

جزوه آموزشی لاتین نرم افزار Minitab در سری های زمانی



گروه دانش آماری

مهرماه ۸۶

منبع : www.Gigapedia.com

با توجه به پیشرفت روز افزون علم آمار در زمینه سری های زمانی ، لزوم یادگیری نرم افزارهای موجود در این شاخه برای دانشجویان این رشته به شدت احساس می شود. لذا بایستی تمامی دانشجویان رشته های آمار و علوم اقتصادی برای یادگیری این گونه نرم افزار ها اقدام کنند . این جزوه آموزشی در حل مسائل سری زمانی با کمک نرم افزار **Minitab** ، کمک بسیاری را به شما خواننده محترم می کند . امیدوارم به کمک این جزوه بتوانید به دانسته های خود مطالب مفیدی را بیفزایید .

Table Of Contents

Time Series.....	5
Time Series Overview	5
Time Series Plot.....	8
Trend Analysis.....	15
Decomposition.....	20
Moving Average.....	29
Single Exponential Smoothing	33
Double Exponential Smoothing	37
Winters' Method.....	42
Differences	48
Lag.....	49
Autocorrelation	49
Partial Autocorrelation	52
Cross Correlation.....	53
ARIMA	54
References - Time Series.....	60
Index	61

Time Series

Time Series Overview

Minitab's time series procedures can be used to analyze data collected over time, commonly called a time series. These procedures include simple forecasting and smoothing methods, correlation analysis methods, and ARIMA modeling. Although correlation analysis may be performed separately from ARIMA modeling, we present the correlation methods as part of ARIMA modeling.

Simple forecasting and smoothing methods are based on the idea that reliable forecasts can be achieved by modeling patterns in the data that are usually visible in a time series plot, and then extrapolating those patterns to the future. Your choice of method should be based upon whether the patterns are *static* (constant in time) or *dynamic* (changes in time), the nature of the trend and seasonal components, and how far ahead that you wish to forecast. These methods are generally easy and quick to apply.

ARIMA modeling also makes use of patterns in the data, but these patterns may not be easily visible in a plot of the data. Instead, ARIMA modeling uses differencing and the autocorrelation and partial autocorrelation functions to help identify an acceptable model. ARIMA stands for Autoregressive Integrated Moving Average, which represent the filtering steps taken in constructing the ARIMA model until only random noise remains. While ARIMA models are valuable for modeling temporal processes and are also used for forecasting, fitting a model is an iterative approach that may not lend itself to application speed and volume.

Simple forecasting and smoothing methods

The simple forecasting and smoothing methods model components in a series that are usually easy to see in a time series plot of the data. This approach decomposes the data into its component parts, and then extends the estimates of the components into the future to provide forecasts. You can choose from the static methods of trend analysis and decomposition, or the dynamic methods of moving average, single and double exponential smoothing, and Winters' method. *Static* methods have patterns that do not change over time; *dynamic* methods have patterns that do change over time and estimates are updated using neighboring values.


You may use two methods in combination. That is, you may choose a static method to model one component and a dynamic method to model another component. For example, you may fit a static trend using trend analysis and dynamically model the seasonal component in the residuals using Winters' method. Or, you may fit a static seasonal model using decomposition and dynamically model the trend component in the residuals using double exponential smoothing. You might also apply a trend analysis and decomposition together so that you can use the wider selection of trend models offered by trend analysis (see Example of trend analysis and Example of decomposition). A disadvantage of combining methods is that the confidence intervals for forecasts are not valid.

For each of the methods, the following table provides a summary and a graph of fits and forecasts of typical data.

Command	Forecast	Example
---------	----------	---------

Trend Analysis		
----------------	--	--

Fits a general trend model to time series data. Choose among the linear, quadratic, exponential growth or decay, and S-curve models. Use this procedure to fit trend when there is no seasonal component in your series.	Length: long Profile: extension of trend line	
--	--	--




Time Series

Decomposition

Separates the times series into linear trend and seasonal components, as well as error. Choose whether the seasonal component is additive or multiplicative with the trend. Use this procedure to forecast when there is a seasonal component in your series or if you simply want to examine the nature of the component parts


Length: long
Profile: trend with seasonal pattern



Moving Average

Smooths your data by averaging consecutive observations in a series. This procedure can be a likely choice when your data do not have a trend or seasonal component. There are ways, however, to use moving averages when your data possess trend and/or seasonality.


Length: short
Profile: flat line



Single Exp Smoothing

Smooths your data using the optimal one-step ahead ARIMA (0,1,1) forecasting formula. This procedure works best without a trend or seasonal component. The single dynamic component in a moving average model is the level.


Length: short
Profile: flat line



Double Exp Smoothing


Smooths your data using the optimal one-step-ahead ARIMA (0,2,2) forecasting formula. This procedure can work well when trend is present but it can also serve as a general smoothing method. Double Exponential Smoothing calculates dynamic estimates for two components: level and trend.

Length: short
Profile: straight line with slope equal to last trend estimate

**Winters' Method**

Smooths your data by Holt-Winters exponential smoothing. Use this procedure when trend and seasonality are present, with these two components being either additive or multiplicative. Winters' Method calculates dynamic estimates for three components: level, trend, and seasonal.

Length: short to medium
Profile: trend with seasonal pattern

**Correlation analysis and ARIMA modeling**

Examining correlation patterns within a time series or between two time series is an important step in many statistical analyses. The correlation analysis tools of differencing, autocorrelation, and partial autocorrelation are often used in ARIMA modeling to help identify an appropriate model.

ARIMA modeling can be used to model many different time series, with or without trend or seasonal components, and to provide forecasts. The forecast profile depends upon the model that is fit. The advantage of ARIMA modeling compared to the simple forecasting and smoothing methods is that it is more flexible in fitting the data. However, identifying and fitting a model may be time-consuming, and ARIMA modeling is not easily automated.

- Differences computes and stores the differences between data values of a time series. If you wish to fit an ARIMA model but there is trend or seasonality present in your data, differencing data is a common step in assessing likely ARIMA models. Differencing is used to simplify the correlation structure and to reveal any underlying pattern.
- Lag computes and stores the lags of a time series. When you lag a time series, Minitab moves the original values down the column, and inserts missing values at the top of the column. The number of missing values inserted depends on the length of the lag.
- Autocorrelation computes and plots the autocorrelations of a time series. Autocorrelation is the correlation between observations of a time series separated by k time units. The plot of autocorrelations is called the autocorrelation function or acf. View the acf to guide your choice of terms to include in an ARIMA model.
- Partial Autocorrelation computes and plots the partial autocorrelations of a time series. Partial autocorrelations, like autocorrelations, are correlations between sets of ordered data pairs of a time series. As with partial correlations in the regression case, partial autocorrelations measure the strength of relationship with other terms being accounted for. The partial autocorrelation at a lag of k is the correlation between residuals at time t from an autoregressive model and observations at lag k with terms for all intervening lags present in the autoregressive model. The plot of partial autocorrelations is called the partial autocorrelation function or pacf. View the pacf to guide your choice of terms to include in an ARIMA model.
- Cross Correlation computes and graphs correlations between two time series.
- ARIMA fits a Box-Jenkins ARIMA model to a time series. ARIMA stands for Autoregressive Integrated Moving Average. The terms in the name—Autoregressive, Integrated, and Moving Average—represent filtering steps taken in

constructing the ARIMA model until only random noise remains. Use ARIMA to model time series behavior and to generate forecasts.

Time Series Plot

Time Series Plots


Graph > Time Series Plot

Stat > Time Series > Time Series Plot

Use to evaluate patterns in data over time. Minitab can generate calendar values, clock values, or index values for the time scale, or you can use your own column of stamp values. See Time Overview.

Minitab plots the data in worksheet order, in equally-spaced time intervals. If your data are not entered in chronological order or were not collected at regular intervals, you may want to use Graph > Scatterplot.

Dialog box items



Data - Time Series Plot

You need one or more columns of time series data in chronological order. You may also use up to three columns of categorical data for grouping. For multiple overlaid time series with groups, all columns (series and grouping variables) must be of equal length.

Tip You can use Data > Sort to reorder data.

To use a stamp for x-axis labels, you need one to three columns of numeric, text, or date/time columns. See Time overview.

If data in your series are missing, those worksheet rows must contain missing values in order to maintain the proper spacing in the time scale.

More For other data options (subsetting and handling missing data), see Data Options Overview.

Time Series Plot - Simple

Graph > Time Series Plot > choose *Simple* > OK

Stat > Time Series > Time Series Plot > choose *Simple* > OK

Use to display a simple time series plot.

Dialog box items

Series: Enter one or more columns of time series data. Minitab displays a separate graph for each column. (To panel all graphs on the same page, use Multiple Variables.)

Tip If your time series has many values, you may want to split data across panels.

To display a simple time series plot

- 1 Choose **Graph > Time Series Plot** or **Stat > Time Series > Time Series Plot**.
- 2 Choose **Simple**, then click **OK**.
- 3 In **Series**, enter one or more columns of time series data. Minitab displays a separate graph for each column.


- If you like, use any dialog box options, then click **OK**.

Example of a simple time series plot

You are a sales manager and you want to view your company's quarterly sales for 2001 to 2003. Create a time series plot.

- Open the worksheet NEWMARKET.MTW.
- Choose **Graph > Time Series Plot** or **Stat > Time Series > Time Series Plot**.
- Choose **Simple**, then click **OK**.
- In **Series**, enter *SalesB*.
- Click **Time/Scale**.
- Under **Time Scale**, choose **Calendar**. Then, choose **Quarter Year**.
- For start values, under **Quarter**, type *1*. Under **Year**, type *2000*.
- Click **OK** in each dialog box.

Graph window output



Interpreting the results

Overall sales increased over the three years. Sales may be cyclical, with lower sales in the first quarter of each year.

Tip To see the y- and x-values for a symbol in a series, hover your cursor over the symbol.

Time Series Plot - With Groups

Graph > Time Series Plot > choose *With Groups* > OK
Stat > Time Series > Time Series Plot > choose *With Groups* > OK

Use to display a time series plot with groups.

Dialog box items

Series: Enter one or more columns of time series data. Minitab displays a separate graph for each column. (To panel all graphs on the same page, use *Multiple Variables*.)

Categorical variables for grouping (1-3): Enter one to three columns of grouping (categorical) variables. Minitab displays all groups on the same graph. (To panel by one or more variables, use *By Variables*.)

Tip If your time series has many values, you may want to split data across panels.

To display a time series plot with groups

- Choose **Graph > Time Series Plot** or **Stat > Time Series > Time Series Plot**.
- Choose **With Groups**, then click **OK**.
- In **Series**, enter one or more columns of time series data. Minitab displays a separate graph for each column.

Time Series


- 4 In **Categorical variables for grouping (1-3)**, enter one to three columns of grouping (categorical) variables.
- 5 If you like, use any dialog box options, then click **OK**.

Example of a time series plot with groups

The ABC company used two advertising agencies in 2000-2001– the Alpha Advertising Agency in 2000 and the Omega Advertising Agency in 2001. You want to compare the sales data for the past two years. Create a time series plot with groups.

- 1 Open the worksheet ABCSALES.MTW.
- 2 Choose **Graph > Time Series Plot** or **Stat > Time Series > Time Series Plot**.
- 3 Choose **With Groups**, then click **OK**.
- 4 In **Series**, enter *Sales*.
- 5 In **Categorical variables for grouping (1-3)**, enter *AdAgency*.
- 5 Click **Time/Scale**.
- 6 Under **Time Scale**, choose **Calendar**, then choose **Month Year**.
- 7 For start values, under **Month**, type *1*. Under **Year**, type *2000*.
- 8 Click **OK** in each dialog box.

Graph window output



Interpreting the results

Tip To see the y- and x-values for a symbol in a series, hover your cursor over the symbol.

Sales increased both years. Sales for the Alpha ad agency increased 161, from 210 to 371. Subsequently, sales for the Omega ad agency rose somewhat less dramatically from 368 to 450, an increase of 82. However, the effects of other factors, such as amount of advertising dollars spent and the economic conditions, are unknown.

Time Series Plot - Multiple

Graph > Time Series Plot > choose Multiple > OK
Stat > Time Series > Time Series Plot > choose Multiple > OK

Use to overlay multiple time series plots on the same graph.

Dialog box items

Series: Enter one or more columns of time series data. Minitab displays all columns on the same graph.

Tip If your time series has many values, you may want to split data across panels.

To display multiple overlaid time series plots

- 1 Choose **Graph > Time Series Plot** or **Stat > Time Series > Time Series Plot**.


- 2 Choose **Multiple**, then click **OK**.
- 3 In **Series**, enter one or more columns of time series data. Minitab displays all columns on the same graph.
- 4 If you like, use any dialog box options, then click **OK**.

Example of multiple overlaid time series plots

You own stocks in two companies (ABC and XYZ) and you want to compare their monthly performance for two years. Create an overlaid time series plot of share prices for ABC and XYZ.

- 1 Open the worksheet SHAREPRICE.MTW.
- 2 Choose **Graph > Time Series Plot** or **Stat > Time Series > Time Series Plot**.
- 3 Choose **Multiple**, then click **OK**.
- 4 In **Series**, enter *ABC XYZ*.
- 5 Click **Time/Scale**.
- 6 Under **Time Scale**, choose **Calendar**, then choose **Month Year**.
- 7 For start values, under **Month**, enter *1*. Under **Year**, enter *2001*.
- 8 Click **OK** in each dialog box.

Graph window output



Interpreting the results

Tip To see the y- and x-values for a symbol in a series, hover your cursor over the symbol.

The solid line for ABC share price shows a slow increase over the two-year period. The dashed line for XYZ share price also shows an overall increase for the two years, but it fluctuates more than that of ABC. The XYZ share price starts lower than ABC (30 vs. 36.25 for ABC). By the end of 2002, the XYZ price surpasses the ABC price by 14.75 (44.50 to 60.25).

Time Series Plot - Multiple With Groups

Graph > Time Series Plot > choose *Multiple with Groups* > OK
Stat > Time Series > Time Series Plot > choose *Multiple with Groups* > OK

Use to overlay multiple time series plots with groups on the same graph.

Dialog box items

Series: Enter multiple, equal-length columns of time series data. Minitab displays all columns on the same graph.

Categorical variables for grouping (1-3): Enter one to three columns of grouping (categorical) variables. Minitab displays all groups on the same graph. (To panel by one or more variables, use *By Variables*.) Columns must be the same length as the series columns.

Tip If your time series has many values, you may want to split data across panels.

To display multiple overlaid time series plots with groups

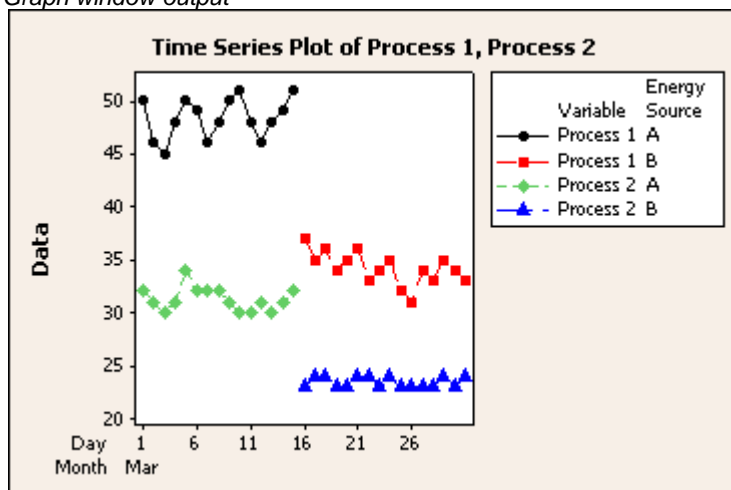
- 1 Choose **Graph > Time Series Plot** or **Stat > Time Series > Time Series Plot**.
- 2 Choose **Multiple with Groups**, then click **OK**.
- 3 In **Series**, enter multiple columns of time series data. Minitab displays all columns on the same graph.
- 4 In **Categorical variables for grouping (1-3)**, enter one to three columns of grouping (categorical) variables.
- 5 If you like, use any dialog box options, then click **OK**.

Example of multiple overlaid time series plots with groups

Your company uses two different processes to manufacture plastic pellets. Energy is a major cost, and you want to try a new source of energy. You use energy source A (your old source) for the first half of the month, and energy source B (your new source) for the second half. Create a time series plot to illustrate the energy costs of two processes from the two sources.

