1. INTRODUCTION TO TIME SERIES ANALYSIS

A time series is any series of data that varies over time e.g. the quarterly sales for a company during the past five years or daily returns of a security.

Time-series models are used to:

1. explain the past
2. predict the future of a time-series

2. CHALLENGES OF WORKING WITH TIMES SERIES

When assumptions of the regression model are not met, we need to transform the time series or modify the specifications of the regression model.

Problems in time series:

1. When the dependent and independent variables are distinct, presence of serial correlation of the errors does not affect the consistency of estimates of intercept or slope coefficients.

But in an autoregressive time-series regression, presence of serial correlation in the error term makes estimates of the intercept (b_0) and slope coefficient (b_1) to be inconsistent.

2. When mean and/or variance of the time series model change over time and is not constant, then using an autoregressive model will provide invalid regression results.

Because of these problems in time series, time series model is needed to be transformed for the purpose of forecasting.

3. TREND MODELS

3.1 Linear Trend Models

In a linear trend model, the dependent variable changes at a ***constant rate*** with time.

$$y_t = b_0 + b_1 t + \varepsilon_t$$

where,

- y_t = *value of time series at time t (value of dependent variable)*
- b_0 = *y-intercept term*
- b_1 = *slope coefficient or trend coefficient*
- t = *time, the independent or explanatory variable*
- ε_t = *random error term*

The predicted or fitted value of y_t in period 1 is:

$$\hat{y}_1 = \hat{b}_0 + \hat{b}_1(1)$$

The predicted or fitted value of y_t in period 5 is:

$$\hat{y}_5 = \hat{b}_0 + \hat{b}_1(5)$$

The predicted or fitted value of y_t in period T + 1 is:

$$\hat{y}_{T+1} = \hat{b}_0 + \hat{b}_1(T+1)$$

NOTE:

Each consecutive observation in the time series increases by $\hat{b}_1$ in a linear trend model.

Practice: Example 1
Volume 1, Reading 11.

3.2 Log-Linear Trend Models

When time series has exponential growth rate, it is more appropriate to use log-linear trend model instead of linear trend model. Exponential growth rate refers to a ***constant growth*** at a particular rate.

$$y_t = e^{b_0 + b_1 t}$$

where,

t = 1, 2, 3,…,T

Taking natural log on both sides we have:

$\ln y_t = b_0 + b_1 t + \varepsilon_t$

where,

t = 1, 2, 3, …,T

Linear trend model	Log-linear trend model
Predicted trend value of y_t is $\hat{b}_0 + \hat{b}_1 t$,	Predicted trend value of y_t is $e^{\hat{b}_0 + \hat{b}_1 t}$ because $e^{\ln y_t} = y_t$.
The model predicts that y_t grows by a ***constant amount*** from one period to the next.	The model predicts a ***constant growth rate*** in y_t of $e^{b_1} - 1$.
A linear trend model is appropriate to use when the residuals from a model are equally distributed above and below the regression line e.g. inflation rate.	A log-linear model is appropriate to use when the residuals of the model exhibit a persistent trend i.e. either positive or negative for a period of time e.g. financial data i.e. stock prices, sales, and stock indices.

Practice: Example 2 & 3 Volume 1, Reading 9.

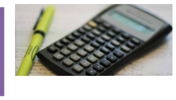

Limitation of Trend Models: Trend model is based on only one independent variable i.e. time; therefore, it does not adequately incorporate the underlying dynamics of the model.

3.3 Trend Models and Testing for Correlated Errors

In case of presence of serial correlation, both the linear trend model and the log-linear trend model are not appropriate to use. In case of serial correlation, autoregressive time series models represent better forecasting models.

4. AUTOREGRESSIVE (AR) TIME-SERIES MODELS

An autoregressive (AR) model is a time series regression in which the independent variable is a lagged (past) value of the dependent variable i.e.

$$x_t = b_0 + b_1 x_{t-1} + \varepsilon_t$$

First order autoregressive AR (1) for the variable x_t is:

$$x_t = b_0 + b_1 x_{t-1} + \varepsilon_t$$

A pth-order autoregressive AR (p) for the variable x_t is:

$$x_t = b_0 + b_1 x_{t-1} + b_2 x_{t-2} + \ldots + b_p x_{t-p} + \varepsilon_t$$

4.1 Covariance-Stationary Series

In order to obtain a valid statistical inference from a time-series analysis, the time series must be covariance stationary.

Time series is covariance stationary when:

1. The expected value of the time series is constant and finite in all periods.
2. The variance of the time series is constant and finite in all periods.
3. The covariance of the time series with past or future values of itself is constant and finite in all periods.

NOTE:

Weakly stationary also refers to covariance stationary.

