

Time Series Components

David Gerbing
School of Business Administration
Portland State University

January 31, 2016

Contents

1	Time Series Models	1
1.1	The components of a time series	1
1.2	Types of models	4
1.3	R Example	6
1.3.1	The Data	6
1.3.2	Decomposition into Trend and Seasonality	8
1.3.3	Forecasting	10

List of Figures

1	Linear trend component.	2
2	Nonlinear trend component.	2
3	Cyclical component (imposed on the underlying trend).	3
4	Seasonal trend component.	3
5	Data based only on a trend component with accompanying random error.	5
6	Data based on trend and seasonality.	6
7	Data based on underlying trend, a cycle and seasonality.	6
8	Time series plot using <code>lessR</code> function <code>LineChart</code>	8
9	Time series plot using R function <code>plot</code>	8
10	Time series decomposition into trend and seasonal additive components.	10
11	Time series plot and fitted Holt-Winter triple smoothing values.	12
12	Time series plot with forecasted values from the triple exponential smoothing analysis.	13

1 Time Series Models

1.1 The components of a time series

A primary use of time-ordered data is to detect trends and other patterns that occur over time. Managers who become adept at analyzing time-ordered data are able to uncover underlying trends in the presence of random fluctuations. Are gross revenues rising? Is March the lowest month for sales each year? Is customer satisfaction increasing over time? Or, conversely, the interest might focus on establishing the lack of an underlying pattern, the demonstration that the time-ordered data exhibit no pattern whatsoever. When the correctly identified pattern is removed from the data, all that should remain is random, chance variation, a stable process, often called a stationary process in the context of a time series. Several different types of relatively sophisticated statistical analyses can accomplish this search for patterns, but often a simple plot of the data reveals the underlying patterns.

Time series plot

Displays the values of the process output in the order in which the values occurred.

A key to analyzing a time series is to understand the form of any underlying pattern of the data ordered over time. The underlying patterns are presented next. This pattern potentially consists of several different components, all of which combine to yield the observed values of the time series. A time series analysis can isolate each component and quantify the extent to which each component influences the form of the observed data. If one or more individual components of a time series are isolated and identified, a forecast can project the underlying pattern into the future.

The most basic time series component is long-term growth or decline.

Trend component T of a time series Y

Long-term, gradual increase or decrease of the variable Y.

The form of the trend pattern may be linear or non-linear, as illustrated in Figure 1. A linear trend pattern such as in (a) and (b) is not only the simplest, but also the most commonly encountered trend pattern.

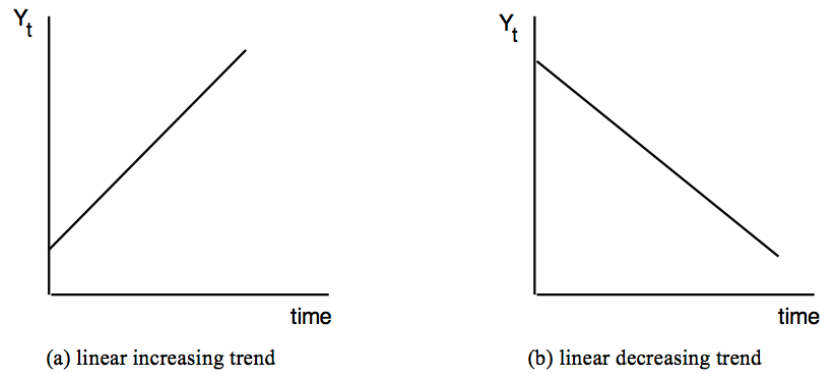


Figure 1: Linear trend component.

Trend can also be non-linear as illustrated in Figure 2.