- 1 Open the worksheet ENERGYCOST.MTW.
- 2 Choose **Graph > Time Series Plot** or **Stat > Time Series > Time Series Plot**.
- 3 Choose **Multiple with Groups**, then click **OK**.
- 4 In **Series**, enter 'Process 1' 'Process 2'.
- 5 In **Categorical variables for grouping (1-3)**, enter 'Energy Source'.
- 6 Click **Time/Scale**.
- 6 Under **Time Scale**, choose **Calendar**, then choose **Day Month**.
- 7 For start values, under **Day**, type 1. Under **Month**, type 3.
- 8 Click **OK** in each dialog box.

Graph window output



Interpreting the results

Energy costs for Process 1 are generally greater than those for Process 2. In addition, costs for both processes were less using source B.

Therefore, using Process 2 and energy source B appears to be more cost effective than using Process 1 and energy source A.

Tip To see the y- and x-values for a symbol in a series, hover your cursor over the symbol.

Time/Scale - Time - Time Series Plot

Time Series Plot > Time/Scale > Time

Use to specify a time scale for the x-axis. After you create a graph you can:

- Edit the time units of the scale.
- Edit the tick positions and range of the scale.

- Edit the tick labels.
- Edit the attributes, font, alignment, and display of the scale.

Dialog box items

Time Scale

Index: Choose to label the x-axis with integers.

Calendar: Choose to label the x-axis with a calendar scale, then choose the calendar units.

Clock: Choose to label the x-axis with a clock scale, then choose the clock units.

Stamp: Choose to label the x-axis with values from one or more stamp columns.

Stamp columns (1-3): Enter up to three columns containing stamp values.

Start Values (not available with all graphs)

One set for all variables: Choose to use the same start values for each variable. Type the integer start values.

One set for each variable: Choose to use a different set of start values for each variable. For each variable, type the integer start values.

Increment: Type the interval between successive worksheet observations on the innermost time scale.

To customize the time scale

- 1 In the graph dialog box, choose **Time/Scale**.
- 2 Click the **Time** tab.
- 3 Under **Time Scale**, choose one of the following:
 - **Index** to label the x-axis with integers.
 - **Calendar** to label the x-axis with a calendar scale.
 - Choose the calendar units.
 - **Clock** to label the x-axis with a clock scale.
 - Choose the clock units.
 - **Stamp** to label the x-axis with values from one or more stamp columns.
 - In **Stamp columns (1-3)**, enter up to three columns containing scale values.
- 4 Under **Start Values**, choose one of the following:
 - **One set for all variables** to use the same start values for each variable.
 - Type the start value for each component of the time scale.
 - **One set for each variable** to use a different set of start values for each variable, if there are overlaid variables.
 - For each variable, type the start values for each component of the time scale.
- 5 In **Increment**, type a value to increment the innermost time scale.
- 6 Click **OK**.

Example of using a stamp for the x-axis


You need to plot your sales using the date format qQyy for the time scale; for example, 4Q00 stands for the fourth quarter of 2000. You enter a column in Minitab called Date and use this column for a stamp on the time series plot.

Note You can also use a stamp with area graphs or control charts.

- 1 Open the worksheet NEWMARKET.MTW.
- 2 Choose **Graph > Time Series Plot**.
- 3 Choose **Simple**, then click **OK**.
- 4 In **Series**, enter *SalesB*.
- 5 Click **Time/Scale**.
- 6 Under **Time Scale**, choose **Stamp**.
- 7 In **Stamp columns (1-3)**, enter *Date*. Click **OK** in each dialog box.

Time Series

Graph window output



Interpreting the results

The tick marks on the x-axis are labeled with values from the stamp column.

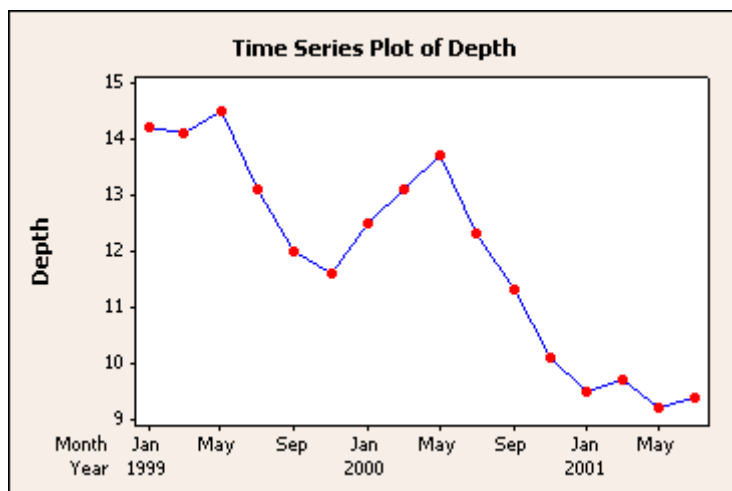
Example of using custom start times

Suppose you want to see how the depth of a reservoir fluctuates over time. You collect measurements every other month for 16 months, starting in January of 1999. Create a time series plot with custom start times.

Note You can also use custom start times with area graphs.

- 1 Open the worksheet EXH_GRP.MTW.
- 2 Choose **Graph > Time Series Plot**.
- 3 Choose **Simple**, then click **OK**.
- 4 In **Series**, enter *Depth*.
- 5 Click **Time/Scale**.
- 6 Under **Time Scale**, choose **Calendar**, then choose **Month Year**.
- 7 Under **Start Values**, choose **One set for all variables**.
- 8 Under **Month**, type 1. Under **Year**, type 1999.
- 9 In **Increment**, type 2. Click **OK** in each dialog box.

Graph window output



More If a graph has too many tick labels, you can hide some. See [Axes and Ticks](#).

Interpreting the results

The time series plot shows data points for every other month, starting with January 1999.

Trend Analysis

Trend Analysis

Stat > Time Series > Trend Analysis

Trend analysis fits a general trend model to time series data and provides forecasts. Choose among the linear, quadratic, exponential growth or decay, and S-curve models. Use this procedure to fit trend when there is no seasonal component to your series.

Dialog box items

Variable: Enter the column containing the time series.

Model Type: Select the model that you want. Use care when interpreting the coefficients from the different models, as they have different meanings. See [4] for details.

Linear: Click to fit the linear trend model.

Quadratic: Click to fit the quadratic trend model.

Exponential growth: Click to fit the exponential growth trend model.

S-Curve (Pearl-Reed logistic): Click to fit the Pearl-Reed logistic S-curve trend model. You cannot have missing data when fitting the S-curve model.

Generate forecasts: Check to generate forecasts.

Number of forecasts: Enter an integer to indicate how many forecasts that you want.

Starting from origin: Enter a positive integer to specify a starting point for the forecasts. For example, if you specify 4 forecasts and 48 for the origin, Minitab computes forecasts for periods 49, 50, 51, and 52. If you leave this space blank, Minitab generates forecasts from the end of the data. Minitab uses data up to the origin for fitting the trend model used to generate forecasts.


Data – Trend Analysis

The time series must be in one numeric column. If you choose the S-curve trend model, you must delete missing data from the worksheet before performing the trend analysis. Minitab automatically omits missing values from the calculations when you use one of the other three trend models.

To perform a Trend Analysis

- 1 Choose **Stat > Time Series > Trend Analysis**.
- 2 In **Variable**, enter the column containing the series.
- 3 If you like, use any dialog box options, then click **OK**.

Trend Analysis : When to Use



Use for:

- Data with constant trend, and
- Data with no seasonal pattern
- Long range forecasting

Forecast profile:

- Continuation of trend line fit to data

ARIMA equivalent: none

Measures of Accuracy

Minitab computes three measures of accuracy of the fitted model: MAPE, MAD, and MSD for each of the simple forecasting and smoothing methods. For all three measures, the smaller the value, the better the fit of the model. Use these statistics to compare the fits of the different methods.

MAPE, or Mean Absolute Percentage Error, measures the accuracy of fitted time series values. It expresses accuracy as a percentage.

$$\text{MAPE} = \frac{\sum |y_t - \hat{y}_t| / y_t}{n} \times 100 \quad (y_t \neq 0)$$

where y_t equals the actual value, \hat{y}_t equals the fitted value, and n equals the number of observations.

MAD, which stands for Mean Absolute Deviation, measures the accuracy of fitted time series values. It expresses accuracy in the same units as the data, which helps conceptualize the amount of error.

$$\text{MAD} = \frac{\sum_{t=1}^n |y_t - \hat{y}_t|}{n}$$

where y_t equals the actual value, \hat{y}_t equals the fitted value, and n equals the number of observations.

MSD stands for Mean Squared Deviation. MSD is always computed using the same denominator, n , regardless of the model, so you can compare MSD values across models. MSD is a more sensitive measure of an unusually large forecast error than MAD.

$$\text{MSD} = \frac{\sum_{t=1}^n |y_t - \hat{y}_t|^2}{n}$$

where y_t equals the actual value, \hat{y}_t equals the forecast value, and n equals the number of forecasts.

Forecasts – Trend Analysis

Forecasts are extrapolations of the trend model fits. Data prior to the forecast origin are used to fit the trend.

Trend models

There are four different trend models you can choose from: linear (default), quadratic, exponential growth curve, or S-curve (Pearl-Reed logistic). Use care when interpreting the coefficients from the different models, as they have different meanings. See [4] for details.

Trend analysis by default uses the *linear trend* model:

$$Y_t = \beta_0 + (\beta_1 * t) + e_t$$

In this model, β_1 represents the average change from one period to the next.

The *quadratic trend model* which can account for simple curvature in the data, is:

$$Y_t = \beta_0 + \beta_1 * t + (\beta_2 * t^2) + e_t$$

The *exponential growth trend model* accounts for exponential growth or decay. For example, a savings account might exhibit exponential growth. The model is:

$$Y_t = \beta_0 * \beta_1^t * e_t$$

The *S-curve model* fits the Pearl-Reed logistic trend model. This accounts for the case where the series follows an S-shaped curve. The model is:

$$Y_t = (10^n) / (\beta_0 + \beta_1 \beta_2^t)$$

Time

Stat > Time Series > menu command > Time

Specify the time scale.

Dialog box items

Time Scale: Change the x-axis of the scale you specify.

Index: Choose to number the x-axis with a single scale from 1 to n by ones (where n is the number of observations in the column containing the time series).

Calendar: Choose to label the x-axis with different time unit scales. For example, if you choose Month Quarter Year, Minitab assumes that your data are in intervals of 1 month and generates 3 scales on the x-axis: 1 for months, 1 for quarters, and 1 for years. At each 4th tick on the month scale, Minitab generates a tick on the quarter scale. At each 4th tick on the quarter scale, Minitab generates a tick on the Year scale.

Clock: Choose to label the x-axis with different time unit scales. For example, if you choose Day Hour Minute, Minitab assumes that your data are in minute intervals and generates 3 scales on the x-axis: 1 for days, 1 for hours, and 1 for minutes. At each 60th tick on the minute scale, Minitab generates a tick on the hour scale. At each 24th tick on the hour scale, Minitab generates a tick on the Day scale.

Start value(s): Enter the integer start values.

Increment: Enter a value to increment the time scale.

Stamp: Choose to add a date/time stamp on the x-axis. Enter a date/time column in the text box. The values are stamped on the plot in the same format as they appear in the column.

Trend Analysis – Options

Stat > Time Series > Trend Analysis > Options

Specify a customized title and the trend analysis fit to be a weighted average of a fit of prior data and the current data.

Dialog box items

Title: To replace the default title with your own custom title, type the desired text in this box.

Prior parameter values: Specify the parameter values of the trend analysis fit to the prior data for the weighted average fit. See Weighted average trend analysis.

Weights for blending priors with new estimates: Enter weights of coefficients of current data for the weighted average fit. Enter numbers between 0 and 1 for weights of each trend model parameter estimate. See Weighted average trend analysis.

To perform a weighted average trend analysis

- 1 Choose **Stat > Time Series > Trend Analysis > Options**.
- 2 Enter the coefficient estimates from the prior trend analysis in the order in which they are given in the Session window or the graph.
- 3 Optionally enter weights between 0 and 1 for each new coefficient, in the same order as for coefficients. Default weights of 0.2 will be used for each coefficient if you don't enter any. If you do enter weights, the number that you enter must be equal to the number of coefficients.
- 4 If you like, use any dialog box options, then click **OK**.

Minitab generates a time series plot of the data, plus a second time series plot that shows trend lines for three models. The Session window displays the coefficients and accuracy measures for all three models.

Weighted average trend analysis

You can perform a weighted average trend analysis to incorporate knowledge learned from fitting the same trend model to prior data in order to obtain an "improved" fit to the present data. The smoothed trend line combines prior and new information in much the same way that exponential smoothing works. In a sense, this smoothing of the coefficients filters out some of the noise from the model parameters estimated in successive cycles.

If you supply coefficients from a prior trend analysis fit, Minitab performs a weighted trend analysis. If the weight for a particular coefficient is a, Minitab estimates the new coefficient by:

$\alpha p_1 + (1 - \alpha) p_2$, where p_1 is the coefficient estimated from the current data and p_2 is the prior coefficient.

Trend Analysis – Storage

Stat > Time Series > Trend Analysis > Storage

Stores fits, residuals, and forecasts in the worksheet.

Dialog box items

Fits (trend line): Check to store the fitted values. These are the values used to plot the trend line. You should store the fitted values if you want to generate diagnostic residual plots.

Residuals (detrended data): Check to store the residuals. If you store the residuals you can generate diagnostic plots, using Autocorrelation. For the linear, quadratic, S-curve, and exponential growth models, the detrended data equal the original data minus the fits.

Forecasts: Check to store the forecasts. This option is available only if you generated forecasts in the main Trend Analysis dialog box.

Trend Analysis – Graphs

Stat > Time Series > Trend Analysis > Graphs

Displays time series plot and residual plots for diagnosis of model fit.

Dialog box items

Time series plot (including optional forecasts)

Display plot: Choose to display the time series plot..

Do not display plot: Choose to suppress the time series plot.

Residual Plots

Individual plots: Choose to display one or more plots.

Histogram of residuals: Check to display a histogram of the residuals.

Normal plot of residuals: Check to display a normal probability plot of the residuals.

Residuals versus fits: Check to plot the residuals versus the fitted values.

Residuals versus order: Check to plot the residuals versus the order of the data. The row number for each data point is shown on the x-axis—for example, 1 2 3 4... n.

Four in one: Choose to display a layout of a histogram of residuals, a normal plot of residuals, a plot of residuals versus fits, and a plot of residuals versus order.

Residuals versus the variables: Enter one or more columns containing the variables against which you want to plot the residuals. Minitab displays a separate graph for each column.

Trend Analysis – Results

Stat > Time Series > Trend Analysis > Results

Controls Session window output.

Dialog box items

Control the Display of Results

Display nothing: Choose to suppress output.

Summary table: Choose to display the default output – fitted trend equation and accuracy measures.

Summary table and results table: Choose to display the default output plus a table of the original series, the trend and the detrended data.

Example of Trend Analysis

You collect employment data in a trade business over 60 months and wish to predict employment for the next 12 months. Because there is an overall curvilinear pattern to the data, you use trend analysis and fit a quadratic trend model. Because there is also a seasonal component, you save the fits and residuals to perform decomposition of the residuals (see Example of Decomposition).

- 1 Open the worksheet EMPLOY.MTW.
- 2 Choose **Stat > Time Series > Trend Analysis**.

- 3 In **Variable**, enter *Trade*.
- 4 Under **Model Type**, choose **Quadratic**.
- 5 Check **Generate forecasts** and enter 12 in **Number of forecasts**.
- 6 Click **Storage**.
- 7 Check **Fits (Trend Line)**, **Residuals (detrended data)**, and **Forecasts**. Click **OK** in each dialog box.

