that the drop was from 15262 in December 1985, to 13937 in January 1986, to 12002 in February, to 10834 in March. Since the drop took place over several periods, you could use another ramp type intervention. However, this example represents the drop as a sudden event using a step intervention and uses February 1986 as the approximate time of the drop.

Select the table row for February 1986 to set the `Date` field. Select `Step` as the intervention type. The window should now appear as shown in Display 27.19.

Display 27.19. Step Intervention Specified



Select the `OK` button to add this intervention to the list for the series.

Since the trend reverses again after the drop, add a ramp intervention for the same date as the step intervention. Select `Add` from the Interventions for Series window. Enter `FEB86` in the `Date` field, select `Ramp`, and then select the `OK` button.

Modeling Complex Intervention Effects

You have now defined three interventions to model the changes in trend and level. The excursion near the end of the series remains to be dealt with.

Select `Add` from the Interventions for Series window. Scroll through the data values and select the date on which the excursion began, August 1990. Leave the intervention type as `Point`.

Part 3. General Information

The pattern of the series from August 1990 through January 1991 is more complex than a simple shift in level or trend. For this pattern, you need a complex intervention model for an event that causes a sharp rise followed by a rapid return to the previous trend line. To specify this model, use the `Effect Time Window` and `Effect Decay Rate` options.

The `Effect Time Window` option controls the number of lags of the intervention's indicator variable used to model the effect of the intervention on the dependent series. For a simple intervention, the number of lags is zero, which means that the effect of the intervention is modeled by fitting a single regression coefficient to the intervention's indicator variable.

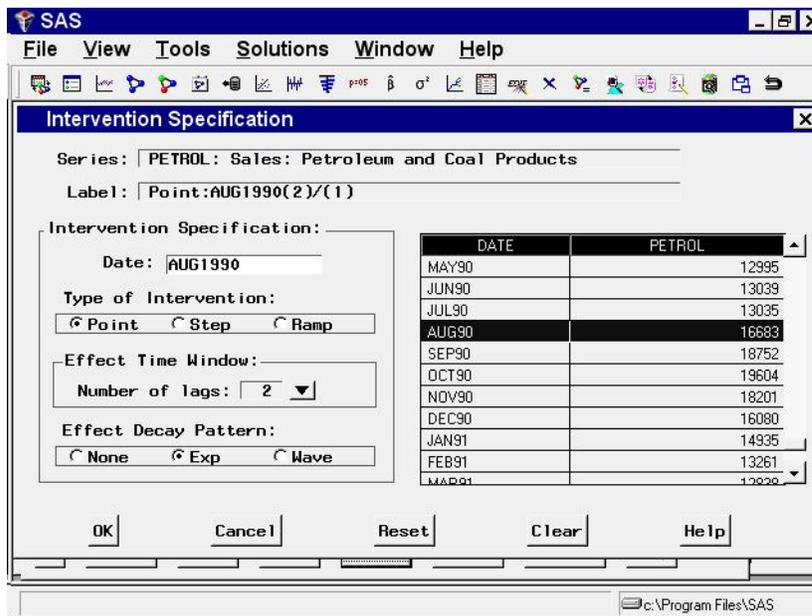
When you set the number of lags greater than zero, regression coefficients are fit to lags of the indicator variable. This allows you to model interventions whose effects take place gradually, or to model rebound effects. For example, severe weather may reduce production during one period but cause an increase in production in the following period as producers struggle to catch up. You could model this using a point intervention with an effect time window of 1 lag. This would fit two coefficients for the intervention, one for the immediate effect and one for the delayed effect.

The `Effect Decay Pattern` option controls how the effect of the intervention dissipates over time. `None` specifies that there is no gradual decay: for point interventions, the effect ends immediately; for step and ramp interventions, the effect continues indefinitely. `Exp` specifies that the effect declines at an exponential rate. `Wave` specifies that the effect declines like an exponentially damped sine wave (or as the sum of two exponentials, depending on the fit to the data).

If you are familiar with time series analysis, these options may be clearer if you note that together the `Effect Time Window` and `Effect Decay Pattern` options define the numerator and denominator orders of a transfer function or dynamic regression model for the indicator variable of the intervention. See the section "Dynamic Regressor" later in this chapter for more information.

For this example, select 2 lags as the value of the `Event Time Window` option, and select `Exp` as the `Effect Decay Pattern` option. The window should now appear as shown in Display 27.20.