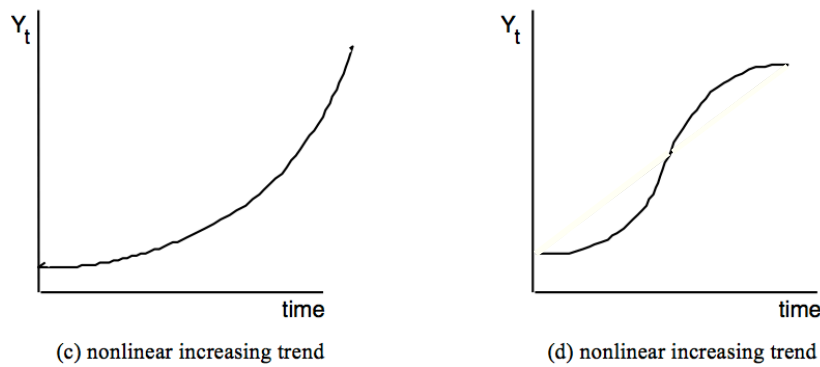


Figure 2: Nonlinear trend component.

The cyclical component modifies the trend component.

Cyclical component of a time series Y

A gradual, long-term, up-and-down potentially irregular swings of the variable Y .

For most business and economic data, the cyclical component is measured in periods of many years, usually decades, and so is usually not present in the typical time series analysis. This component reflects broad swings about either side of the trend line, as illustrated in Figure 3.

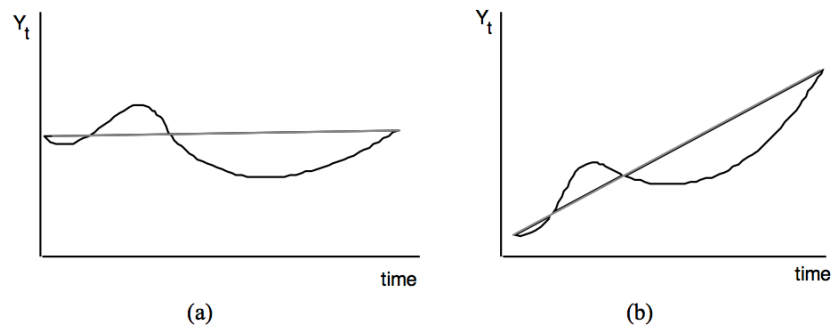


Figure 3: Cyclical component (imposed on the underlying trend).

In addition to the trend and cyclical components, time series data may include a seasonal component.

Seasonal component S of a time series Y .

Regular, relatively short-term repetitive up-and-down fluctuations of the variable Y .

For much business and economic data, the seasonal component is measured in quarters of the year, literally the four seasons of Winter, Spring, Summer, and Fall. Two seasonal components in the absence of cycles are illustrated in Figure 4.




Figure 4: Seasonal trend component.

The trend, cyclical, and seasonal components combine to form the pattern underlying the time series data, the values of Y ordered over time, though not all components characterize every data set. An important purpose of time series analysis is to isolate each of these three components and demonstrate how each affects the value of Y over time, including forecasts for the future. The identification of this pattern, the predictable aspect of a time series, however, is complicated by the presence of random error.

Error component E of a time series Y

A random increase or decrease of the variable Y for a specific time period.

Because all time series include random error, pure trend, cyclical, and seasonal components or combinations of these components are observed together with the error component.

1.2 Types of models

The specific combination of trend, cyclical, or seasonal components that defines the underlying pattern of a time series is observed in the presence of random error. The forecast is based on the underlying pattern delineated from the random error or noise that obscures this pattern. Isolating the pattern is a central goal of a time series analysis.

The stable, predictable component of a time-series is the specific combination of trend, cyclical, and seasonal components that characterize the particular time series. How do these components combine to result in a value of Y_t . One of two models accounts for the underlying pattern, an additive model or a multiplicative model. The additive model expresses Y_t as the sum of the trend, cyclical, seasonal, and error components.

Additive Model of Time Series Data			
$\boxed{Y_t} = \boxed{T_t + C_t + S_t} + \boxed{E_t}$			
data	predictable		random

The cyclical, seasonal, and error components are each a fixed number of units above or below the underlying trend. For example, if the seasonal component at time t is $S_t = -\$12.89$ and there is no cyclical component, then the effect of that season is to lower the value of Y_t the amount of \$12.89 below the affect of the trend component. If the error component is \$4.23 at time t , then the effect of error on Y_t is an increase of \$4.23 above the combined influence of trend and seasonality.