Session Window Output

Trend Analysis for Trade

Data Trade
 Length 60
 NMissing 0

Fitted Trend Equation

$$Y_t = 320.762 + 0.509373*t + 0.0107456*t**2$$

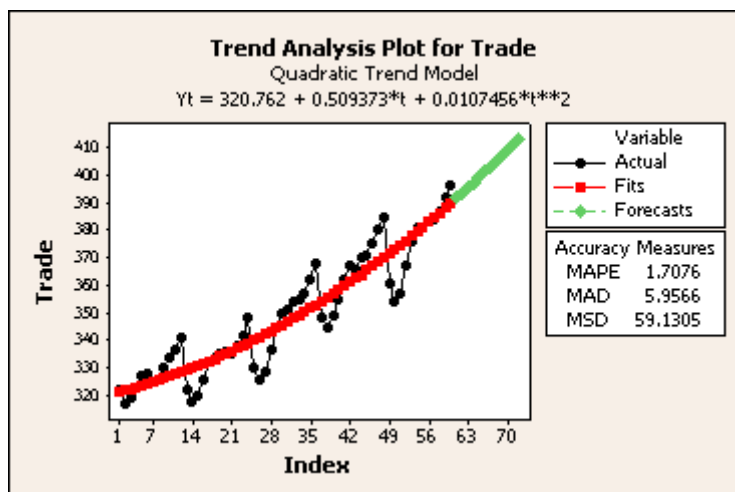
Accuracy Measures

MAPE 1.7076
 MAD 5.9566
 MSD 59.1305

Forecasts

Period	Forecast
61	391.818
62	393.649
63	395.502
64	397.376
65	399.271
66	401.188
67	403.127
68	405.087
69	407.068
70	409.071
71	411.096
72	413.142

Graph window output



Interpreting the results

The trend plot that shows the original data, the fitted trend line, and forecasts. The Session window output also displays the fitted trend equation and three measures to help you determine the accuracy of the fitted values: MAPE, MAD, and MSD. The trade employment data show a general upward trend, though with an evident seasonal component. The trend model appears to fit well to the overall trend, but the seasonal pattern is not well fit. To better fit these data, you also use decomposition on the stored residuals and add the trend analysis and decomposition fits and forecasts (see Example of decomposition).

Decomposition

Decomposition

Stat > Time Series > Decomposition

You can use decomposition to separate the time series into linear trend and seasonal components, as well as error, and provide forecasts. You can choose whether the seasonal component is additive or multiplicative with the trend. Use this procedure when you wish to forecast and there is a seasonal component to your series, or if you simply want to examine the nature of the component parts. See [6] for a discussion of decomposition methods.

Dialog box items

Variable: Enter the column containing the time series.

Seasonal Length: Enter a positive integer greater than or equal to 2. This is the length of the seasonal component. For example, if you have monthly data, you might use a seasonal length of 12.

Model Type:

Multiplicative: Choose to use the multiplicative model.

Additive: Choose to use the additive model.

Model Components:

Trend plus seasonal: Choose to include the trend component in the decomposition.

Seasonal only: Choose to omit the trend component from the decomposition. You might want to do this if you have already detrended your data with Trend Analysis.

Caution If the data contain a trend component but you omit it from the decomposition, the estimates of the seasonal indices may be affected.

Generate forecasts: Check if you want to generate forecasts.

Number of forecasts: Enter an integer to indicate how the number of forecasts.

Starting from origin: Enter a positive integer to specify a starting point for the forecasts. For example, if you specify 4 forecasts and 48 for the origin, Minitab computes forecasts for periods 49, 50, 51, and 52. If you leave this space blank, Minitab generates forecasts from the end of the data.

Data – Decomposition

The time series must be in one numeric column. Minitab automatically omits missing data from the calculations.

The data that you enter depends upon how you use this procedure. Usually, decomposition is performed in one step by simply entering the time series. Alternatively, you can perform a decomposition of the trend model residuals. This process may improve the fit of the model by combining the information from the trend analysis and the decomposition. See Decomposition of trend model residuals .

To perform a Decomposition

- 1 Choose **Stat > Time Series > Decomposition**.
- 2 In **Variable**, enter the column containing the series.
- 3 In **Seasonal length**, enter the seasonal length or period.
- 4 If you like, use any dialog box options, then click **OK**.


To combine trend analysis and decomposition

- 1 Perform a Trend Analysis and store the fits, residuals, and forecasts (see Example of a trend analysis).
- 2 Choose **Stat > Time Series > Decomposition**.
- 3 In **Variable**, enter the column containing trend analysis residuals.

- 4 Under **Model Type**, choose **Additive**.
- 5 Under **Model Components**, choose **Seasonal only**.
- 6 Click **Storage** and check **Fits**. Click **OK** in each dialog box.
- 7 Next, you need to calculate the fits from the combined procedure. If you want these components to be additive, add the respective fits together.

Note The MAPE, MAD, MSD accuracy measures from decomposition used in this manner are not comparable to these statistics calculated from other procedures, but you can calculate the comparable values fairly easily. We demonstrate this with MSD in the decomposition example.

Decomposition, Additive Model : When to Use



Use for:


- Data with either no trend or constant trend, and
- Data with constant seasonal pattern
- Size of seasonal pattern not proportional to data
- Long range forecasting

Forecast profile:

- Straight line with seasonal pattern added

ARIMA equivalent: none

Decomposition, Multiplicative Model : When to Use



Use for:

- Data with either no trend or constant trend, and
- Data with constant seasonal pattern

Time Series

- Size of seasonal pattern proportional to data
- Long range forecasting

Forecast profile:

- Straight line multiplied by seasonal pattern

ARIMA equivalent: none

How Minitab Does Decomposition

Decomposition model

By default, Minitab uses a *multiplicative model*. Use the multiplicative model when the size of the seasonal pattern in the data depends on the level of the data. This model assumes that as the data increase, so does the seasonal pattern. Most time series exhibit such a pattern. The multiplicative model is

$$Y_t = \text{Trend} * \text{Seasonal} * \text{Error}$$

where Y_t is the observation at time t .

The *additive model* is:

$$Y_t = \text{Trend} + \text{Seasonal} + \text{Error}$$

where Y_t is the observation at time t .

You can also omit the trend component from the decomposition.

Method

Decomposition involves the following steps:

- 1 Minitab smoothes the data using a centered moving average with a length equal to the length of the seasonal cycle. When the seasonal cycle length is an even number, a two-step moving average is required to synchronize the moving average correctly.
- 2 Minitab either divides the moving average into (multiplicative model) or subtracts it from (additive model) the data to obtain what are often referred to as raw seasonal values.
- 3 For corresponding time periods in the seasonal cycles, Minitab determines the median of the raw seasonal values. For example, if you have 60 consecutive months of data (5 years), Minitab determines the median of the 5 raw seasonal values corresponding to January, to February, and so on.
- 4 Minitab adjusts the medians of the raw seasonal values so that their average is one (multiplicative model) or zero (additive model). These adjusted medians constitute the seasonal indices.
- 5 Minitab uses the seasonal indices to seasonally adjust the data.
- 6 Minitab fits a trend line to the seasonally adjusted data using least squares regression.

The data can be detrended by either dividing the data by the trend component (multiplicative model) or subtracting the trend component from the data (additive model).

Forecasts – Decomposition

Decomposition calculates the forecast as the linear regression line multiplied by (multiplicative model) or added to (additive model) the seasonal indices. Data prior to the forecast origin are used for the decomposition.

Time

Stat > Time Series > menu command > Time

Specify the time scale.

Dialog box items

Time Scale: Change the x-axis of the scale you specify.

Index: Choose to number the x-axis with a single scale from 1 to n by ones (where n is the number of observations in the column containing the time series).

Calendar: Choose to label the x-axis with different time unit scales. For example, if you choose Month Quarter Year, Minitab assumes that your data are in intervals of 1 month and generates 3 scales on the x-axis: 1 for months, 1 for quarters, and 1 for years. At each 4th tick on the month scale, Minitab generates a tick on the quarter scale. At each 4th tick on the quarter scale, Minitab generates a tick on the Year scale.

Clock: Choose to label the x-axis with different time unit scales. For example, if you choose Day Hour Minute, Minitab assumes that your data are in minute intervals and generates 3 scales on the x-axis: 1 for days, 1 for hours, and 1 for

minutes. At each 60th tick on the minute scale, Minitab generates a tick on the hour scale. At each 24th tick on the hour scale, Minitab generates a tick on the Day scale.

Start value(s): Enter the integer start values.

Increment: Enter a value to increment the time scale.

Stamp: Choose to add a date/time stamp on the x-axis. Enter a date/time column in the text box. The values are stamped on the plot in the same format as they appear in the column.

Decomposition – Options

Stat > Time Series > Decomposition > Options

Specify a customized title and the first observation in the seasonal period.

Dialog box items

Title: To replace the default title with your own custom title, type the desired text in this box. This title will be used for all three sets of plots.

First obs. is in seasonal period: Enter a number to specify a starting value. For example, if you have monthly data and the first observation is in June, then enter 6 to set the seasonal period correctly. By default, the starting value is 1 because Minitab assumes that the first data value in the series corresponds to the first seasonal period.

Decomposition – Storage

Stat > Time Series > Decomposition > Storage

Stores various values in the worksheet.

Dialog box items

Trend line: Check to store the trend component data. This is the trend component; it does not contain the error or seasonal component.

Detrended data: Check to store the detrended data (has seasonal and error components, no trend)

Seasonals: Check to store the seasonal component data (has seasonal component, no error or trend).

Seasonally adjusted data: Check to store the data which has had the seasonal component removed (has trend and error components, no seasonal).

Fits: Check to store the fitted values. These are the combination of trend and seasonal components, with no error component.

Residuals: Check to store the residuals. If you store the residuals you can generate diagnostic plots using Autocorrelation.

Forecasts: Check to store the forecasts. This option is available only if you generated forecasts in the main Decomposition dialog box.

Decomposition – Graphs

Stat > Time Series > Decomposition > Graphs

Displays time series plot and residual plots for diagnosis of model fit.

Dialog box items

Time series plot (including optional forecasts)

Display plot: Choose to display the time series plot.

Do not display plot: Choose to suppress the time series plot.

Residual Plots

Individual plots: Choose to display one or more plots.

Histogram of residuals: Check to display a histogram of the residuals.

Normal plot of residuals: Check to display a normal probability plot of the residuals.

Residuals versus fits: Check to plot the residuals versus the fitted values.

Residuals versus order: Check to plot the residuals versus the order of the data. The row number for each data point is shown on the x-axis—for example, 1 2 3 4... n.

Four in one: Choose to display a layout of a histogram of residuals, a normal plot of residuals, a plot of residuals versus fits, and a plot of residuals versus order.

Residuals versus the variables: Enter one or more columns containing the variables against which you want to plot the residuals. Minitab displays a separate graph for each column.

Decomposition – Results

Stat > Time Series > Decomposition > Results

Controls Session window output.

Dialog box items

Control the Display of Results

Display nothing: Choose to suppress output.

Summary table: Choose to display the default output – fitted trend equation, the seasonal indices, and accuracy measures.

Summary table and results table: Choose to display the default output plus a table of the original series, the trend and seasonal components, the detrended data, the seasonally adjusted data, the residuals, and the fits.

Decomposition of trend model residuals

You can use trend analysis and decomposition in combination when your data have a trend that is fit well by the quadratic, exponential growth curve, or S-curve models of trend analysis and possess seasonality that can be fit well by decomposition.

Example of Fits and Forecasts of Combined Trend Analysis and Decomposition

You collect employment data in a trade business over 60 months and wish to predict employment for the next 12 months. Because there is an overall curvilinear pattern to the data, you use trend analysis and fit a quadratic trend model (see Example of Trend). Because a seasonal component also exists, you save the fits and forecasts to perform decomposition of the residuals (see Example of Decomposition). You now want to combine the trend analysis and decomposition results.

Step 1: Calculate the fits and forecasts of the combined trend analysis and decomposition

- 1 Open the worksheet EMPLOY.MTW.
- 2 Choose **Calc > Calculator**.
- 3 In **Store result in variable**, enter *NewFits*.
- 4 In **Expression**, add the fits from trend analysis to the fits from decomposition. Click **OK**.
- 5 Choose **Calc > Calculator**.
- 6 Clear the **Expression** box by selecting the contents and pressing [Delete].
- 7 In **Store result in variable**, enter *NewFore*.
- 8 In **Expression**, add the forecasts from trend analysis to the forecasts from decomposition. Click **OK**.

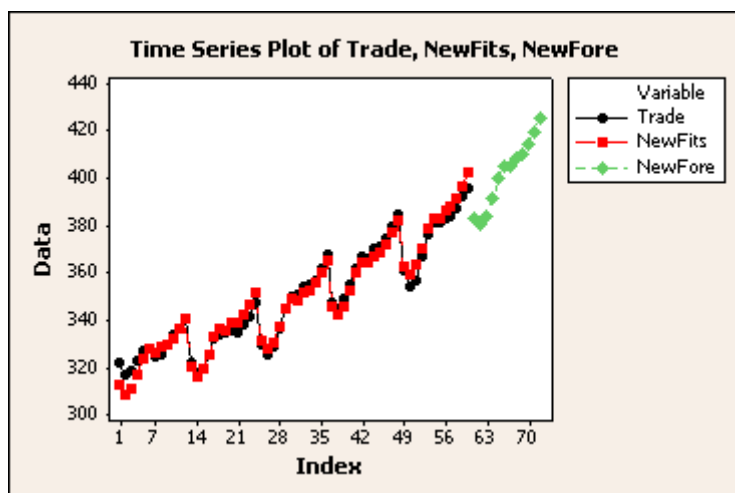
Step 2: Plot the fits and forecasts of the combined trend analysis and decomposition

- 1 Choose **Stat > Time Series > Time Series Plot > choose Multiple > OK**.
- 2 Under **Series**, enter *Trade NewFits NewFore*.
- 3 Click **Time/Scale**.
- 4 Under **Start Values**, choose **One set for each variable**, then enter 1 1 61 in rows 1–3, respectively.
- 5 Click **OK** in each dialog box.

Step 3: Calculate MSD

- 1 Choose **Calc > Calculator**.
- 2 In **Store result in variable**, enter *MSD*.
- 3 Clear the **Expression** box by selecting the contents and pressing [Delete].
- 4 In **Functions**, double-click **Sum**.
- 5 Within the parentheses in **Expression**, enter $((Trade - NewFits)**2) / 60$. Click **OK**.
- 6 Choose **Data > Display Data**.
- 7 In **Columns, constants, and matrices to display**, enter *MSD*. Click **OK**.

Graph window output



Session window output

Trend Analysis for Trade

Time Series Decomposition for RESI1

Data Display

MSD
11.8989

Interpreting the results

In the time series plot, the combined fits from the trend analysis and decomposition (*NewFits*) are close to the actual Trade graphed values (*Trade*).

You can compare fits of different models using MSD. MSD for the quadratic trend model was 59.13. Additive and multiplicative decomposition models with a linear trend (not shown) give MSD values of 20.39 and 18.54, respectively. The MSD value of 11.8989 for the combined quadratic trend and decomposition of residuals indicates a better fit using the additive trend analysis and decomposition models. You might also check the fit to these data of the multiplicative trend analysis and decomposition models.

Example of Decomposition

You wish to predict trade employment for the next 12 months using data collected over 60 months. Because the data have a trend that is fit well by trend analysis' quadratic trend model and possess a seasonal component, you use the residuals from trend analysis example (see Example of a trend analysis) to combine both trend analysis and decomposition for forecasting.

- 1 Do the trend analysis example.
- 2 Choose **Stat > Time Series > Decomposition**.
- 3 In **Variable**, enter the name of the residual column you stored in the trend analysis.
- 4 In **Seasonal length**, enter 12.
- 5 Under **Model Type**, choose **Additive**. Under **Model Components**, choose **Seasonal only**.
- 6 Check **Generate forecasts** and enter 12 in **Number of forecasts**.
- 7 Click **Storage**. Check **Forecasts** and **Fits**.
- 8 Click **OK** in each dialog box.