Display 27.20. Complex Intervention Model

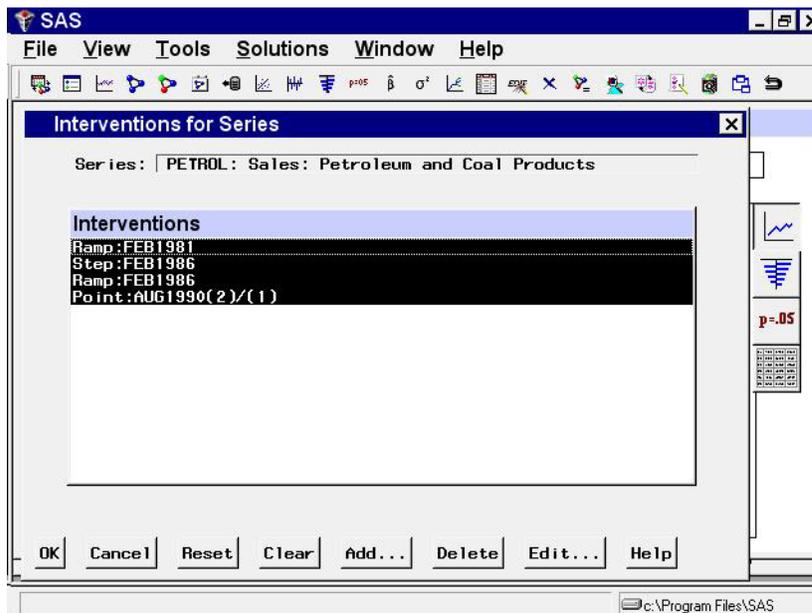


Select the OK button to add the intervention to the list.

Fitting the Intervention Model

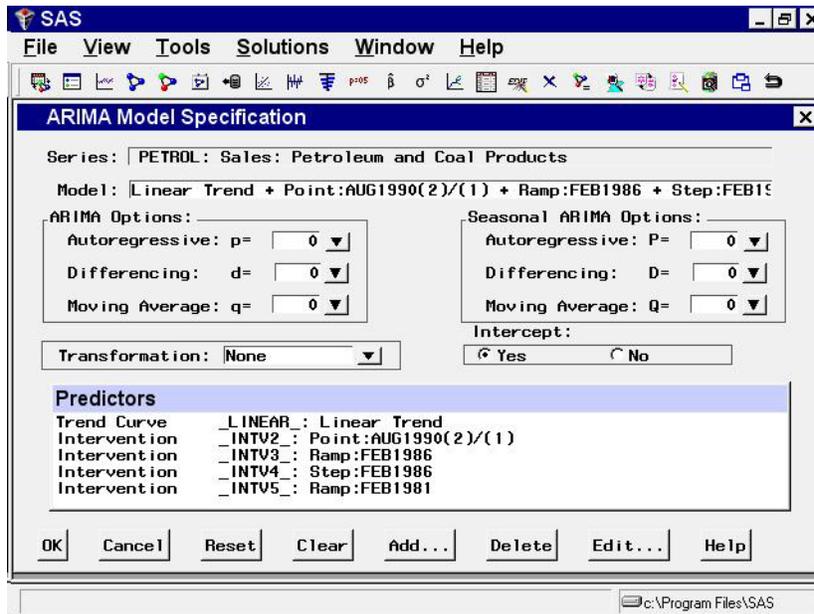
The Interventions for Series window now contains definitions for four intervention predictors. Select all four interventions, as shown in Display 27.21.

Display 27.21. Interventions for Series Window



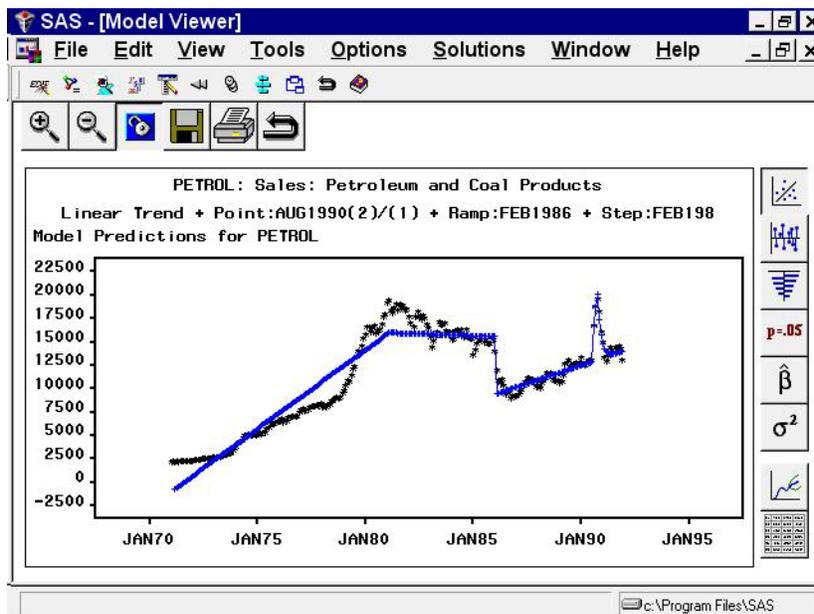
Select the OK button. This returns you to the ARIMA Model Specification window, which now lists items in the Predictors list, as shown in Display 27.22.