The multiplicative model accounts for the value of Y at time t as a product of the individual components.

Multiplicative Model of Time Series Data			
$\boxed{Y_t} = \boxed{T_t \cdot C_t \cdot S_t} \cdot \boxed{E_t}$			
data	predictable		random

The multiplicative model expresses the C , S , and E components as percentages above or below the underlying trend. Each percentage is a ratio above or below the baseline of 1. The trend component T is expressed in the original units of Y , which is multiplied by each of the C , S , and E ratios to yield the actual data. With the multiplicative model,

larger trends magnify the influence of the remaining components. For the additive model, the influence of a given value of C, S, or E is the same for all values of T.

To illustrate, consider a time series with no cyclical component, C, and a seasonal component for the first third quarter of $S_3 = 1.2$, which means that Y is 20% above the annual average during this season. Assume that the trend at Time 3 is $T_3 = 5$. The stable component of Y_3 from the multiplicative model is the combination of trend and seasonality, or $(5)(1.2) = 6.0$. The next third quarter is at Time 7, so $S_3 = S_7 = 1.2$. What if the trend has increased to $T_7 = 10$? The stable aspect of Y_7 is $(10)(1.2) = 12.0$. The seasonal component of a 20% increase resulted in an increase of trend at Time 3 by 1 unit, and at Time 7 the same seasonal component increased the larger trend by 2 units.

The forecast is based on the identification of the underlying stable components, free of the obscuring error.

Key Concept

The forecast extends into the future the stable pattern of trend, cyclical, and seasonal components underlying the historical data.

For any given time series, each of the three stable components may or may not be present, though typically the cyclical component is not present and usually some trend is exhibited.

Figure 5 presents examples of only a trend component in the presence of random errors, first in the context of little random error (a), and then considerably more random error (b).




Figure 5: Data based only on a trend component with accompanying random error.

A more complex time series involves a combination of trend and seasonality in the presence of random error as shown in Figure 6. Again (a) provides the stable components in the presence of little random error, and then a larger amount imposed over the same stable components in (b).




Figure 6: Data based on trend and seasonality.

The most general type of time series is influenced by all four components, a stable pattern consisting of trend, cyclical, and seasonal components observed in the presence of the random error component, as shown in Figure 7. Again (a) provides the stable components in the presence of little random error, and then a larger amount imposed over the same stable components in (b).

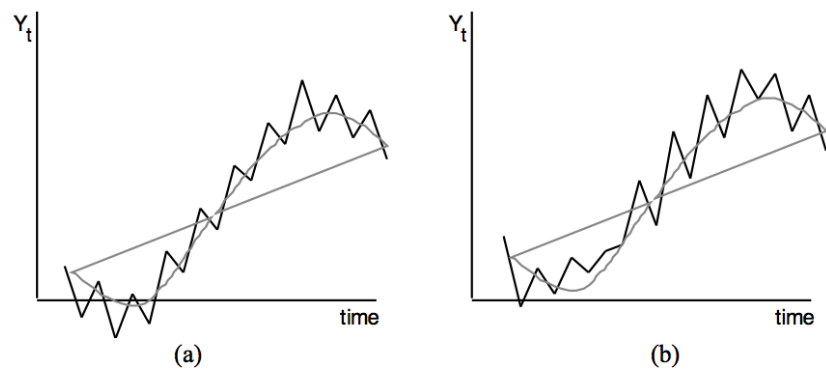


Figure 7: Data based on underlying trend, a cycle and seasonality.

1.3 R Example

A key consideration in time series analysis is how is the stable pattern detected in the presence of the random error, and how are the specific components that contribute to the pattern isolated from each other?