Time Series

Session window output

Trend Analysis for Trade

Data Trade
Length 60
NMissing 0

Fitted Trend Equation

$$Y_t = 320.762 + 0.509373*t + 0.0107456*t**2$$

Accuracy Measures

MAPE 1.7076
MAD 5.9566
MSD 59.1305

Forecasts

Period	Forecast
61	391.818
62	393.649
63	395.502
64	397.376
65	399.271
66	401.188
67	403.127
68	405.087
69	407.068
70	409.071
71	411.096
72	413.142

Time Series Decomposition for RESI1

Additive Model

Data RESI1
Length 60
NMissing 0

Seasonal Indices

Period	Index
1	-8.4826
2	-13.3368
3	-11.4410
4	-5.8160
5	0.5590
6	3.5590
7	1.7674
8	3.4757
9	3.2674
10	5.3924
11	8.4965
12	12.5590

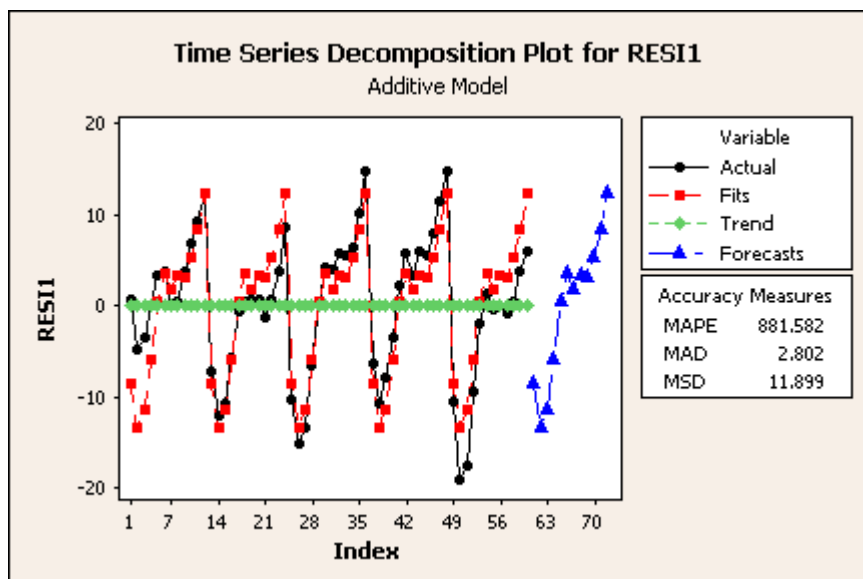
Accuracy Measures



MAPE 881.582
 MAD 2.802
 MSD 11.899

Forecasts

Period	Forecast
61	-8.4826
62	-13.3368
63	-11.4410
64	-5.8160
65	0.5590
66	3.5590
67	1.7674
68	3.4757
69	3.2674
70	5.3924
71	8.4965
72	12.5590

Graph window output





Interpreting the results

Decomposition generates three sets of plots:

- A time series plot that shows the original series with the fitted trend line, predicted values, and forecasts.
- A component analysis – in separate plots are the series, the detrended data, the seasonally adjusted data, the seasonally adjusted and detrended data (the residuals).
- A seasonal analysis – charts of seasonal indices and percent variation within each season relative to the sum of variation by season and boxplots of the data and of the residuals by seasonal period.

In addition, Minitab displays the fitted trend line, the seasonal indices, the three accuracy measures– MAPE, MAD, and MSD (see Measures of accuracy) – and forecasts in the Session window.

In the example, the first graph shows that the detrended residuals from trend analysis are fit fairly well by decomposition, except that part of the first annual cycle is underpredicted and the last annual cycle is overpredicted. This is also evident in the lower right plot of the second graph; the residuals are highest in the beginning of the series and lowest at the end.

Moving Average

Moving Average

Stat > Time Series > Moving Average

Moving Average smoothes your data by averaging consecutive observations in a series and provides short-term forecasts. This procedure can be a likely choice when your data do not have a trend or seasonal component. There are ways, however, to use moving averages when your data possess trend and/or seasonality.

Dialog box items

Variable: Enter the column containing the time series.

MA Length: Enter a positive integer to indicate desired length for the moving average. With non-seasonal time series, it is common to use short moving averages to smooth the series, although the length you select may depend on the amount of noise in the series. A longer moving average filters out more noise, but is also less sensitive to changes in the series. With seasonal series, it is common to use a moving average of length equal to the length of an annual cycle.

Center the moving averages: If you check this option, Minitab places the moving average values at the period which is in the center of the range rather than at the end of the range. This is called centering the moving average, and is done to position the moving average values at their central positions in time. Click here for more information on centering.

Generate forecasts: Check to generate forecasts. Forecasts appear in green on the time series plot with 95% prediction interval bands.

Number of forecasts: Enter an integer to indicate how many forecasts that you want.

Starting from origin: Enter a positive integer to specify a starting point for the forecasts. For example, if you specify 4 forecasts and 48 for the origin, Minitab computes forecasts for periods 49, 50, 51, and 52, based on the moving average at period 48. If you leave this space blank, Minitab generates forecasts from the end of the data.


Data – Moving Average

The time series must be in one numeric column. Minitab automatically omits missing data from the calculations.

To perform a Moving Average

- 1 Choose **Stat > Time Series > Moving Average**.
- 2 In **Variable**, enter the column containing the time series.
- 3 In **MA length**, enter a number to indicate the moving average length. See Determining the moving average length.
- 4 If you like, use any dialog box options, then click **OK**.

Moving Average : When to Use



Use for:

- Data with no trend, and
- Data with no seasonal pattern
- Short term forecasting

Forecast profile:

- Flat line

ARIMA equivalent: none

Linear Moving Averages Method

You can use this method with a time series that exhibits a trend as well as with moving average schemes involving more than two moving averages. First, compute and store the moving average of the original series. Then compute and store the moving average of the previously stored column to obtain a second moving average.

To compute and store the moving average, choose Stat > Time Series > Moving Average, complete that dialog box, choose Storage, and check the "Moving averages" box.

Determining the Moving Average Length

With non-seasonal time series, it is common to use short moving averages to smooth the series, although the length you select may depend on the amount of noise in the series. A longer moving average filters out more noise, but is also less sensitive to changes in the series. With seasonal series, it is common to use a moving average of length equal to the length of the period. For example, you might choose a moving average length of 12 for monthly data with an annual cycle.

Centering Moving Average Values

By default, moving average values are placed at the period in which they are calculated. For example, for a moving average length of 3, the first numeric moving average value is placed at period 3, the next at period 4, and so on.

When you center the moving averages, they are placed at the center of the range rather than the end of it. This is done to position the moving average values at their central positions in time.

- **If the moving average length is odd:** Suppose the moving average length is 3. In that case, Minitab places the first numeric moving average value at period 2, the next at period 3, and so on. In this case, the moving average value for the first and last periods is missing (*).
- **If the moving average length is even:** Suppose the moving average length is 4. The center of that range is 2.5, but you cannot place a moving average value at period 2.5. This is how Minitab works around the problem. Calculate the average of the first four values, call it MA1. Calculate the average of the next four values, call it MA2. Average those two numbers (MA1 and MA2), and place that value at period 3. Repeat throughout the series. In this case, the moving average values for the first two and last two periods are missing (*).

Forecasts – Moving Average

The fitted value at time t is the uncentered moving average at time $t - 1$. The forecasts are the fitted values at the forecast origin. If you forecast 10 time units ahead, the forecasted value for each time will be the fitted value at the origin. Data up to the origin are used for calculating the moving averages.

You can use the linear moving average method by performing consecutive moving averages. This is often done when there is a trend in the data. First, compute and store the moving average of the original series. Then compute and store the moving average of the previously stored column to obtain a second moving average.

In naive forecasting, the forecast for time t is the data value at time $t - 1$. Using moving average procedure with a moving average of length one gives naive forecasting.

See [1], [4], and [6] for a discussion.

Time

Stat > Time Series > *menu command* > Time

Specify the time scale.

Dialog box items

Time Scale: Change the x-axis of the scale you specify.

Index: Choose to number the x-axis with a single scale from 1 to n by ones (where n is the number of observations in the column containing the time series).

Calendar: Choose to label the x-axis with different time unit scales. For example, if you choose Month Quarter Year, Minitab assumes that your data are in intervals of 1 month and generates 3 scales on the x-axis: 1 for months, 1 for quarters, and 1 for years. At each 4th tick on the month scale, Minitab generates a tick on the quarter scale. At each 4th tick on the quarter scale, Minitab generates a tick on the Year scale.

Clock: Choose to label the x-axis with different time unit scales. For example, if you choose Day Hour Minute, Minitab assumes that your data are in minute intervals and generates 3 scales on the x-axis: 1 for days, 1 for hours, and 1 for minutes. At each 60th tick on the minute scale, Minitab generates a tick on the hour scale. At each 24th tick on the hour scale, Minitab generates a tick on the Day scale.

Start value(s): Enter the integer start values.

Increment: Enter a value to increment the time scale.

Stamp: Choose to add a date/time stamp on the x-axis. Enter a date/time column in the text box. The values are stamped on the plot in the same format as they appear in the column.

Moving Average – Options

Stat > Time Series > Moving Average > Options

Specify a customized title.

Dialog box items

Title: To replace the default title with your own custom title, type the desired text in this box.

Moving Average – Storage

Stat > Time Series > Moving Average > Storage

Stores various statistics from the moving average fit in the worksheet.

Dialog box items

Moving averages: Check to store the moving averages, which are averages of consecutive groups of data in a time series.

Fits (one-period-ahead forecasts): Check to store the fitted values. The uncentered moving average at time T is the fitted value for time T+1. You should store the fitted values if you want to generate diagnostic residual plots.

Residuals: Check to store the residuals. The residual at time T is the difference between the actual data at time T and the fitted value at time T. These residuals are used to calculate MAPE, MAD and MSD. If you store the residuals you can generate diagnostic plots using Autocorrelation.

Forecasts: Check to store the forecasts. This option is available only if you generated forecasts in the initial Moving Average dialog box.

Upper 95% Prediction Interval: Check to store the upper 95% prediction limits for the forecasts.

Lower 95% Prediction Interval: Check to store the lower 95% prediction limits for the forecasts.

Moving Average – Graphs

Stat > Time Series > Moving Average > Graphs

Displays time series plot and residual plots for diagnosis of model fit.

Dialog box items

Time series plot (including optional forecasts)

Plot predicted vs. actual: Choose to generate a time series plot which displays the data and one-period-ahead forecasts, or fitted values.

Plot smoothed vs. actual: Choose to generate a time series plot which displays the data and the smoothed, or moving average, values.

Do not display plot: Choose to suppress the time series plot.

Residual Plots

Individual plots: Choose to display one or more plots.

Histogram of residuals: Check to display a histogram of the residuals.

Normal plot of residuals: Check to display a normal probability plot of the residuals.

Residuals versus fits: Check to plot the residuals versus the fitted values.

Residuals versus order: Check to plot the residuals versus the order of the data. The row number for each data point is shown on the x-axis—for example, 1 2 3 4... n.

Four in one: Choose to display a layout of a histogram of residuals, a normal plot of residuals, a plot of residuals versus fits, and a plot of residuals versus order.

Residuals versus the variables: Enter one or more columns containing the variables against which you want to plot the residuals. Minitab displays a separate graph for each column.

Moving Average – Results

Stat > Time Series > Moving Average > Results

Controls Session window output.

Dialog box items

Control the Display of Results

Display nothing: Choose to suppress output.

Summary table: Choose to display the default non-graphical output – summary information about the data, the moving average length, and the accuracy measures.

Summary table and results table: Choose to display the default output plus a table of the original series, moving averages, the predicted values and the errors. If you generate forecasts, they are also listed, with corresponding lower and upper 95% prediction limits.

Example of Moving Average

You wish to predict employment over the next 6 months in a segment of the metals industry using data collected over 60 months. You use the moving average method as there is no well-defined trend or seasonal pattern in the data.

- 1 Open the worksheet EMPLOY.MTW.
- 2 Choose **Stat > Time Series > Moving Average**.
- 3 In **Variable**, enter *Metals*. In **MA length**, enter 3.
- 4 Check **Center the moving averages**.
- 5 Check **Generate forecasts**, and enter 6 in **Number of forecasts**. Click **OK**.

Session window output

Moving Average for Metals

```
Data      Metals
Length    60
NMissing  0
```

```
Moving Average
```

```
Length 3
```

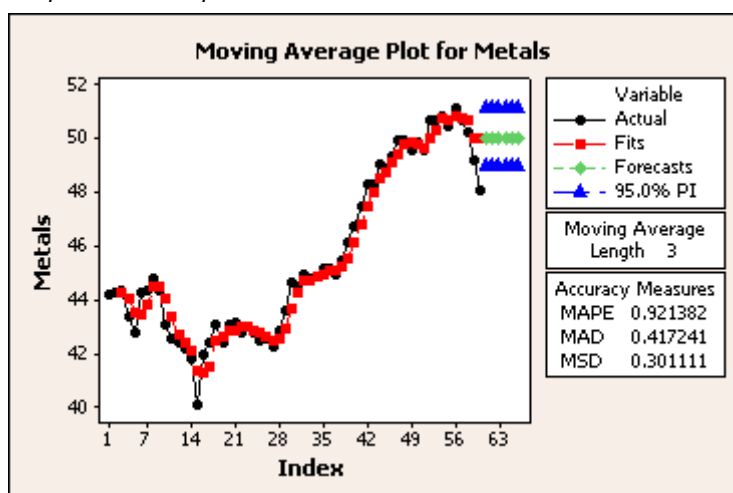
```
Accuracy Measures
```

```
MAPE  0.921382
MAD   0.417241
MSD   0.301111
```

```
Forecasts
```

Period	Forecast	Lower	Upper
61	49.2	47.4865	50.9135
62	49.2	47.4865	50.9135
63	49.2	47.4865	50.9135
64	49.2	47.4865	50.9135
65	49.2	47.4865	50.9135
66	49.2	47.4865	50.9135

Graph window output



Interpreting the results

Minitab generated the default time series plot which displays the series and fitted values (one-period-ahead forecasts), along with the six forecasts. Notice that the fitted value pattern lags behind the data pattern. This is because the fitted values are the moving averages from the previous time unit. If you wish to visually inspect how moving averages fit your data, plot the smoothed values rather than the predicted values.

To see exponential smoothing methods applied to the same data, see Example of single exponential smoothing and Example of double exponential smoothing.

In the Session window, Minitab displays three measures to help you determine the accuracy of the fitted values: MAPE, MAD, and MSD. See Measures of accuracy. Minitab also displays the forecasts along with the corresponding lower and upper 95% prediction limits.

Single Exponential Smoothing

Single Exponential Smoothing

Stat > Time Series > Single Exp Smoothing

Single exponential smoothing smoothes your data by computing exponentially weighted averages and provides short-term forecasts. This procedure works best for data without a trend or seasonal component. See [1], [4], and [6] for a discussion of exponential smoothing methods.

Dialog box items

Variable: Enter the column containing the time series.

Weight to Use in Smoothing: Use the options below to specify which weight to use. See Computing Weights, or Smoothed Values for details.

Optimal ARIMA: Choose to use the default weight, which Minitab computes by fitting an ARIMA (0, 1, 1) model to the data. With this option, Minitab calculates the initial smoothed value by backcasting.

Use: Choose to enter a specific weight, then type a number greater than or equal to 0 and less than 2. With this option, Minitab uses the average of the first six observations (or all the observations if there are less than six observations) for the initial smoothed value by default. You can change this default by specifying a different value in the Single Exponential Smoothing - Options dialog box.

Generate forecasts: Check to generate forecasts. Forecasts appear in green on the time series plot with 95% prediction interval bands.

Number of forecasts: Enter an integer to indicate how many forecasts you want.

Starting from origin: Enter a positive integer to specify a starting point for the forecasts. For example, if you specify 4 forecasts and 48 for the origin, Minitab computes forecasts for periods 49, 50, 51, and 52, based on the smoothed value at period 48. If you leave this space blank, Minitab generates forecasts from the end of the data.

Data – Single Exponential Smoothing

Your time series must be in a numeric column.

Time Series

The time series *cannot* include any missing values. If you have missing values, you may want to provide estimates of the missing values. If you

- Have seasonal data, estimate the missing values as the fitted values from the decomposition procedure. Replace the missing values in the series with the corresponding fitted values computed by Decomposition.
- Do not have seasonal data, estimate the missing values as the fitted values from the moving average procedure. Replace the missing value with the fitted value computed by Moving Average.

To perform single exponential smoothing

- 1 Choose **Stat > Time Series > Single Exp Smoothing**.
- 2 In **Variable**, enter the column containing the time series.
- 3 If you like, use any dialog box options, then click **OK**.

Computing Weights, or Smoothed Values

The weight is the smoothing parameter. You can have Minitab supply some optimal weight (the default) or you can specify a weight.

Large weights result in more rapid changes in the fitted line; small weights result in less rapid changes in the fitted line. Therefore, the larger the weights, the more the smoothed values follow the data; the smaller the weights, the less jagged the pattern is in the smoothed values. Thus, small weights are usually recommended for a series with a high noise level around the signal or pattern.


To specify your own weight – Single Exponential Smoothing

In the main Single Exponential Smoothing dialog box, choose **Use** under **Weight to use in smoothing**, and enter a value between 0 and 2, although the usual choices are between 0 and 1.

You can use a rule of thumb for choosing a weight.

- A weight α will give smoothing that is approximately equivalent to an unweighted moving average of length $(2 - \alpha) / \alpha$.
- Conversely, if you want a weight to give a moving average of approximate length l , specify the weight to be $2 / (l + 1)$.

Single Exponential Smoothing : When to Use



Use for:

- Data with no trend, and
- Data with no seasonal pattern
- Short term forecasting

Forecast profile:

- Flat line

ARIMA equivalent: (0,1,1) model

Forecasts – Single Exponential Smoothing

The fitted value at time t is the smoothed value at time $t-1$. The forecasts are the fitted value at the forecast origin. If you forecast 10 time units ahead, the forecasted value for each time will be the fitted value at the origin. Data up to the origin are used for the smoothing.