Display 27.22. Linear Trend with Interventions Specified



Select the OK button to fit this model. After the model is fit, bring up the Model Viewer. You will see a plot of the model predictions, as shown in Display 27.23.

Display 27.23. Linear Trend with Interventions Model



You may wish to use the Zoom In feature to take a closer look at how the complex intervention effect fits the excursion in the series starting in August 1990.

Limitations of Intervention Predictors

Note that the model you have just fit is intended only to illustrate the specification of interventions. It is not intended as an example of good forecasting practice.

The use of continuing (step and ramp type) interventions as predictors has some limitations that you should consider. If you model a change in trend with a simple ramp intervention, then the trend in the data before the date of the intervention has no influence on the forecasts. Likewise, when you use a step intervention, the average level of the series before the intervention has no influence on the forecasts.

Only the final trend and level at the end of the series are extrapolated into the forecast period. If a linear trend is the only pattern of interest, then instead of specifying step or ramp interventions, it would be simpler to adjust the period of fit so that the model ignores the data before the final trend or level change.

Step and ramp interventions are valuable when there are other patterns in the data—such as seasonality, autocorrelated errors, and error variance—that are stable across the changes in level or trend. Step and ramp interventions allow you to fit seasonal and error autocorrelation patterns to the whole series while fitting the trend only to the latter part of the series.

Point interventions are a useful tool for dealing with outliers in the data. A point intervention will fit the series value at the specified date exactly, and it has the effect of removing that point from the analysis. When you specify an effect time window, a point intervention will exactly fit as many additional points as the number of lags specified.

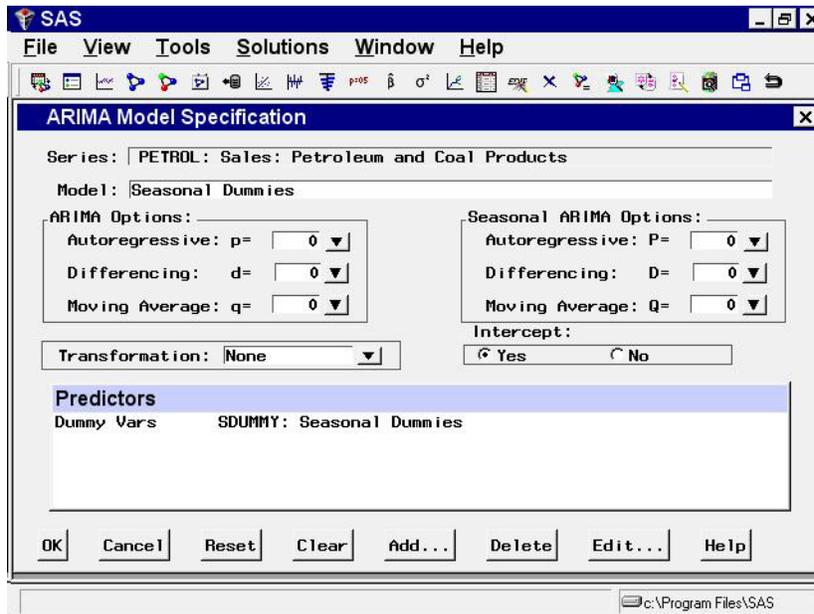
Seasonal Dummies

A *Seasonal Dummies* predictor is a special feature that adds to the model seasonal indicator or "dummy" variables to serve as regressors for seasonal effects.

From the Develop Models window, select `Fit ARIMA Model`. From the ARIMA Model Specification window, select `Add` and then select `Seasonal Dummies` from the menu (shown in Display 27.1).

A Seasonal Dummies input is added to the Predictors list, as shown in Display 27.24.