1.3.1 The Data

Consider the following data.

```
> mydata <- Read("http://web.pdx.edu/~gerbing/data/Gujararti/Seasonal10-9.csv")
```


A partial listing of the data follows for the variable FRIG, the quarterly sales of a brand of refrigerators. Other variables in the data table can be ignored for purposes of this example.

```
> mydata
  FRIG
1 1317
2 1615
3 1662
...
31 1764
32 1328
```

To work with a time series in R, convert the data values for variable to an R time series object with the R time series function `ts`. This transformation structures the data according to the specified dates. Even if a second column of the data table included the dates as another variable, the `ts` function would still be invoked because the dates need to be a property of the data values themselves.

The functions introduced here are R functions. Unlike `lessR`, the name of the data table that contains the variable must be explicitly stated, such as with the variable name followed by a \$ sign. For a variable in the data table `mydata`, refer to the variable by its name preceded by: `mydata$`.

Here create a time series of the data with `ts`, store the result in the R object called `F.ts`, and then list the result. The data are quarterly, with the first value from the first quarter of 1978. Specify the `start` value with the year followed by the time period, here Quarter 1. The value of `frequency` in the call to `ts` specifies 4 periods a year.

```
> F.ts <- ts(mydata$FRIG, start=c(1978,1), frequency=4)
```

Creating the time series object, here `F.ts`, does not display any output. As always with R, enter the name of the object to list its values, here to view the imposed time organization of the data.

```
> F.ts

      Qtr1 Qtr2 Qtr3 Qtr4
1978 1317 1615 1662 1295
1979 1271 1555 1639 1238
1980 1277 1258 1417 1185
1981 1196 1410 1417  919
1982  943 1175 1269  973
1983 1102 1344 1641 1225
1984 1429 1699 1749 1117
1985 1242 1684 1764 1328
```

Time series data are typically plotted, here with the `lessR` function `LineChart`, abbreviated `lc`, after doing `set(colors="gray")`.

```
> lc(F.ts)
```




Figure 8: Time series plot using `lessR` function `LineChart`.

The `lessR` function for line plots, `LineChart` can be compared to the corresponding R function `plot`. This original R function is sometimes needed in our analysis of time series.

```
> plot(F.ts)
```




Figure 9: Time series plot using R function `plot`.

1.3.2 Decomposition into Trend and Seasonality

A visual inspection of the plot reveals an apparent seasonality. To quantify this pattern, invoke the R function for time series decompositions, `stl`, an acronym for seasonal and trend decomposition using what is called the loess curve plot for the overall trend. The loess curve estimation procedure does not impose linearity, so it tends to follow the pattern of the data, linear or not.

Save the results in this example in the R object `decomp`.

```
> decomp <- stl(F.ts, s.window="periodic")
```

The contents of `decomp` can be listed.

```
> decomp
```

The listing reveals that affect of seasonality and trend at each time point, separated from the underlying noise, indicated in the column labeled `remainder`. This `remainder` is the obtained modeling error, the variation in the data not accounted for by the model.

```
Components
      seasonal      trend      remainder
1978 Q1 -129.6638 1470.738 -24.073878
1978 Q2  114.5366 1467.280  33.182938
1978 Q3  215.8296 1462.578 -16.407256
1978 Q4 -200.7027 1453.586  42.116792
1979 Q1 -129.6638 1442.039 -41.375378
...
1984 Q4 -200.7027 1438.661 -120.957857
1985 Q1 -129.6638 1442.361 -70.697598
1985 Q2  114.5366 1482.810  86.653061
1985 Q3  215.8296 1525.977  22.193285
1985 Q4 -200.7027 1569.272 -40.568929
```

The seasonal decomposition indicates that the first quarter suppresses refrigerator sales on average by -129.66 . And, for example, the average impact of the third quarter is to increase sales by 215.83 units.

Each set of the three components in the `stl` additive decomposition – seasonal, trend and remainder – together sum to the corresponding data value. To illustrate, consider the first data value, from the first quarter of 1978, $Y_{1978,1} = 1317$.

$$Y_{1978,1} = 1317 = -129.6638 + 1470.738 - 24.073878$$

To further illustrate the meaning of the decomposition, plot the data and the corresponding three components.

```
> plot(decomp)
```




Figure 10: Time series decomposition into trend and seasonal additive components.