In naive forecasting, the forecast for time t is the data value at time $t-1$. Perform single exponential smoothing with a weight of one to give naive forecasting.

Time

Stat > Time Series > menu command > Time

Specify the time scale.

Dialog box items

Time Scale: Change the x-axis of the scale you specify.

Index: Choose to number the x-axis with a single scale from 1 to n by ones (where n is the number of observations in the column containing the time series).

Calendar: Choose to label the x-axis with different time unit scales. For example, if you choose Month Quarter Year, Minitab assumes that your data are in intervals of 1 month and generates 3 scales on the x-axis: 1 for months, 1 for quarters, and 1 for years. At each 4th tick on the month scale, Minitab generates a tick on the quarter scale. At each 4th tick on the quarter scale, Minitab generates a tick on the Year scale.

Clock: Choose to label the x-axis with different time unit scales. For example, if you choose Day Hour Minute, Minitab assumes that your data are in minute intervals and generates 3 scales on the x-axis: 1 for days, 1 for hours, and 1 for minutes. At each 60th tick on the minute scale, Minitab generates a tick on the hour scale. At each 24th tick on the hour scale, Minitab generates a tick on the Day scale.

Start value(s): Enter the integer start values.

Increment: Enter a value to increment the time scale.

Stamp: Choose to add a date/time stamp on the x-axis. Enter a date/time column in the text box. The values are stamped on the plot in the same format as they appear in the column.

Single Exponential Smoothing – Options

Stat > Time Series > Single Exp Smoothing > Options

Specify a customized title and the initial smoothed value.

Dialog box items

Title: To replace the default title with your own custom title, type the desired text in this box.

Set initial smoothed value

Use average of first [] observations: By default, if you specify a weight, Minitab uses the average of the first six observations for the initial smoothed value. You can change this default by entering a different value here. If you instructed Minitab to use an optimal weight, this option will not be available since in that case Minitab computes the initial smoothed value by backcasting.

Single Exponential Smoothing – Storage

Stat > Time Series > Single Exp Smoothing > Storage

Stores the smoothed values, the fits or predicted values (smoothed value at time $t-1$), the residuals (data - fits), forecasts, and upper and lower 95% prediction limits.

Dialog box items

Smoothed data: Check to store the smoothed data. The smoothed value at time T is the fitted value for time $T+1$.

Fits (one-period-ahead forecasts): Check to store the fitted values. You should store the fitted values if you want to generate diagnostic residual plots.

Residuals: Check to store the residuals. The residual at time T is the difference between the actual data at time T and the fitted value at time T . These residuals are used to calculate MAPE, MAD and MSD. If you store the residuals you can generate diagnostic plots using Autocorrelation.

Forecasts: Check to store the forecasts. This option is available only if you generated forecasts in the main Single Exponential Smoothing dialog box.

Upper 95% Prediction Interval: Check to store the upper 95% prediction limits for the forecasts.

Lower 95% Prediction Interval: Check to store the lower 95% prediction limits for the forecasts.

Single Exponential Smoothing – Graphs

Stat > Time Series > Single Exp Smoothing > Graphs

Displays time series plot and residual plots for diagnosis of model fit.

Dialog box items

Time series plot (including optional forecasts)

Plot predicted vs. actual: Choose to generate a time series plot which displays the data and one-period-ahead forecasts, or fitted values.

Plot smoothed vs. actual: Choose to generate a time series plot which displays the data and the smoothed, or moving average, values.

Do not display plot: Choose to suppress the time series plot.

Residual Plots

Individual plots: Choose to display one or more plots.

Histogram of residuals: Check to display a histogram of the residuals.

Normal plot of residuals: Check to display a normal probability plot of the residuals.

Residuals versus fits: Check to plot the residuals versus the fitted values.

Residuals versus order: Check to plot the residuals versus the order of the data. The row number for each data point is shown on the x-axis—for example, 1 2 3 4... n.

Four in one: Choose to display a layout of a histogram of residuals, a normal plot of residuals, a plot of residuals versus fits, and a plot of residuals versus order.

Residuals versus the variables: Enter one or more columns containing the variables against which you want to plot the residuals. Minitab displays a separate graph for each column.

Single Exponential Smoothing – Results

Stat > Time Series > Single Exp Smoothing > Results

Controls Session window output.

Dialog box items

Control the Display of Results

Display nothing: Choose to suppress output.

Summary table: Choose to display the default non-graphical output – the smoothing constant and accuracy measures.

Summary table and results table: Choose to display the default output plus a table of the original series, smoothed data, the one-period-ahead forecasts (or fitted values), and the one-period-ahead errors. If you generated forecasts, the table also includes the forecasts and their upper and lower 95% prediction limits.

Example of Single Exponential Smoothing

You wish to predict employment over 6 months in a segment of the metals industry using data collected over 60 months. You use single exponential smoothing because there is no clear trend or seasonal pattern in the data.

- 1 Open the worksheet EMPLOY.MTW.
- 2 Choose **Stat > Time Series > Single Exp Smoothing**.
- 3 In **Variable**, enter *Metals*.
- 4 Check **Generate forecasts**, and enter 6 in **Number of forecasts**. Click **OK**.

Session window output

Single Exponential Smoothing for Metals

```
Data      Metals
Length    60
```

Smoothing Constant

Alpha 1.04170

Accuracy Measures

MAPE 1.11648

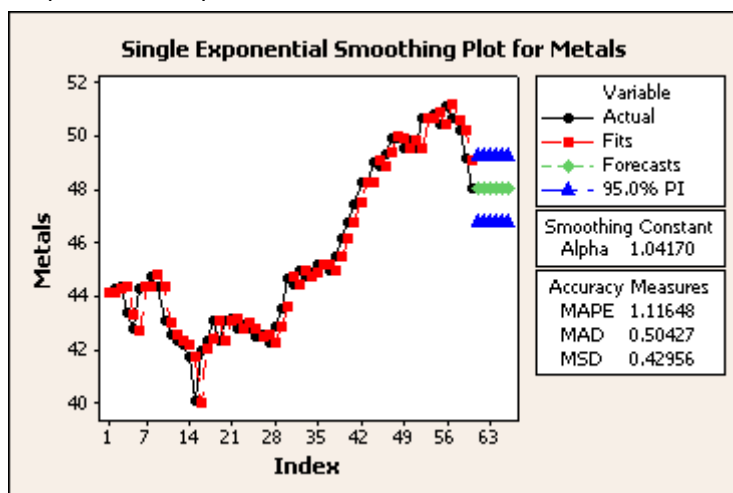
MAD 0.50427

MSD 0.42956

Forecasts

Period	Forecast	Lower	Upper
61	48.0560	46.8206	49.2914
62	48.0560	46.8206	49.2914
63	48.0560	46.8206	49.2914
64	48.0560	46.8206	49.2914
65	48.0560	46.8206	49.2914
66	48.0560	46.8206	49.2914

Graph window output



Interpreting the results

Minitab generated the default time series plot which displays the series and fitted values (one-period-ahead forecasts), along with the six forecasts.

In both the Session and Graph windows, Minitab displays the smoothing constant (weight) used and three measures to help you to determine the accuracy of the fitted values: MAPE, MAD, and MSD (see Measures of accuracy). The three accuracy measures, MAPE, MAD, and MSD, were 1.12, 0.50, and 0.43, respectively for the single exponential smoothing model, compared to 1.55, 0.70, and 0.76, respectively, for the moving average fit (see Example of moving average). Because these values are smaller for single exponential smoothing, you can judge that this method provides a better fit to these data.

Double Exponential Smoothing

Double Exponential Smoothing

Stat > Time Series > Double Exp Smoothing

Double exponential smoothing smoothes your data by Holt (and Brown as a special case) double exponential smoothing and provides short-term forecasts. This procedure can work well when a trend is present but it can also serve as a general smoothing method. Dynamic estimates are calculated for two components: level and trend.

Dialog box items

Variable: Enter the column containing the time series.

Weight to Use in Smoothing: The method Minitab uses to calculate level and trend components is determined by the option chosen below. See Calculating Level and Trend Components for details.

Optimal ARIMA: Choose to use the default weights, or smoothing parameters, which Minitab computes by fitting an ARIMA (0,2,2) model to the data.

Use: Choose to enter specific values for the smoothing parameters. You must specify the appropriate weights greater than 0 and less than 2 for the level component and greater than 0 and less than $[4 / (\text{weight for level component}) - 2]$ for the trend component.

Generate forecasts: Check to generate forecasts. Forecasts appear in green on the time series plot with 95% prediction interval bands.

Number of forecasts: Enter an integer to indicate how many forecasts that you want.

Starting from origin: Enter a positive integer to specify a starting point for the forecasts. For example, if you specify 4 forecasts and 48 for the origin, Minitab computes forecasts for periods 49, 50, 51, and 52, based on the level and trend components at period 48. If you leave this space blank, Minitab generates forecasts from the end of the data.

Data – Double Exponential Smoothing

Your time series must be in a numeric column.


The time series cannot include any missing values. If you have missing values, you may want to provide estimates of the missing values. If you:

- Have seasonal data, estimate the missing values as the fitted values from the decomposition procedure. Replace the missing values in the series with the corresponding fitted values computed by Decomposition.
- Do not have seasonal data, estimate the missing values as the fitted values from the moving average procedure. Replace the missing value with the fitted value computed by Moving Average.

To perform double exponential smoothing

- 1 Choose **Stat > Time Series > Double Exp Smoothing**.
- 2 In **Variable**, enter the column containing the time series.
- 3 If you like, use any dialog box options, then click **OK**.

Double Exponential Smoothing : When to Use



Use for:

- Data with constant or non-constant trend, and
- Data with no seasonal pattern
- Short term forecasting

Forecast profile:

- Straight line with slope equal to last trend estimate

ARIMA equivalent: (0,2,2) model

Level and Trend Components – Double Exponential Smoothing

Double exponential smoothing employs a level component and a trend component at each period. It uses two weights, or smoothing parameters, to update the components at each period. The double exponential smoothing equations are:

$$L_t = \alpha Y_t + (1 - \alpha) [L_{t-1} + T_{t-1}]$$

$$T_t = \gamma [L_t - L_{t-1}] + (1 - \gamma) T_{t-1}$$

$$\hat{Y}_t = L_{t-1} + T_{t-1}$$

where L_t is the level at time t , α is the weight for the level, T_t is the trend at time t , γ is the weight for the trend, Y_t is the data value at time t , and \hat{Y}_t is the fitted value, or one-step-ahead forecast, at time t .

....You must provide the initial level and trend to proceed in one of two ways..... the first observation is numbered one, then level and trend estimates at time zero must be initialized in order to proceed. The initialization method used to determine how the smoothed values are obtained in one of two ways: with Minitab generated weights or with specified weights.

Optimal ARIMA weights

- 1 Minitab fits an ARIMA (0,2,2) model to the data, in order to minimize the sum of squared errors.
- 2 The trend and level components are then initialized by backcasting.

Specified weights

- 1 Minitab fits a linear regression model to time series data (y variable) versus time (x variable).
- 2 The constant from this regression is the initial estimate of the level component, the slope coefficient is the initial estimate of the trend component.

When you specify weights that correspond to an equal-root ARIMA (0, 2, 2) model, Holt's method specializes to Brown's method. See [4] for more information on Holt's and Brown's methods.

Choosing weights – Double Exponential Smoothing

The weights are the smoothing parameters. You can have Minitab supply some optimal weights (the default) or you can specify weights between 0 and 2 for the level component and between 0 and $[4 / \text{weight for the level component}] - 2$ for the trend component. See Method for more information.

Regardless of the component, large weights result in more rapid changes in that component; small weights result in less rapid changes. Therefore, the larger the weights the more the smoothed values follow the data; the smaller the weights the smoother the pattern in the smoothed values. The components in turn affect the smoothed values and the predicted values. Thus, small weights are usually recommended for a series with a high noise level around the signal or pattern.

To specify your own weights

In the main Double Exponential Smoothing dialog box, choose **Use under Weight** to use in smoothing, and enter a value between 0 and 2 for the level component and between 0 and $[4 / \text{weight for the level component}] - 2$ for the trend component. Values between 0 and 1 are commonly used for both smoothing parameters.

Forecasts – Double Exponential Smoothing

Double exponential smoothing uses the level and trend components to generate forecasts. The forecast for m periods ahead from a point at time t is

$L_t + mT_t$, where L_t is the level and T_t is the trend at time t .

Data up to the forecast origin time will be used for the smoothing.

Time

Stat > Time Series > menu command > Time

Specify the time scale.

Dialog box items

Time Scale: Change the x-axis of the scale you specify.

Index: Choose to number the x-axis with a single scale from 1 to n by ones (where n is the number of observations in the column containing the time series).

Calendar: Choose to label the x-axis with different time unit scales. For example, if you choose Month Quarter Year, Minitab assumes that your data are in intervals of 1 month and generates 3 scales on the x-axis: 1 for months, 1 for

quarters, and 1 for years. At each 4th tick on the month scale, Minitab generates a tick on the quarter scale. At each 4th tick on the quarter scale, Minitab generates a tick on the Year scale.

Clock: Choose to label the x-axis with different time unit scales. For example, if you choose Day Hour Minute, Minitab assumes that your data are in minute intervals and generates 3 scales on the x-axis: 1 for days, 1 for hours, and 1 for minutes. At each 60th tick on the minute scale, Minitab generates a tick on the hour scale. At each 24th tick on the hour scale, Minitab generates a tick on the Day scale.

Start value(s): Enter the integer start values.

Increment: Enter a value to increment the time scale.

Stamp: Choose to add a date/time stamp on the x-axis. Enter a date/time column in the text box. The values are stamped on the plot in the same format as they appear in the column.

Double Exponential Smoothing – Options

Stat > Time Series > Double Exp Smoothing > Options

Specify a customized title.

Dialog box items

Title: To replace the default title with your own custom title, type the desired text in this box.

Double Exponential Smoothing – Storage

Stat > Time Series > Double Exp Smoothing > Storage

Stores various items in the worksheet.

Dialog box items

Smoothed data: Check to store the smoothed data.

Level estimates: Check to store the level components. The level component at time T equals the smoothed value at time T, while adding the level and trend components at time T equals the one-period-ahead forecast for time T+1.

Trend estimates: Check to store the trend components.

Fits (one-period-ahead-forecasts): Check to store the fitted values. You should store the fitted values if you want to generate diagnostic residual plots.

Residuals: Check to store the residuals. The residual at time T is the difference between the actual data at time T and the fitted value at time T. These residuals are used to calculate MAPE, MAD and MSD. If you store the residuals you can generate diagnostic plots using Autocorrelation.

Forecasts: Check to store the forecasts. This option is available only if you generated forecasts in the main Double Exponential Smoothing dialog box.

Upper 95% Prediction Interval: Check to store the upper 95% prediction limits for the forecasts.

Lower 95% Prediction Interval: Check to store the lower 95% prediction limits for the forecasts.

Double Exponential Smoothing – Graphs

Stat > Time Series > Double Exp Smoothing > Graphs

Displays time series plot and residual plots for diagnosis of model fit.

Dialog box items

Time series plot (including optional forecasts)

Plot predicted vs. actual: Choose to generate a time series plot which displays the data and one-period-ahead forecasts, or fitted values.

Plot smoothed vs. actual: Choose to generate a time series plot which displays the data and the smoothed, or moving average, values.

Do not display plot: Choose to suppress the time series plot.

Residual Plots

Individual plots: Choose to display one or more plots.

Histogram of residuals: Check to display a histogram of the residuals.

Normal plot of residuals: Check to display a normal probability plot of the residuals.

Residuals versus fits: Check to plot the residuals versus the fitted values.

Residuals versus order: Check to plot the residuals versus the order of the data. The row number for each data point is shown on the x-axis—for example, 1 2 3 4... n.

Four in one: Choose to display a layout of a histogram of residuals, a normal plot of residuals, a plot of residuals versus fits, and a plot of residuals versus order.

Residuals versus the variables: Enter one or more columns containing the variables against which you want to plot the residuals. Minitab displays a separate graph for each column.

Double Exponential Smoothing – Results

Stat > Time Series > Double Exp Smoothing > Results

Controls Session window output.

Dialog box items

Control the Display of Results

Display nothing: Choose to suppress output.

Summary table: Choose to display the default non-graphical output – the smoothing constants for level and trend, and accuracy measures.

Summary table and results table: Choose to display the default output plus a table containing the original series, the smoothed data, the fitted data and the errors. If you generated forecasts, the table also includes the forecasts and their upper and lower 95% prediction limits.