Stationary Data: When a time series variable does not exhibit any significant upward or downward trend over time.

Nonstationary Data: When a time series variable exhibits a significant upward or downward trend over time.

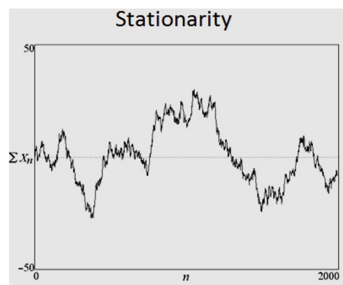

Non-stationarity (upward trend)

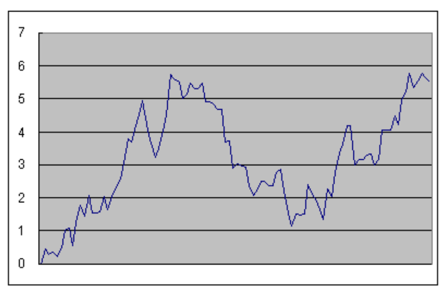

Consequence of Covariance Non-Stationarity: When time series is not covariance stationary, the regression estimation results are invalid because:

- The "*t*-ratios" will not follow a *t*-distribution.
- The estimate of b_1 will be biased and any hypothesis tests will be invalid.

NOTE:

Stationarity in the past does not guarantee stationarity in the future because state of the world may change over time.

4.2 Detecting Serially Correlated Errors in an Autoregressive Model

An Autoregressive model can be estimated using ordinary least squares model (OLS) when the time series is covariance stationary and the errors are uncorrelated.

Detecting Serial Correlation in AR models: In AR models, Durbin-Watson statistic cannot be used to test serial correlation in errors. In such cases, t-test is used.

The autocorrelations of time series refer to the correlations of that series with its own past values.

- When autocorrelations of the error term are zero, the model can be specified correctly.
- When autocorrelations of the error term are significantly different from zero, the model cannot be specified correctly.

Example:

Suppose a sample has 59 observations and one independent variable. Then,

S.D = 1 / $\sqrt{T}$ = 1 / $\sqrt{59}$ = 0.1302
Critical value of t (at 5% significant level with df = 59 – 2 = 57) is 2.

Suppose autocorrelations of the Residual are as follows:

Lag	Autocorrelation	Standard Error	t-statistic*
1	0.0677	0.1302	0.5197
2	-0.1929	0.1302	-1.4814
3	0.0541	0.1302	0.4152
4	-0.1498	0.1302	-1.1507

* **t-statistic = Autocorrelations / Standard Error**

It can be seen from the table that none of the first four autocorrelations has t-statistic > 2 in absolute value.

Conclusion: None of these autocorrelations differ significantly from 0 thus, residuals are not serially correlated and model is specified correctly and OLS can be used to estimate the parameters and the standard errors of the parameters in the autoregressive model.

Correcting Serial Correlation in AR models: The serial correlation among the residuals in AR models can be removed by estimating an autoregressive model by adding more lags of the dependent variable as explanatory variables.

4.3 Mean Reversion

A time series shows mean reversion if it tends to move towards its mean i.e. decrease when its current value is above its mean and increase when its current value is below its mean.

- When a time series equals its mean-reverting value, then the model predicts that the value of the time series will be the same in the next period i.e. $x_{t+1} = x_t$.

Mean reverting level of $x_t = \dfrac{b_0}{1-b_1}$

- Time series will remain the same if its current value $= \dfrac{b_0}{1-b_1}$
- Time series will Increase if its current value $< \dfrac{b_0}{1-b_1}$
- Time series will Decrease if its current value $> \dfrac{b_0}{1-b_1}$

4.4 Multiperiod Forecasts and the Chain Rule of Forecasting

The **chain rule of forecasting** is a process in which a predicted value two periods ahead is estimated by first predicting the next period's value and substituting it into the equation of a predicted value two periods ahead i.e.

The one-period ahead forecast of x_t from an AR (1) model is as follows:

$$\hat{x}_{t+1} = \hat{b}_0 + \hat{b}_1 x_t$$

Two-period ahead forecast is:

$$\hat{x}_{t+2} = \hat{b}_0 + \hat{b}_1 x_{t+1}$$

NOTE:

Multiperiod forecast is more uncertain than single-period forecast because the uncertainty increases when number of periods in the forecast increase.

Example:

The one-period ahead forecast of x_t from an AR (1) model when x_t = 0.65 is as follows:

$$\hat{x}_{t+1} = 0.0834 + 0.8665(0.65) = 0.6466$$

Two-period ahead forecast is:

$$\hat{x}_{t+2} = 0.0834 + 0.8665(0.6466) = 0.6437$$

Practice: Example 6
Volume 1, Reading 9.