The graph provides a convenient visual summary of the estimated components and the error at each time point inconsistent with these components. The impact of the trend component is to first decrease and then increase as time goes on, from 1470.74 units in the first quarter of 1978 down to 1467.28 units for the next quarter, but then rebounding up to 1569.27 units in the fourth quarter of 1985.

A remainder above the mid-line indicates that a data value is larger than that consistent with the stable component of trend and seasonality. A remainder below the line indicates the data value is smaller than the consistent stable component. Here the largest error (remainder) is for 1980 Q2 where the data value is 1258, but the trend and seasonal components are 1279.5 and 114.5, which sum to 1394. The data value is 136.1 units less than the value consistent with the stable trend and seasonal components.

1.3.3 Forecasting

The `stl` function estimates trend and seasonality for each time point in the data, but does not provide any forecasts for the data value at future time points. Forecasts can be obtained from several different time series analyses. Presented here is what is called exponential smoothing with trend and seasonal components, or triple exponential smoothing. This method is often referred to as Holt-Winters after the names of its authors.

Exponential smoothing means that the model expresses the current value as a weighted

average of all past observations with exponentially decreasing weights over time. The triple smoothing method also accounts for any trend and seasonality.

The corresponding R function is `HoltWinters`, which analyzes an R time series data object, here the previously created `F.ts`. The analysis can apply an additive or multiplicative model, with the additive model as the default, and can also do exponential smoothing with or without trend and seasonality, with the default of full triple smoothing.

Here save the results of the analysis in the R object called `hw`.

```
> hw <- HoltWinters(F.ts)
```

The contents of `hw` can be listed.

```
> hw
```

Part of the output of `HoltWinters` is the estimate of the trend influence and the seasonal influences, here quarterly. This data is not linear, so the `a` and `b` estimated values of the linear trend are ignored. The Holt-Winters estimate of the seasonal influences is considerably different from the averaging used in the simpler `stl` decomposition. The values are similar but do not agree completely.

```
a 1504.4125
b -11.1250
s1 -130.8680
s2 120.0038
s3 210.8371
s4 -174.5925
```

Applying the R `plot` function to a Holt-Winters object, here `hw`, by default provides both a plot of the data and of the fitted values provided by the analysis, shown in Figure 11.

```
plot(hw)
```




Figure 11: Time series plot and fitted Holt-Winter triple smoothing values.

The plot of the Holt-Winters object in Figure 11 nicely illustrates the match of the fitted values from the model to the data, but does not provide a forecast of future values. To obtain the forecast invoke the R function `predict`, here applied to the previously calculated object `hw`. Specify the number of time periods to project into the future with the `n.ahead` parameter. Here 12 periods is 3 years based on quarterly data. Store the results in the `hw.p` object.

```
> hw.p <- predict(hw, n.ahead=12)
```

Obtain the output by listing the contents of the `hw.p` object.

```
> hw.p
```

The forecasted values are organized by year and then quarter.

```
      Qtr1    Qtr2    Qtr3    Qtr4
1986 1362.419 1602.166 1681.875 1285.320
1987 1317.919 1557.666 1637.375 1240.820
1988 1273.419 1513.166 1592.875 1196.320
```

These forecasted values can also be plotted on the same graph as the data, but as extension into the future. The data time series `F.ts` goes until 1986, so to obtain a graph of 12 quarters ahead, specify the plot out 3 years to 1989, beginning with the first data value from 1978.

```
> plot(F.ts, xlim=c(1978,1989))
```

Enhance the obtained plot with the forecasted values, here in red with the R function `lines`.




Figure 12: Time series plot with forecasted values from the triple exponential smoothing analysis.

```
> lines(hw.p, col="red")
```

The graph provides the visual extension of the data. The output of the forecasted values provides the exact forecasts. The forecast is only accurate to the extent that the non-linear trend will repeat in the future.