Example of Double Exponential Smoothing

You wish to predict employment over six months in a segment of the metals industry. You use double exponential smoothing as there is no clear trend or seasonal pattern in the data, and you want to compare the fit by this method with that from single exponential smoothing (see Example of single exponential smoothing).

- 1 Open the worksheet EMPLOY.MTW.
- 2 Choose **Stat > Time Series > Double Exp Smoothing**.
- 3 In **Variable**, enter *Metals*.
- 4 Check **Generate forecasts** and enter 6 in **Number of forecasts**. Click **OK**.

Session window output

Double Exponential Smoothing for Metals

```
Data      Metals
Length   60
```

Smoothing Constants

```
Alpha (level)  1.03840
Gamma (trend)  0.02997
```

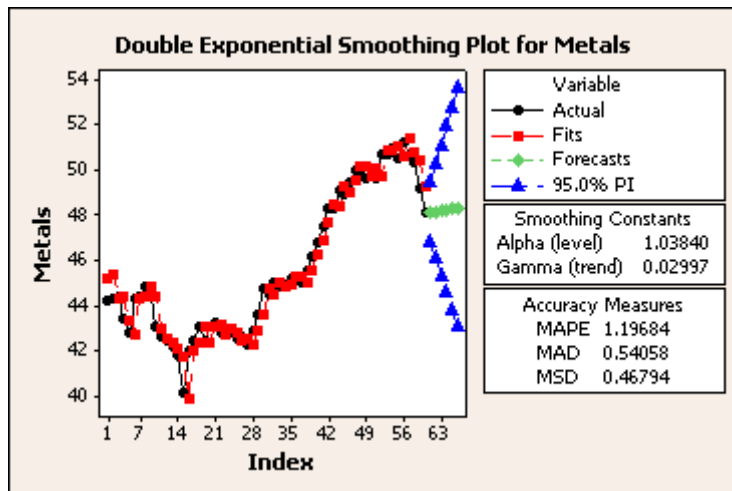
Accuracy Measures

```
MAPE  1.19684
MAD    0.54058
MSD    0.46794
```

Forecasts

Period	Forecast	Lower	Upper
61	48.0961	46.7718	49.4205
62	48.1357	46.0600	50.2113
63	48.1752	45.3135	51.0368
64	48.2147	44.5546	51.8747
65	48.2542	43.7899	52.7184
66	48.2937	43.0221	53.5652

Graph window output



Interpreting the results

Minitab generated the default time series plot which displays the series and fitted values (one-step-ahead forecasts), along with the six forecasts.

In both the Session and Graph windows, Minitab displays the smoothing constants (weights) for the level and trend components and three measures to help you determine the accuracy of the fitted values: MAPE, MAD, and MSD (see Measures of accuracy).

The three accuracy measures, MAPE, MAD, and MSD, were respectively 1.19684, 0.54058, and 0.46794 for double exponential smoothing fit, compared to 1.11648, 0.50427, and 0.42956 for the single exponential smoothing fit (see Example of single exponential smoothing). Because these values are smaller for single exponential smoothing, you can judge that this method provides a slightly better fit to these data.

Because the difference in accuracy measures for the two exponential smoothing methods are small, you might consider the type of forecast (horizontal line versus line with slope) in selecting between methods. Double exponential smoothing forecasts an employment pattern that is slightly increasing though the last four observations are decreasing. A higher weight on the trend component can result in a prediction in the same direction as the data, which may be more realistic, but the measured fit may not be as good as when Minitab generated weights are used.

Winters' Method

Winters' Method

Stat > Time Series > Winters' Method

Winters' Method smoothes your data by Holt-Winters exponential smoothing and provides short to medium-range forecasting. You can use this procedure when both trend and seasonality are present, with these two components being either additive or multiplicative. Winters' Method calculates dynamic estimates for three components: level, trend, and seasonal.

Dialog box items

Variable: Select the column containing the time series.

Seasonal Length: Enter the length of the seasonal pattern. This must be a positive integer greater than or equal to 2.

Model Type:

Multiplicative: Choose the multiplicative model when the seasonal pattern in the data depends on the size of the data. In other words, the magnitude of the seasonal pattern increases as the series goes up, and decreases as the series goes down.

Additive: Choose the additive model when the seasonal pattern in the data does not depend on the size of the data. In other words, the magnitude of the seasonal pattern does not change as the series goes up or down.

Weights to Use in Smoothing: By default, all three weights, or smoothing parameters, are set to 0.2. Since an equivalent ARIMA model exists only for a very restricted form of the Holt-Winters model, optimal parameters are not found for Winters' Method as they are for Single Exponential Smoothing and Double Exponential Smoothing.

Level: Specify the level component weight; must be a number from 0 to 1.

Trend: Specify the trend component weight; must be a number from 0 to 1.

Seasonal: Specify the seasonal component weight; must be a number from 0 to 1.

Generate forecasts: Check to generate forecasts. Forecasts appear in green on the time series plot with 95% prediction interval bands.

Number of forecasts: Enter an integer to indicate how many forecasts you want.

Starting from origin: Enter a positive integer to specify a starting point for the forecasts. For example, if you specify 4 forecasts and 48 for the origin, Minitab computes forecasts for periods 49, 50, 51, and 52, based on the level and trend components at period 48, and the corresponding seasonal components. If you leave this space blank, Minitab generates forecasts from the end of the data.

Data – Winters' Method

Your time series must be in one numeric column.


The time series *cannot* include any missing values. If you have missing values, you may want to provide estimates of the missing values. If you:

- Have seasonal data, estimate the missing values as the fitted values from the decomposition procedure. Replace the missing values in the series with the corresponding fitted values computed by Decomposition.
- Do not have seasonal data, estimate the missing values as the fitted values from the moving average procedure. Replace the missing value with the fitted value computed by Moving Average.

To perform an exponential smoothing by Winters' Method

- 1 Choose **Stat > Time Series > Winters' Method**.
- 2 In **Variable**, enter the column containing the time series.
- 3 In **Seasonal length**, enter a number ≥ 2 for the period or seasonal length.
- 4 If you like, use any dialog box options, then click **OK**.

Winters' Method, Additive Model : When to Use



Use for:


- Data with or without trend
- Data with seasonal pattern
- Size of seasonal pattern not proportional to data
- Short to medium range forecasting

Forecast profile:

- Straight line with seasonal pattern added on

ARIMA equivalent: none

Winters' Method, Multiplicative Model : When to Use



Use for:

- Data with or without trend
- Data with seasonal pattern
- Size of seasonal pattern proportional to data
- Short to medium range forecasting

Forecast profile:

- Straight line multiplied by seasonal pattern

ARIMA equivalent: none

Level, Trend, and Seasonal Components – Winters' Method

Winters' method employs a level component, a trend component, and a seasonal component at each period. It uses three weights, or smoothing parameters, to update the components at each period. Initial values for the level and trend components are obtained from a linear regression on time. Initial values for the seasonal component are obtained from a dummy-variable regression using detrended data. The Winters' method smoothing equations are:

- Additive model:

$$L_t = \alpha (Y_t - S_{t-p}) + (1 - \alpha) [L_{t-1} + T_{t-1}]$$

$$T_t = \gamma [L_t - L_{t-1}] + (1 - \gamma) T_{t-1}$$

$$S_t = \delta (Y_t - L_t) + (1 - \delta) S_{t-p}$$

$$\hat{Y}_t = L_{t-1} + T_{t-1} + S_{t-p}$$

- Multiplicative model:

$$L_t = \alpha (Y_t / S_{t-p}) + (1 - \alpha) [L_{t-1} + T_{t-1}]$$

$$T_t = \gamma [L_t - L_{t-1}] + (1 - \gamma) T_{t-1}$$

$$S_t = \delta (Y_t / L_t) + (1 - \delta) S_{t-p}$$

$$\hat{Y}_t = (L_{t-1} + T_{t-1}) S_{t-p}$$

where

- L_t is the level at time t
- α is the weight for the level
- T_t is the trend at time t
- γ is the weight for the trend
- S_t is the seasonal component at time t
- δ is the weight for the seasonal component
- p is the seasonal period
- Y_t is the data value at time t
- \hat{Y}_t is the fitted value, or one-period-ahead forecast, at time t

Choosing weights – Winters' Method

You can enter weights, or smoothing parameters, for the level, trend, and seasonal components. The default weights are 0.2 and you can enter values between 0 and 1. Since an equivalent ARIMA model exists only for a very restricted form of the Holt-Winters model, Minitab does not compute optimal parameters for Winters' method as it does for single and double exponential smoothing.

Regardless of the component, large weights result in more rapid changes in that component; small weights result in less rapid changes. The components in turn affect the smoothed values and the predicted values. Thus, small weights are usually recommended for a series with a high noise level around the signal or pattern. Large weights are usually recommended for a series with a small noise level around the signal.

Winters' Method – An additive or a multiplicative model?

The Holt-Winters' model is multiplicative when the level and seasonal components are multiplied together and it is additive when they are added together. Choose the multiplicative model when the magnitude of the seasonal pattern in the data depends on the magnitude of the data. In other words, the magnitude of the seasonal pattern increases as the data values increase, and decreases as the data values decrease.

Choose the additive model when the magnitude of the seasonal pattern in the data does not depend on the magnitude of the data. In other words, the magnitude of the seasonal pattern does not change as the series goes up or down.

Forecasts – Winters' Method

Winters' method uses the level, trend, and seasonal components to generate forecasts. The forecast for m periods ahead from a point at time t is:

$$L_t + mT_t$$

where L_t is the level and T_t is the trend at time t , multiplied by (or added to for an additive model) the seasonal component for the same period from the previous year.

Winters' Method uses data up to the forecast origin time to generate the forecasts.

Time

Stat > Time Series > menu command > Time

Specify the time scale.

Dialog box items

Time Scale: Change the x-axis of the scale you specify.

Index: Choose to number the x-axis with a single scale from 1 to n by ones (where n is the number of observations in the column containing the time series).

Calendar: Choose to label the x-axis with different time unit scales. For example, if you choose Month Quarter Year, Minitab assumes that your data are in intervals of 1 month and generates 3 scales on the x-axis: 1 for months, 1 for quarters, and 1 for years. At each 4th tick on the month scale, Minitab generates a tick on the quarter scale. At each 4th tick on the quarter scale, Minitab generates a tick on the Year scale.

Clock: Choose to label the x-axis with different time unit scales. For example, if you choose Day Hour Minute, Minitab assumes that your data are in minute intervals and generates 3 scales on the x-axis: 1 for days, 1 for hours, and 1 for minutes. At each 60th tick on the minute scale, Minitab generates a tick on the hour scale. At each 24th tick on the hour scale, Minitab generates a tick on the Day scale.

Start value(s): Enter the integer start values.

Increment: Enter a value to increment the time scale.

Stamp: Choose to add a date/time stamp on the x-axis. Enter a date/time column in the text box. The values are stamped on the plot in the same format as they appear in the column.

Winters' Method – Options

Stat > Time Series > Winters' Method > Options

Specify a customized title and the first observation in the seasonal period.

Dialog box items

Title: To replace the default title with your own custom title, type the desired text in this box.

First obs. is in seasonal period: By default this value is 1 because Minitab assumes that the first data value in the series corresponds to the first seasonal period. Enter a different number to specify a different starting value. For example, if you have monthly data and the first observation is in June, then enter 6 to set the seasonal period correctly.

Winters' Method – Storage

Stat > Time Series > Winters' Method > Storage

Stores various values in the worksheet.

Dialog box items

Smoothed data: Check to store the smoothed data. The smoothed value for time T is equal to the seasonal estimate for time $T - p$ (where p is the seasonal period) multiplied by (or added to for an additive model) the level estimate for time $T - 1$.

Level estimates: Check to store the level components.

Trend estimates: Check to store the trend components.

Seasonal estimates: Check to store the seasonal components.

Fits (one-period-ahead forecasts): Check to store the fitted values. You should store the fitted values if you want to generate diagnostic residual plots. The fitted value for time T is equal to the sum of the level and trend estimates for time $T - 1$, multiplied by (or added to for an additive model) the seasonal estimate for time $T - p$ (where p is the seasonal period).

Residuals: Check to store the one-period-ahead forecast errors. These residuals are used to calculate MAPE, MAD and MSD. If you store the residuals you can generate diagnostic plots using Autocorrelation.

Forecasts: Check to store the forecasts. This option is available only if you generated forecasts in the main Winters' Method dialog box.

Upper 95% Prediction Interval: Check to store the upper 95% prediction limits for the forecasts.

Lower 95% Prediction Interval: Check to store the lower 95% prediction limits for the forecasts.

Winters' Method – Graphs

Stat > Time Series > Winters' Method > Graphs

Displays time series plot and residual plots for diagnosis of model fit.

Dialog box items

Time series plot (including optional forecasts)

Plot predicted vs. actual: Choose to generate a time series plot which displays the data and one-period-ahead forecasts, or fitted values.

Plot smoothed vs. actual: Choose to generate a time series plot which displays the data and the smoothed, or moving average, values.

Do not display plot: Choose to suppress the time series plot.

Residual Plots

Individual plots: Choose to display one or more plots.

Histogram of residuals: Check to display a histogram of the residuals.

Normal plot of residuals: Check to display a normal probability plot of the residuals.

Residuals versus fits: Check to plot the residuals versus the fitted values.

Residuals versus order: Check to plot the residuals versus the order of the data. The row number for each data point is shown on the x-axis—for example, 1 2 3 4... n.

Four in one: Choose to display a layout of a histogram of residuals, a normal plot of residuals, a plot of residuals versus fits, and a plot of residuals versus order.

Residuals versus the variables: Enter one or more columns containing the variables against which you want to plot the residuals. Minitab displays a separate graph for each column.

Winters' Method – Results

Stat > Time Series > Winters' Method > Results

Controls Session window output.

Dialog box items

Control the Display of Results

Display nothing: Choose to suppress output.

Summary table: Choose to display the default non-graphical output – the smoothing constants for level and trend, and accuracy measures.

Summary table and results table: Choose to display the default output plus a table containing the original series, the smoothed data, the fitted data and the errors. If you generated forecasts, the table also includes the forecasts and their upper and lower 95% prediction limits.

Example Winters' Method

You wish to predict employment for the next six months in a food preparation industry using data collected over the last 60 months. You use Winters' method with the default multiplicative model, because there is a seasonal component, and possibly trend, apparent in the data.

- 1 Open the worksheet EMPLOY.MTW.
- 2 Choose **Stat > Time Series > Winters' Method**.
- 3 In **Variable**, enter *Food*. In **Seasonal length**, enter 12.
- 4 Under **Model Type**, choose **Multiplicative**.
- 5 Check **Generate forecasts** and enter 6 in **Number of forecasts**. Click **OK**.

Session window output

Winters' Method for Food

Multiplicative Method

```
Data      Food
Length    60
```

Smoothing Constants

```
Alpha (level)      0.2
Gamma (trend)      0.2
Delta (seasonal)   0.2
```

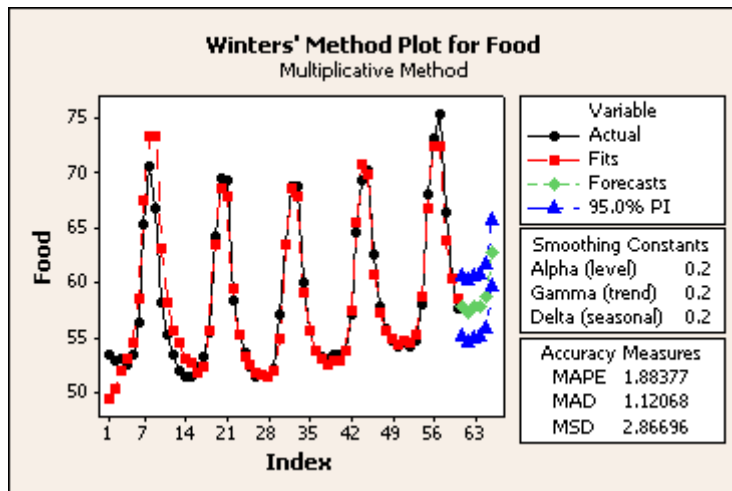
Accuracy Measures

```
MAPE  1.88377
MAD    1.12068
MSD    2.86696
```

Forecasts

Period	Forecast	Lower	Upper
61	57.8102	55.0646	60.5558
62	57.3892	54.6006	60.1778
63	57.8332	54.9966	60.6698
64	57.9307	55.0414	60.8199
65	58.8311	55.8847	61.7775
66	62.7415	59.7339	65.7492

Graph window output



Interpreting the results

Minitab generated the default time series plot which displays the series and fitted values (one-period-ahead forecasts), along with the six forecasts.

