

Chapter 466

Exponential Smoothing – Trend

Introduction

This module forecasts series with upward or downward trends. Three techniques are available: least squares trend, double smoothing, and Holt's linear trend algorithm.

Least Squares Trend

Least squares trend computes a straight-line trend equation through the data using standard least squares techniques in which the dependent variable is the time series, and the independent variable is the row (sequence) number. The forecasting equation is

$$F_t = a + bt$$

where F_t is the forecast at time period t , a is the y-intercept, and b is the slope of the trend. The slope indicates how much is added (or subtracted if b is negative) from each time period to the next.

This method is useful for series that show a stable, long-term trend. It places the largest weights in estimation on the two ends of the series, while the rows near the middle with an insignificant impact on the estimates.

Double Exponential Smoothing

Double exponential smoothing computes a trend equation through the data using a special weighting function that places the greatest emphasis on the most recent time periods. The forecasting equation changes from period to period.

The forecasting algorithm makes use of the following formulas:

$$F_t = a_t + b_t$$

$$a_t = X_t + (1 - \alpha)^2 e_t$$

$$b_t = b_{t-1} + \alpha^2 e_t$$

$$e_t = F_t - X_t$$

The smoothing constant, α , dictates the amount of smoothing that takes place. It ranges from zero to one.

The forecast at time period T for the value at time period $T+k$ is $a_T + b_T k$. Double smoothing is discussed in detail in Thomopoulos (1980).

This method is included more for its historical significance since Holt's algorithm is usually preferred to it.

Holt's Linear Trend

Holt's Linear Trend computes an evolving trend equation through the data using a special weighting function that places the greatest emphasis on the most recent time periods. Instead of the global trend equation of the least squares trend algorithm, this technique uses a local trend equation. The trend equation is modified from period to period. The forecasting equation changes from period to period.

The forecasting algorithm makes use of the following formulas:

$$a_t = \alpha X_t + (1 - \alpha)(a_{t-1} + b_{t-1})$$

$$b_t = \beta(a_t - a_{t-1}) + (1 - \beta)b_{t-1}$$

Here α and β are smoothing constants which are each between zero and one. Again, a_t gives the y-intercept (or level) at time t , while b_t is the slope at time t .

The forecast at time T for the value at time $T+k$ is $a_T + b_T k$.

Smoothing Constants

Notice that in both double smoothing and Holt's linear trend, the *smoothing constant(s)* determines how fast the weights of the series decays. The values may be chosen either subjectively or objectively. Values of a smoothing constant near one put almost all weight on the most recent observations. Values of a smoothing constant near zero allow the distant past observations to have a large influence.

When selecting the smoothing constant *subjectively*, you use your own experience with this, and similar, series. Also, specifying the smoothing constant yourself lets you tune the forecast to your own beliefs about the future of the series. If you believe that the mechanism generating the series has recently gone through some fundamental changes, use a smoothing constant value of 0.9 which will cause distant observations to be ignored. If, however, you think the series is fairly stable and only going through random fluctuations, use a value of 0.1.

To select the value of the smoothing constant(s) *objectively*, you search for values that are best in some sense. Our program searches for those values that minimizes the size of the combined forecast errors of the currently available series. Three methods of summarizing the amount of error in the forecasts are available: the mean square error (MSE), the mean absolute error (MAE), and the mean absolute percent error (MAPE). The forecast error is the difference between the forecast of the current period made at the last period and the value of the series at the current period. This is written as

$$e_t = X_t - F_{t-1}$$

Using this formulation, we can define the three error-size criterion as follows:

$$MSE = \frac{1}{n} \sum e_t^2$$

$$MAE = \frac{1}{n} \sum |e_t|$$

$$MAPE = \frac{100}{n} \sum \left| \frac{e_t}{X_t} \right|$$

Exponential Smoothing – Trend

To find the value of the smoothing constants objectively, we select one of these criterion and search for those values of α and β that minimize this function. The program conducts a search for the appropriate values using an efficient grid-searching algorithm.

Initial Values

Both double smoothing and Holt's linear trend require initialization since the forecast for period one requires the forecast at period zero, which we do not, by definition, have. Several methods have been proposed for generating starting values. We have adopted the backcasting method which is currently considered to be one of the best methods. Backcasting is simply reversing the series so that we forecast into the past instead of into the future. This produces the required starting values for the slope and intercept. Once we have done this, we can then switch the series back and apply the algorithm in the regular manner.

Relationship to ARIMA Method

It can be shown that both double exponential smoothing and Holt's linear trend technique are equivalent to the ARIMA(0,2,2) model (see Kendall and Ord (1990) page 133). This is why backcasting is recommended for initial values.

Assumptions and Limitations

These algorithms are useful for forecasting non-seasonal time series with (local or global) trend.