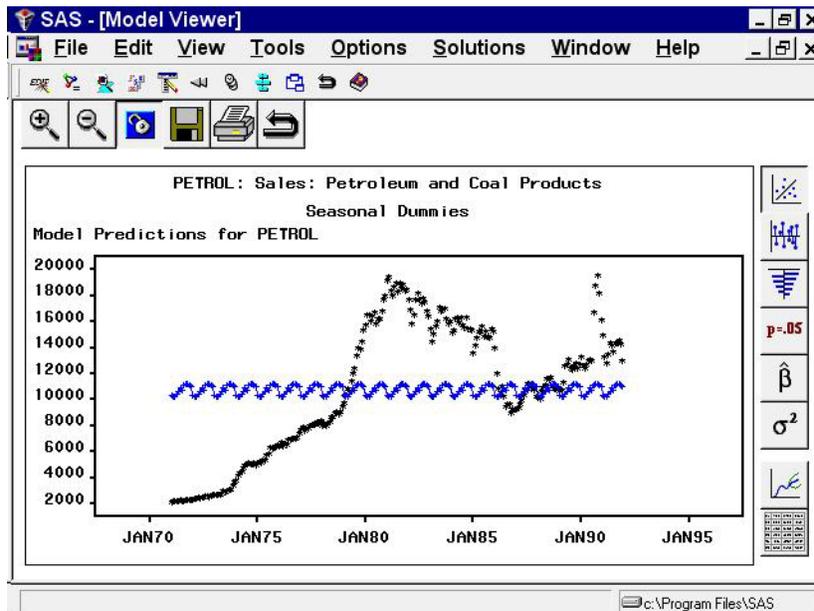
Display 27.24. Seasonal Dummies Specified



Select the OK button. A model consisting of an intercept and 11 seasonal dummy variables is fit and added to the model list in the Develop Models window. This is effectively a mean model with a separate mean for each month.

Now return to the Model Viewer, which displays a plot of the model predictions and actual series values, as shown in Display 27.25. This is obviously a poor model for this series, but it serves to illustrate how seasonal dummy variables work.

Display 27.25. Seasonal Dummies Model



Now select the parameter estimates icon, the fifth from the top on the vertical toolbar. This displays the Parameter Estimates table, as shown in Display 27.26.

Display 27.26. Parameter Estimates for Seasonal Dummies Model

Parameter Estimates
PETROL: Sales: Petroleum and Coal Products
Seasonal Dummies

Model Parameter	Estimate	Std. Error	T	Prob> T
Intercept	11064	1138	9.7229	<.0001
Seasonal Dummy 1	-737.76190	1609	-0.4584	0.6471
Seasonal Dummy 2	-845.71429	1609	-0.5255	0.5997
Seasonal Dummy 3	-867.71429	1609	-0.5392	0.5903
Seasonal Dummy 4	-660.42857	1609	-0.4104	0.6819
Seasonal Dummy 5	-405.61905	1609	-0.2520	0.8012
Seasonal Dummy 6	-111.47619	1609	-0.0693	0.9448
Seasonal Dummy 7	-339.00000	1609	-0.2106	0.8333
Seasonal Dummy 8	29.23810	1609	0.0182	0.9855
Seasonal Dummy 9	160.66667	1609	0.0998	0.9206
Seasonal Dummy 10	115.38095	1609	0.0717	0.9429

Fit Range: JAN1971 to DEC1991

Since the data for this example are monthly, the Seasonal Dummies option added 11 seasonal dummy variables. These include a dummy regressor variable that is 1.0 for January and 0 for other months, a regressor that is 1.0 only for February, and so forth through November.

Because the model includes an intercept, no dummy variable is added for December. The December effect is measured by the intercept, while the effect of other seasons is measured by the difference between the intercept and the estimated regression coefficient for the season's dummy variable.

The same principle applies for other data frequencies: the "Seasonal Dummy 1" parameter will always refer to the first period in the seasonal cycle; and, when an intercept is present in the model, there will be no seasonal dummy parameter for the last period in the seasonal cycle.

References

- Box, G.E.P. and Jenkins, G.M. (1976), *Time Series Analysis: Forecasting and Control*, San Francisco: Holden-Day.
- Pankratz, Alan (1991), *Forecasting with Dynamic Regression Models*, New York: John Wiley & Sons, Inc.

The correct bibliographic citation for this manual is as follows: SAS Institute Inc., *SAS/ETS User's Guide, Version 8*, Cary, NC: SAS Institute Inc., 1999. 1546 pp.

SAS/ETS User's Guide, Version 8

Copyright © 1999 by SAS Institute Inc., Cary, NC, USA.

ISBN 1-58025-489-6

All rights reserved. Printed in the United States of America. 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.

U.S. Government Restricted Rights Notice. Use, duplication, or disclosure of the software by the government is subject to restrictions as set forth in FAR 52.227-19 Commercial Computer Software-Restricted Rights (June 1987).

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

1st printing, October 1999

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 registered trademarks or trademarks of their respective companies.

The Institute is a private company devoted to the support and further development of its software and related services.