In both the Session and Graph windows, Minitab displays the smoothing constants (weights) used for level, trend, and seasonal components used and three measures to help you determine the accuracy of the fitted values: MAPE, MAD, and MSD (see Measures of accuracy).

For these data, MAPE, MAD, and MSD were 1.88, 1.12, and 2.87, respectively, with the multiplicative model. MAPE, MAD, and MSD were 1.95, 1.15, and 2.67, respectively (output not shown) with the additive model, indicating that the multiplicative model provided a slightly better fit according to two of the three accuracy measures.

Differences

Differences

Stat > Time Series > Differences

Differences computes the differences between data values of a time series. If you wish to fit an ARIMA model but there is trend or seasonality present in your data, differencing data is a common step in assessing likely ARIMA models. Differencing is used to simplify the correlation structure and to help reveal any underlying pattern.

Dialog box items

Series: Enter the column containing the variable for which you want to compute differences.

Store differences in: Enter a storage column for the differences.

Lag: Enter the value for the lag. The default lag value is 1.

Data – Differences

Your time series must be in one numeric column. Minitab stores the difference for missing data as missing (*).

To perform Differences

- 1 Choose **Stat > Time Series > Differences**.
- 2 In **Series**, enter a column containing the series that you wish to difference.
- 3 In **Store differences in**, enter a name for the storage column.
- 4 If you like, change the lag, then click **OK**.

Lag

Lag

Stat > Time Series > Lag

Lag computes lags of a column and stores them in a new column. To lag a time series, Minitab moves the data down the column and inserts missing value symbols, *, at the top of the column. The number of missing values inserted depends upon the length of the lag.

Dialog box items

Series: Enter the column containing the variable that you want to lag.

Store lags in: Enter the storage column for the lags.

Lag: Enter the value for the lag. The default lag value is 1.

Data – Lag

The data must be in one column. Minitab stores the lag for missing data with the missing value symbol (*). Data can be either numeric or text.

To Lag a time series

- 1 Choose **Stat > Time Series > Lag**.
- 2 In **Series**, enter a column containing the series that you wish to lag.
- 3 In **Store lags in**, enter a name for the storage column.
- 4 If you like, change the value in **Lag**, then click **OK**.

Autocorrelation

Autocorrelation Function

Stat > Time Series > Autocorrelation

Autocorrelation computes and plots the autocorrelations of a time series. Autocorrelation is the correlation between observations of a time series separated by k time units. The plot of autocorrelations is called the autocorrelation function or ACF. View the ACF to guide your choice of terms to include in an ARIMA model. See Fitting an ARIMA model.

Dialog box items

Series: Enter the column containing the time series.

Default number of lags: Choose to use the default number of lags, which is $n / 4$ for a series with less than or equal to 240 observations or $\sqrt{n} + 45$ for a series with more than 240 observations, where n is the number of observations in the series.

Number of lags: Choose to enter the number of lags to use instead of the default. The maximum number of lags is $n - 1$.

Store ACF: Check to store the autocorrelation values in the next available column.

Store t statistics: Check to store the t-statistics.

Store Ljung-Box Q statistics: Check to store the Ljung-Box Q statistics.

Title: Enter a new title to replace the default title on the graphical output.

Data – Autocorrelation

Your time series must be entered in one numeric column. You must either estimate or delete missing data before using this procedure.

To perform an autocorrelation function

- 1 Choose **Stat > Time Series > Autocorrelation**.
- 2 In **Series**, enter the column containing the time series.
- 3 If you like, use any dialog box options, then click **OK**.

Using the Ljung-Box Q statistic

You can use the Ljung-Box Q (LBQ) statistic to test the null hypothesis that the autocorrelations for all lags up to lag k equal zero. This is automatically done for you if you use ARIMA. However, you may find it useful to do this for autocorrelation.

See Example of Testing Autocorrelations for more information.

Example of Autocorrelation

You wish to predict employment in a food preparation industry using past employment data. You want to use ARIMA to do this but first you use autocorrelation in order to help identify a likely model. Because the data exhibit a strong 12 month seasonal component, you take a difference at lag 12 in order to induce stationarity and look at the autocorrelation of the differenced series. There may be some long-term trend in these data, but the magnitude of it appears to be small compared to the seasonal component. If the trend was larger, you might consider taking another difference at lag 1 to induce stationarity.

- 1 Open the worksheet EMPLOY.MTW.
- 2 Choose **Stat > Time Series > Differences**.
- 3 In **Series**, enter *Food*.
- 4 In **Store differences in**, enter *Food2*.
- 5 In **Lag**, enter 12. Click **OK**.
- 6 Choose **Stat > Time Series > Autocorrelation**.
- 7 In **Series**, enter *Food2*. Click **OK**.

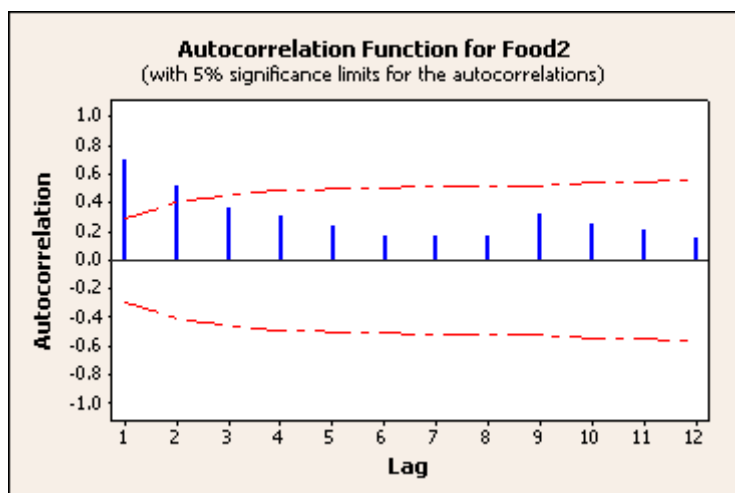
Session window output

Autocorrelation Function: Food2

Lag	ACF	T	LBQ
1	0.701388	4.86	25.12
2	0.512266	2.52	38.81
3	0.366882	1.60	45.99
4	0.310364	1.29	51.24
5	0.234743	0.94	54.32
6	0.173069	0.68	56.03
7	0.162046	0.63	57.57
8	0.170051	0.66	59.30
9	0.322438	1.24	65.70
10	0.252774	0.94	69.74
11	0.208020	0.76	72.54
12	0.150936	0.55	74.06

Autocorrelation for Food2

Graph window output



Interpreting the results

In the Session window, Minitab displays the autocorrelations, associated t-statistics, and Ljung-Box Q statistics. Because you did not specify the lag length, autocorrelation uses the default length of $n / 4$ for a series with less than or equal to 240 observations. Minitab generates an autocorrelation function (ACF) with approximate $\alpha = 0.05$ critical bands for the hypothesis that the correlations are equal to zero.

The ACF for these data shows large positive, significant spikes at lags 1 and 2 with subsequent positive autocorrelations that do not die off quickly. This pattern is typical of an autoregressive process.

See Example of Testing Autocorrelations to test the null hypothesis that the autocorrelations for all lags up to a lag of 6 are zero.

Example of Testing Autocorrelations

Using the results from the Example of Autocorrelation, you can use the Ljung-Box Q (LBQ) statistic to test the null hypothesis that the autocorrelations for all lags up to lag k equal zero. Let's test that all autocorrelations up to a lag of 6 are zero. The LBQ statistic is 56.03.

Step 1: Compute the cumulative probability function

- 1 Choose **Calc > Probability Distributions > Chi-Square**.
- 2 Choose **Cumulative Probability**.
- 3 In **Degrees of freedom**, enter 6 (the lag of your test).
- 4 Choose **Input constant** and enter 56.03 (the LBQ value).
- 5 In **Optional storage**, enter *Cumprob*. This stores the cumulative probability function in a constant named *Cumprob*. Click **OK**.

Step 2: Compute the p-value

- 1 Choose **Calc > Calculator**.
- 2 In **Store result in variable** enter *pvalue*.
- 3 In **Expression**, enter $1 - 'Cumprob'$. Click **OK**.

Interpreting the results

Examine the value in the Data window. In this example, the p-value is 0.000000, which means the p-value is less than 0.0000005. The very small p-value implies that one or more of the autocorrelations up to lag 6 can be judged as significantly different from zero at any reasonable α level.

Partial Autocorrelation

Partial Autocorrelation Function

Stat > Time Series > Partial Autocorrelation

Partial autocorrelation computes and plots the partial autocorrelations of a time series. Partial autocorrelations, like autocorrelations, are correlations between sets of ordered data pairs of a time series. As with partial correlations in the regression case, partial autocorrelations measure the strength of relationship with other terms being accounted for. The partial autocorrelation at a lag of k is the correlation between residuals at time t from an autoregressive model and observations at lag k with terms for all intervening lags present in the autoregressive model. The plot of partial autocorrelations is called the partial autocorrelation function or PACF. View the PACF to guide your choice of terms to include in an ARIMA model. See Fitting an ARIMA model.

Dialog box items

Series: Choose the column containing the time series.

Default number of lags: Choose to use the default number of lags. This is $n / 4$ for a series with less than or equal to 240 observations or $\sqrt{n} + 45$ for a series with more than 240 observations, where n is the number of observations in the series.

Number of lags: Choose to enter the number of lags to use instead of the default. The maximum number of lags is equal to $n - 1$.

Store PACF: Check to store the partial autocorrelations in the next available column.

Store t statistics: Check to store the t-statistics.

Title: Enter a new title to replace the default title on the graphical output.

To perform a partial autocorrelation function

- 1 Choose **Stat > Time Series > Partial Autocorrelation**.
- 2 In **Series**, enter the column containing the time series.
- 3 If you like, use any available dialog box options, then click **OK**.

Data – Partial Autocorrelation

Your time series must be entered in one numeric column. You must either estimate or delete missing data before using this procedure.

Example of partial autocorrelation

You obtain a partial autocorrelation function (PACF) of the food industry employment data, after taking a difference of lag 12, in order to help determine a likely ARIMA model.

- 1 Open the worksheet EMPLOY.MTW.
- 2 Choose **Stat > Time Series > Differences**.
- 3 In **Series**, enter *Food*.
- 4 In **Store differences in**, enter *Food2*.
- 5 In **Lag**, enter 12. Click **OK**.
- 6 Choose **Stat > Time Series > Partial Autocorrelation**.
- 7 In **Series**, enter *Food2*. Click **OK**.

Session window output


Partial Autocorrelation Function: Food2

Lag	PACF	T
1	0.701388	4.86
2	0.039998	0.28
3	-0.012022	-0.08
4	0.092572	0.64
5	-0.034921	-0.24
6	-0.014194	-0.10
7	0.075222	0.52

8	0.049848	0.35
9	0.326936	2.27
10	-0.227678	-1.58
11	0.005302	0.04
12	-0.000979	-0.01

Partial Autocorrelation for Food2

Graph window output



Interpreting the results

Minitab generates a partial autocorrelation function with critical bands at approximately $\alpha = 0.05$ for the hypothesis that the correlations are equal to zero. In the Data window, Minitab stores the partial autocorrelations and associated t-statistics.

In the food data example, there is a single large spike of 0.7 at lag 1, which is typical of an autoregressive process of order one. There is also a significant spike at lag 9, but you have no evidence of a nonrandom process occurring there.

Cross Correlation

Cross Correlation Function

Stat > Time Series > Cross Correlation

Computes and plots the cross correlation between two time series.

Dialog box items

First series: Select the column containing the response variable of first time series.

Second series: Select the column containing the response variable second time series.

Default number of lags: Choose to have Minitab set $K = -(\sqrt{n} + 10)$ to $K = +(\sqrt{n} + 10)$, where K is the number of lags and n is the number of observations in the series.

Number of lags: Choose to specify the number of lags desired. Then enter the number of lags.

Data – Cross Correlation

You must have two time series in separate numeric columns of equal length. You must either estimate or delete missing data before using this procedure.

To perform a cross correlation function

- 1 Choose **Stat > Time Series > Cross Correlation**.
- 2 In **First Series**, enter the column containing one time series.
- 3 In **Second Series**, enter the column containing other time series.

- If you like, specify the number of lags for which to display cross correlations, then click **OK**.

ARIMA

ARIMA

Stat > Time Series > ARIMA

Use ARIMA to model time series behavior and to generate forecasts. ARIMA fits a Box-Jenkins ARIMA model to a time series. ARIMA stands for Autoregressive Integrated Moving Average with each term representing steps taken in the model construction until only random noise remains. ARIMA modeling differs from the other time series methods discussed in this chapter in the fact that ARIMA modeling uses correlational techniques. ARIMA can be used to model patterns that may not be visible in plotted data. The concepts used in this procedure follow Box and Jenkins [2]. For an elementary introduction to time series, see [3], [10]

For information on creating an ARIMA model, see Entering the ARIMA model and ARIMA specifications.

Dialog box items

Series: Enter the column containing the response variable of the time series you want to fit.

Fit seasonal model: Check to fit a seasonal model.

Period: Specify the number of units in a complete cycle.

Autoregressive

Nonseasonal: Enter the order of the autoregressive (AR) component (p).

Seasonal: If you have a seasonal model, enter the order of the seasonal autoregressive component (P) .

Difference

Nonseasonal: Enter the number of differences (d) used to discount trends over time. At least three data points must remain after differencing.

Seasonal: If you have a seasonal model, enter the number of differences for the seasonal component (D).

Moving average

Nonseasonal: Enter the order of the moving average (MA) component (q).

Seasonal: If you have a seasonal model, enter the order of the seasonal moving average component (Q).

Include constant term in model: Check to include a constant term in the ARIMA model.

Starting values for coefficients: Check to specify the initial parameter values and then enter the column containing the values. The values must be entered in the order that the parameters appear in the output: p (AR values), P (seasonal AR values), q (MA values), Q (seasonal MA values), and then (optionally) the constant. If you do not specify the initial parameter values, Minitab uses 0.1 for the parameters with the exception of the constant.

Data – ARIMA

Your time series must be in a numeric column. Missing data in the middle of your series are not allowed. If you have missing values, you may want to provide estimates of the missing values.

To fit an ARIMA model

- Choose **Stat > Time Series > ARIMA**.
- In **Series**, enter the column containing the time series.
- For at least one of **Autoregressive** or **Moving Average** under either **Nonseasonal** or **Seasonal**, enter the number of parameters. See Entering the ARIMA model.
- If you like, use any dialog box options, then click **OK**.

Entering the ARIMA model

After you have identified one or more likely models, you need to specify the model in the main ARIMA dialog box.

- If you want to fit a seasonal model, check **Fit seasonal model** and enter a number to specify the period. The period is the span of the seasonality or the interval at which the pattern is repeated. The default period is 12.

You must check **Fit seasonal model** before you can enter the seasonal autoregressive and moving average parameters or the number of seasonal differences to take.

- To specify autoregressive and moving average parameters to include in nonseasonal or seasonal ARIMA models, enter a value from 0 to 5. The maximum is 5. At least one of these parameters must be nonzero. The total for all

parameters must not exceed 10. For most data, no more than two autoregressive parameters or two moving average parameters are required in ARIMA models.

Suppose you enter 2 in the box for **Moving Average** under **Seasonal**, the model will include first and second order moving average terms.

- To specify the number of nonseasonal and/or seasonal differences to take, enter a number in the appropriate box. If you request one seasonal difference with k as the seasonal period, the k th difference will be taken.
- To include the constant in the model, check **Include constant term in model**.
- You may want to specify starting values for the parameter estimates. You must first enter the starting values in a worksheet column in the following order: AR's (autoregressive parameters), seasonal AR's, MA's (moving average parameters), seasonal MA's, and if you checked **Include constant term in model** enter the starting value for the constant in the last row of the column. This is the same order in which the parameters appear on the output. Check **Starting values for coefficients**, and enter the column containing the starting values for each parameter included in the model. Default starting values are 0.1 except for the constant.

Fitting an ARIMA model

Box and Jenkins [2] present an interactive approach for fitting ARIMA models to time series. This iterative approach involves identifying the model, estimating the parameters, checking model adequacy, and forecasting, if desired. The model identification step generally requires judgment from the analyst.

- 1 First, decide if the data are stationary. That is, do the data possess constant mean and variance.
 - Examine a time series plot to see if a transformation is required to give constant variance.
 - Examine the ACF to see if large autocorrelations do not die out, indicating that differencing may be required to give a constant mean.

A seasonal pattern that repeats every k^{th} time interval suggests taking the k^{th} difference to remove a portion of the pattern. Most series should not require more than two difference operations or orders. Be careful not to overdifference. If spikes in the ACF die out rapidly, there is no need for further differencing. A sign of an overdifferentiated series is the first autocorrelation close to -0.5 and small values elsewhere [10].