Data Structure

The data are entered in a single variable.

Missing Values

When missing values are found in the series, they are either replaced or omitted. The replacement value is the average of the nearest observation in the future and in the past or the nearest non-missing value in the past.

If you do not feel that this is a valid estimate of the missing value, you should manually enter a more reasonable estimate before using the algorithm. These missing value replacement methods are particularly poor for seasonal data. We recommend that you replace missing values manually before using the algorithm.

Example 1 – Trend Exponential Smoothing

This section presents an example of how to generate a forecast of a series using Holt's linear trend. The data in the Intel dataset gives price and volume data for Intel stock during August, 1995. We will forecast values for daily volumes. These values are contained in the variable Intel_Volume.

Setup

To run this example, complete the following steps:

1 Open the Intel example dataset

- From the File menu of the NCSS Data window, select **Open Example Data**.
- Select **Intel** and click **OK**.

2 Specify the Exponential Smoothing – Trend procedure options

- Find and open the **Exponential Smoothing – Trend** procedure using the menus or the Procedure Navigator.
- The settings for this example are listed below and are stored in the **Example 1** settings file. To load these settings to the procedure window, click **Open Example Settings File** in the Help Center or File menu.

Variables Tab

Time Series Variable(s)**Intel_Volume**
Forecast Method.....**Holt's Linear Trend**

Report Tab

Forecast Report.....**Data and Forecasts**

3 Run the procedure

- Click the **Run** button to perform the calculations and generate the output.

Forecast Summary Section

Forecast Summary Section

Variable	Intel_Volume
Number of Rows	20
Mean	10974.54
Missing Values	None
Pseudo R-Squared	0.000000
Mean Square Error	1.801067E+07
Mean Error	3229.632
Mean Percent Error	28.64764
Forecast Method	Holt's Linear Trend
Search Iterations	72
Search Criterion	Mean Square Error
Alpha	0.4151869
Beta	0.1181962
Intercept (A)	9274.478
Slope (B)	211.0687

This report summarizes the forecast equation.

Variable

The name of the variable for which the forecasts are generated.

Number of Rows

The number of rows that were in the series. This is provided to allow you to double-check that the correct series was used.

Missing Values

If missing values were found, this option lists the method used to estimate them.

Mean

The mean of the variable across all time periods.

Pseudo R-Squared

This value generates a statistic that acts like the R-Squared value in multiple regression. A value near zero indicates a poorly fitting model, while a value near one indicates a well-fitting model. The statistic is calculated as follows:

$$R^2 = 100 \left(1 - \frac{SSE}{SST} \right)$$

where *SSE* is the sum of square residuals and *SST* is the total sum of squares after correcting for the mean.

Mean Square Error

The average squared residual (MSE) is a measure of how closely the forecasts track the actual data. The statistic is popular because it shows up in analysis of variance tables. However, because of the squaring, it tends to exaggerate the influence of outliers (points that do not follow the regular pattern).

Exponential Smoothing – Trend

Mean | Error |

The average absolute residual (MAE) is a measure of how closely the forecasts track the actual data without the squaring.

Mean | Percent Error |

The average percent absolute residual (MAPE) is a measure of how closely the forecasts track the actual data put on a percentage basis.

Forecast Method

This line shows which of the three possible trend forecasting algorithms was selected.

Search Iterations

This line shows how many iterations were needed to find the best value(s) for the smoothing constant(s).

Search Criterion

If a search was made to find the best values of the smoothing constants, this row gives the criterion used during the search.

Alpha

The value of the smoothing constant alpha that was used to generate the forecasts.

Beta

The value of the smoothing constant beta that was used to generate the forecasts.

Intercept (A)