Use **Stat > Time Series > Differences** to take and store differences. Then, to examine the ACF and PACF of the differenced series, use **Stat > Time Series > Autocorrelation** and **Stat > Time Series > Partial Autocorrelation**.

- 2 Next, examine the ACF and PACF of your stationary data in order to identify what autoregressive or moving average models terms are suggested.
 - An ACF with large spikes at initial lags that decay to zero or a PACF with a large spike at the first and possibly at the second lag indicates an autoregressive process.
 - An ACF with a large spike at the first and possibly at the second lag and a PACF with large spikes at initial lags that decay to zero indicates a moving average process.
 - The ACF and the PACF both exhibiting large spikes that gradually die out indicates that both autoregressive and moving averages processes are present.

For most data, no more than two autoregressive parameters or two moving average parameters are required in ARIMA models. See [10] for more details on identifying ARIMA models.

- 3 Once you have identified one or more likely models, you are ready to use the ARIMA procedure.
 - Fit the likely models and examine the significance of parameters and select one model that gives the best fit. See **Entering the ARIMA model**.
 - Check that the ACF and PACF of residuals indicate a random process, signified when there are no large spikes. You can easily obtain an ACF and a PACF of residual using ARIMA's **Graphs** subdialog box. If large spikes remain, consider changing the model.
 - You may perform several iterations in finding the best model. When you are satisfied with the fit, go ahead and make forecasts.

The ARIMA algorithm will perform up to 25 iterations to fit a given model. If the solution does not converge, store the estimated parameters and use them as starting values for a second fit. You can store the estimated parameters and use them as starting values for a subsequent fit as often as necessary.

Diagnostic checking

The graphs for the ACF and PACF of the ARIMA residuals include lines representing two standard errors to either side of zero. Values that extend beyond two standard errors are statistically significant at approximately $\alpha = 0.05$, and show evidence that the model has not explained all autocorrelation in the data. For ACF the distance between the lines and zero for the i^{th} autocorrelation are determined by the following formula:

$$2 \sqrt{1 + 2 \sum_{k=1}^{i-1} \Gamma_k^2} / \sqrt{n}$$

where n = the number of observations in the series, and Γ_k = the k , autocorrelation.

The distance between the lines and zero for all partial autocorrelations is $2/\sqrt{n}$.

ARIMA – Graphs

Stat > Time Series > ARIMA > Graphs

Displays a time series plot and various residual plots. You do not have to store the residuals in order to produce these plots.

Dialog box items

Time series plot (including optional forecasts): Check to display a time series plot of the series. When you use the ARIMA - Forecasts subdialog box to generate forecasts, Minitab displays the forecasts and their 95% confidence limits on the plot.

Residual Plots

ACF of residuals: Check to display an autocorrelation function for the standard or raw residuals. The graph includes critical value lines at \pm two standard errors.

PACF of residuals: Check to display a partial autocorrelation function for the standard or raw residuals. The graph includes critical value lines at \pm two standard errors.

Individual plots: Choose to display one or more plots.

Histogram of residuals: Check to display a histogram of the residuals.

Normal plot of residuals: Check to display a normal probability plot of the residuals.

Residuals versus fits: Check to plot the residuals versus the fitted values.

Residuals versus order: Check to plot the residuals versus the order of the data. The row number for each data point is shown on the x-axis--for example, 1 2 3 4... n.

Four in one: Choose to display a layout of a histogram of residuals, a normal plot of residuals, a plot of residuals versus fits, and a plot of residuals versus order.

Residuals versus the variables: Enter one or more columns containing the variables against which you want to plot the residuals. Minitab displays a separate graph for each column.

ARIMA – Forecasts

Stat > Time Series > ARIMA > Forecasts

Predicts up to 150 values for specified time series.

Dialog box items

Lead: Specify the number of forecasts that you want Minitab to generate.

Origin: Specify the origin for the forecasts to begin. If the origin is not specified, it is set to the end of the series and the forecasts are for the future.

Storage: Controls the storage of ARIMA results.

Forecasts: Specify a storage column for the forecasted values (possibly for later plotting).

Lower limits: Specify a storage column for the lower confidence limits for the forecasts.

Upper limits: Specify a storage column for the upper confidence limits for the forecasts.

ARIMA – Results

Stat > Time Series > ARIMA > Results

Control the display of results to the Session window.

Dialog box items

Control the Display of Results

Display nothing: Display nothing in the Session window.

Table of final estimates, differencing information, residual sums of squares, and number of observations:

Display a table of parameter estimates, differencing information, the number of observations, sums of squares statistics, and autocorrelation statistics.

In addition, table of estimates at each iteration (and back forecasts, if they are not dying out rapidly): In addition to the above, display a table of parameter estimates at each iteration and back forecasts if they are not dying out rapidly.

In addition, correlation matrix of estimated parameters: in addition to the above, display a correlation matrix of parameter estimates.

In addition, the back forecasts: In addition to the above, display the back forecasts.

ARIMA – Storage

Stat > Time Series > ARIMA > Storage

Stores residuals, fits, and coefficients in the worksheet.

Dialog box items

Residuals: Check to store the residuals.

Fits: Check to store the fitted values.

Coefficients: Check to store estimated coefficients.

Example of ARIMA

The ACF and PACF of the food employment data (see Example of autocorrelation and Example of partial autocorrelation) suggest an autoregressive model of order 1, or AR(1), after taking a difference of order 12. You fit that model here, examine diagnostic plots, and examine the goodness of fit. To take a seasonal difference of order 12, you specify the seasonal period to be 12, and the order of the difference to be 1. In the subsequent example, you perform forecasting.

- 1 Open the worksheet EMPLOY.MTW.
- 2 Choose **Stat > Time Series > ARIMA**.
- 3 In **Series**, enter *Food*.
- 4 Check **Fit seasonal model**. In **Period**, enter *12*. Under **Nonseasonal**, enter *1* in **Autoregressive**. Under **Seasonal**, enter *1* in **Difference**.
- 5 Click **Graphs**. Check **ACF of residuals** and **PACF of residuals**.
- 6 Click **OK** in each dialog box.

Session window output

ARIMA Model: Food

Estimates at each iteration

Iteration	SSE	Parameters	
0	95.2343	0.100	0.847
1	77.5568	0.250	0.702
2	64.5317	0.400	0.556
3	56.1578	0.550	0.410
4	52.4345	0.700	0.261
5	52.2226	0.733	0.216
6	52.2100	0.741	0.203
7	52.2092	0.743	0.201
8	52.2092	0.743	0.200
9	52.2092	0.743	0.200

Relative change in each estimate less than 0.0010

Final Estimates of Parameters

Type	Coef	SE Coef	T	P
AR 1	0.7434	0.1001	7.42	0.000
Constant	0.1996	0.1520	1.31	0.196


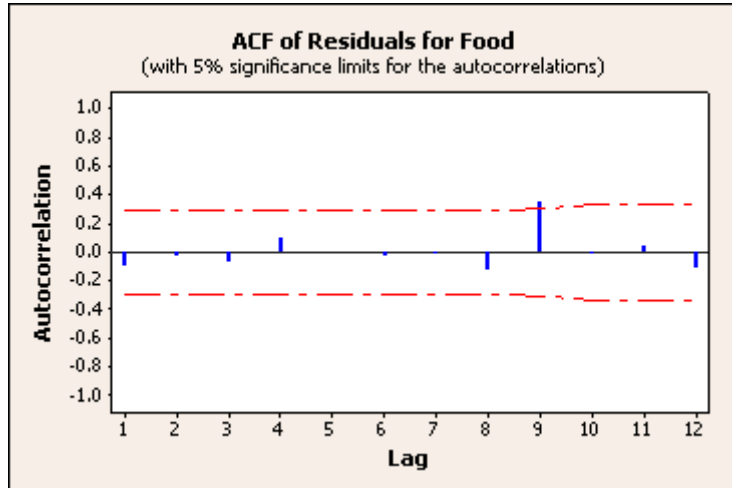
Time Series

Differencing: 0 regular, 1 seasonal of order 12
Number of observations: Original series 60, after differencing 48
Residuals: SS = 51.0364 (backforecasts excluded)
MS = 1.1095 DF = 46

Modified Box-Pierce (Ljung-Box) Chi-Square statistic

Lag	12	24	36	48
Chi-Square	11.3	19.1	27.7	*
DF	10	22	34	*
P-Value	0.338	0.641	0.768	*

Graph window output



Interpreting the results

The ARIMA model converged after nine iterations. The AR(1) parameter had a t-value of 7.42. As a rule of thumb, you can consider values over two as indicating that the associated parameter can be judged as significantly different from zero. The MSE (1.1095) can be used to compare fits of different ARIMA models.

The Ljung-Box statistics give nonsignificant p-values, indicating that the residuals appeared to uncorrelated. The ACF and PACF of the residuals corroborate this. You assume that the spikes in the ACF and PACF at lag 9 are the result of random events. The AR(1) model appears to fit well so you use it to forecast employment in the Example of Forecasting with ARIMA.

Example of Forecasting with ARIMA

In the example of fitting an ARIMA model, you found that an AR(1) model with a twelfth seasonal difference gave a good fit to the food sector employment data. You now use this fit to predict employment for the next 12 months.

Step 1: Refit the ARIMA model without displaying the acf and pacf of the residuals

1 Perform steps 1– 4 of Example of ARIMA.

Step 2: Display a time series plot

1 Click **Graphs**. Check **Time series plot**. Click **OK**.

Step 3: Generate the forecasts

1 Click **Forecast**. In **Lead**, enter 12. Click **OK** in each dialog box.

Session window output

ARIMA Model: Food

Estimates at each iteration

Iteration	SSE	Parameters	
0	95.2343	0.100	0.847
1	77.5568	0.250	0.702
2	64.5317	0.400	0.556
3	56.1578	0.550	0.410
4	52.4345	0.700	0.261
5	52.2226	0.733	0.216
6	52.2100	0.741	0.203
7	52.2092	0.743	0.201
8	52.2092	0.743	0.200
9	52.2092	0.743	0.200

Relative change in each estimate less than 0.0010

Final Estimates of Parameters

Type	Coef	SE Coef	T	P
AR 1	0.7434	0.1001	7.42	0.000
Constant	0.1996	0.1520	1.31	0.196

Differencing: 0 regular, 1 seasonal of order 12

Number of observations: Original series 60, after differencing 48

Residuals: SS = 51.0364 (backforecasts excluded)
MS = 1.1095 DF = 46

Modified Box-Pierce (Ljung-Box) Chi-Square statistic

Lag	12	24	36	48
Chi-Square	11.3	19.1	27.7	*
DF	10	22	34	*
P-Value	0.338	0.641	0.768	*

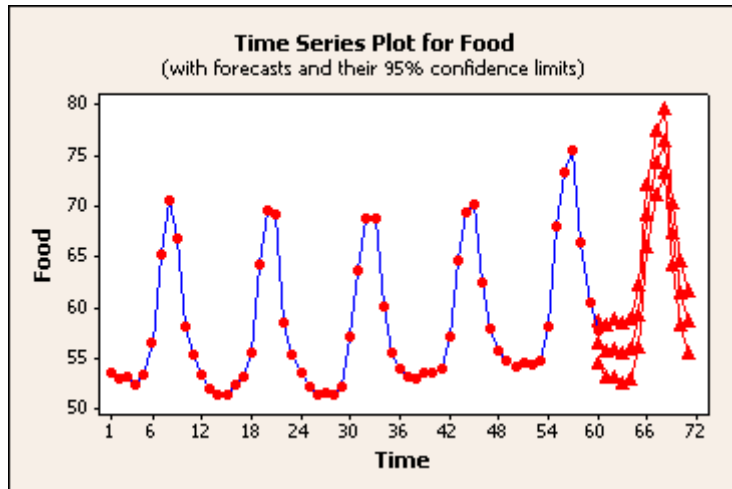
Forecasts from period 60

Period	Forecast	95 Percent Limits		Actual
		Lower	Upper	
61	56.4121	54.3472	58.4770	
62	55.5981	53.0251	58.1711	
63	55.8390	53.0243	58.6537	
64	55.4207	52.4809	58.3605	
65	55.8328	52.8261	58.8394	

Time Series

66	59.0674	56.0244	62.1104
67	69.0188	65.9559	72.0817
68	74.1827	71.1089	77.2565
69	76.3558	73.2760	79.4357
70	67.2359	64.1527	70.3191
71	61.3210	58.2360	64.4060
72	58.5100	55.4240	61.5960

Graph window output



Interpreting the results

ARIMA gives forecasts, with 95% confidence limits, using the AR(1) model in both the Session window and a Graph window. The seasonality dominates the forecast profile for the next 12 months with the forecast values being slightly higher than for the previous 12 months.

References - Time Series

- [1] B.L. Bowerman and R. T. O' Connell (1993). *Forecasting and Time Series: An Applied Approach*, 3rd edition. Duxbury Press.
- [2] G.E.P. Box and G.M. Jenkins (1994). *Time Series Analysis: Forecasting and Control*, 3rd Edition. Prentice Hall.
- [3] J.D. Cryer (1986). *Time Series Analysis*. Duxbury Press.
- [4] N.R. Farnum and L.W. Stanton (1989). *Quantitative Forecasting Methods*. PWS-Kent.
- [5] G.M. Ljung and G.E.P. Box (1978). "On a Measure of Lack of Fit in Time Series Models," *Biometrika*, 65, 67–72.
- [6] S. Makridakis, S.C. Wheelwright, and R. J. Hyndman (1998). *Forecasting: Methods and Applications*. Wiley.
- [7] D.W. Marquardt (1963). "An Algorithm for Least Squares Estimation of Nonlinear Parameters," *Journal Soc. Indust. Applied Mathematics*, 11, 431–441.
- [8] W.Q. Meeker, Jr. (1977). "TSERIES—A User-oriented Computer Program for Identifying, Fitting and Forecasting ARIMA Time Series Models," ASA 1977 Proceedings of the Statistical Computing Section.
- [9] W.Q. Meeker, Jr. (1977). *TSERIES User's Manual*, Statistical Laboratory, Iowa State University.
- [10] W. Vandaele (1983). *Applied Time Series and Box-Jenkins Models*. Academic Press, Inc.

Acknowledgment

The ARIMA algorithm is based on the fitting routine in the TSERIES package written by Professor William Q. Meeker, Jr., of Iowa State University [8], [9]. We are grateful to Professor Meeker for his help in the adaptation of his routine to Minitab.

Index

A

ACF.....	49
ARIMA.....	54
ARIMA (Stat menu)	54
Autocorrelation.....	49
Autocorrelation (Stat menu).....	49

C

Calendar units.....	12
CCF.....	53
Clock units	12
Cross Correlation.....	53
Cross Correlation (Stat menu).....	53

D

Date/time data	
Time series plots	8
Date/time stamp.....	13
Decomposition	20
Decomposition (Stat menu).....	20
Differences.....	48
Differences (Stat menu).....	48
Double Exp Smoothing	37
Double Exp Smoothing (Stat menu).....	37

E

Exponential growth trend model	16
--------------------------------------	----

F

Forecasting	
Decomposition.....	22
Double exponential smoothing	39
Moving average.....	30
Single exponential smoothing	35
Trend model	16
Winters' method.....	45

L

Lag.....	49
Lag (Stat menu).....	49
LBQ statistic.....	50
Level and trend components	39, 44
Double Exponential Smoothing	39
Winters' Method.....	44
Linear moving averages method.....	30
Linear trend model.....	16
Ljung-Box Q statistic.....	50

M

MAD.....	16
MAPE	16
Measures of accuracy	16
Moving Average.....	29
Moving Average (Stat menu)	29
MSD.....	16

P

PACF.....	52
Partial Autocorrelation	52
Partial Autocorrelation (Stat menu).....	52

Q

Quadratic trend model	16
-----------------------------	----

S

Scale (dialog box option)	
Time	12
S-curve trend model	16
Single Exp Smoothing	33
Single Exp Smoothing (Stat menu).....	33
Smoothed values.....	34
Stamp.....	8, 12
Area graph	13
Time series plot.....	8

T

Time series	5
Time Series Plot	8
Time Series Plot (Graph and Stat menus)	8
Time units	12
Time/Scale.....	12, 13, 14
Time	13
Trend Analysis.....	15
Trend Analysis (Stat menu).....	15

W

Weighted average	17
Weights.....	34, 39, 45
Double Exponential Smoothing.....	39
Single Exponential Smoothing	34
Winters' Method	45
Winters' Method.....	42
Winters' Method (Stat menu)	42