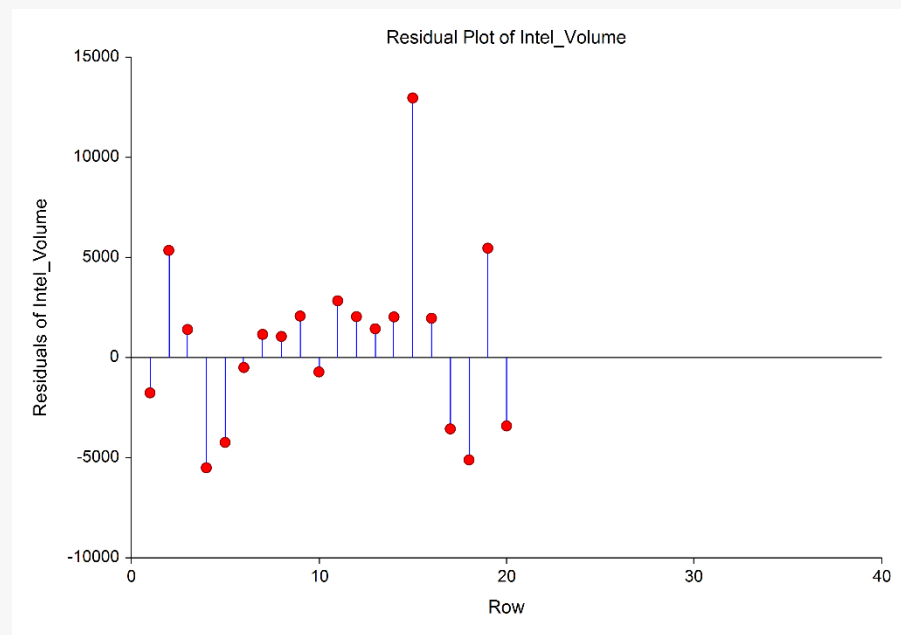
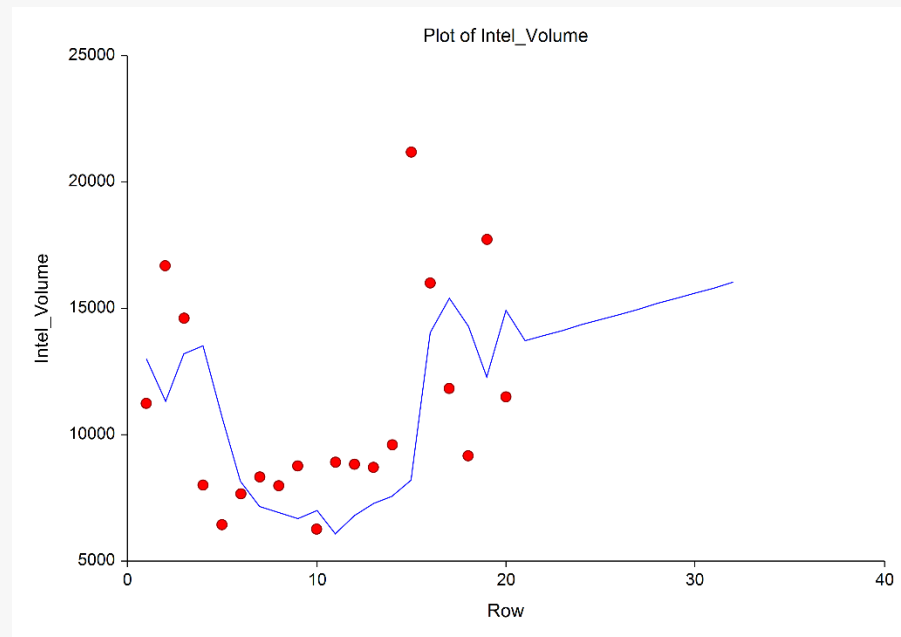
The value of the y-intercept for time period one! Hence, to forecast for time period 21 (the next period after the current period) we would use $9277.523 + 21(210.8949) = 13706.32$.

Slope (B)

The value of the slope.

Forecast and Residuals Plots

Forecast and Residuals Plots



Forecast Plot

The forecast plot lets you analyze how closely the forecasts track the data. The plot also shows the forecasts at the end of the data series.

Residual Plot

This plot lets you analyze the residuals themselves. You are looking for patterns, outliers, or any other information that may help you improve the forecasting model. The first thing to compare is the scale of the Residual Plot versus the scale of the Forecast Plot. If your forecasting is working well, the vertical scale of the Residual Plot will be much less than the scale of the Forecast Plot.

Forecasts Section

Forecasts Section			
Row No.	Forecast Intel_Volume	Actual Intel_Volume	Residuals
1	13002.52	11242.2	-1760.32
2	11335.18	16689.9	5354.722
3	13208.4	14613.3	1404.899
4	13510.65	8009	-5501.653
5	10675.41	6441.8	-4233.61
6	8158.883	7664.5	-494.3835
7	7170.574	8330.3	1159.726
8	6925.94	7983	1057.06
9	6690.555	8767.1	2076.545
10	6980.35	6266.4	-713.9496
11	6076.532	8915.3	2838.768
12	6787.064	8833	2045.936
13	7268.824	8709.7	1440.876
14	7570.08	9603	2032.92
15	8216.907	21185.2	12968.29
16	14040.36	16006.5	1966.143
17	15392.34	11832.4	-3559.945
18	14275.27	9168.1	-5107.174
19	12265.19	17729.3	5464.113
20	14912.3	11500.7	-3411.603
21	13706.92		
22	13917.99		
23	14129.06		
24	14340.13		
25	14551.19		
26	14762.26		
27	14973.33		
28	15184.4		
29	15395.47		
30	15606.54		
31	15817.61		
32	16028.67		

This section shows the values of the forecasts, the actual values, and the residuals.