4.5 Comparing Forecast Model Performance

The accuracy of the model depends on its forecast error variance.

- *The smaller the forecast error variance, the more accurate the model will be.*

In-sample forecast errors: These are the residuals from the fitted time series model i.e. residuals within a sample period.

Out-of-sample forecast errors: These are the residuals outside the sample period. It is more important to have smaller forecast error variance (i.e. high accuracy) for out-of-sample forecasts because the future is always out of sample.

To evaluate the out-of-sample forecasting accuracy of the model, ***Root mean squared error*** (RMSE) is used. RMSE is the square root of the average squared error.

Decision Rule: *The smaller the RMSE, the more accurate the model will be.*

The **RMSE (Root Mean Squared Error)** is used as a criterion for comparing forecasting performance of different forecasting models. To accurately evaluate uncertainty of forecast, both the uncertainty related to the error term and the uncertainty related to the estimated parameters in the time-series model must be considered.

NOTE:

If the model has the lowest RMSE for in-sample data, it does not guarantee that the model will have the lowest RMSE for out-of-sample data as well.

4.6 Instability of Regression Coefficients

When the estimated regression coefficients in one period are quite different from those estimated during another period, this problem is known as instability or nonstationarity.

The estimates of regression coefficients of the time-series model can be different across different sample periods i.e. the estimates of regression coefficients using shorter sample period will be different from using longer sample periods. Thus, sample period selection is one of the important decisions in time series regression analysis.

- Using longer time periods increase statistical reliability but estimates are not stable.
- Using shorter time periods increase stability of the estimates but statistical reliability is decreased.

NOTE:

We cannot select the correct sample period for the regression analysis by simply analyzing the autocorrelations of the residuals from a time-series model. In order to select the correct sample, it is necessary that data should be Covariance Stationary.

5. RANDOM WALKS AND UNIT ROOTS

5.1 Random Walks

A. Random walk without drift: In a random walk without drift, the value of the dependent variable in one period is equal to the value of the series in the previous period plus an unpredictable random error.

$$x_t = x_{t-1} + \varepsilon_t$$

where,

$b_0 = 0$ and $b_1 = 1$.

In other words, the best predictor of the time series in the next period is its current value plus an error term.

The following conditions must hold:

1. Error term has an expected value of zero.
2. Error term has a constant variance.
3. Error term is uncorrelated with previous error terms.

- The equation of a random walk represents a special case of an AR (1) model with $b_0 = 0$ and $b_1 = 1$.

- AR (1) model cannot be used for time series with random walk because random walk has no finite mean, variance and covariance. In random walk $b_0 = 0$ and $b_1 = 1$, so $\frac{b_0}{1-b_1} = 0 / 0 =$ undefined mean reverting level.
- A standard regression analysis cannot be used for a time series that is random walk.

Correcting Random Walk: When time series has a random walk, it must be converted to covariance-stationary time series by taking the first difference between x_t and x_{t-1} i.e. equation becomes:

$$y_t = x_t - x_{t-1} = \varepsilon_t$$

- Thus, best forecast of y_t made in period t-1 is 0. This implies that the best forecast is that the value of the current time series x_{t-1} will not change in future.

After taking the first difference, the first differential variable y_t becomes covariance stationary. It has $b_0 = 0$ and $b_1 = 0$ and mean reverting level = 0/1 = 0.

- The first differential variable y_t can now be modeled using linear regression.
- However, modeling the first differential variable y_t with an AR (1) model is not helpful to predict the future because $b_0 = 0$ and $b_1 = 0$.

Consequences of Random Walk: When the model has random walk, its R^2 will be significantly high and at the same time changes in dependent variable are unpredictable. In other words, the statistical results of the regression will be invalid.

B. Random walk with a drift: In a random walk with a drift, dependent variable increases or decreases by a constant amount in each period.

$$x_t = b_0 + x_{t-1} + \varepsilon_t$$

where,

$b_0 \neq 0$ and $b_1 = 1$.

By taking first difference,

$$y_t = x_t - x_{t-1} = b_0 + \varepsilon_t$$

NOTE:

All random walks (with & without a drift) have unit roots.

5.2 The Unit Root Test of Nonstationarity

AR (1) time series model will be covariance stationary only when the absolute value of the lag coefficients $b_1<1$. (Note that when b_1 is > 1 in absolute value, it is known as explosive root).

Detecting Random Walk: When time series has random walk, the series does not follow t-distribution and t-test will be invalid. Therefore, t-statistic cannot be used to test the presence of random walk because standard errors in an AR model are invalid if the model has a random walk. Thus, **Dickey-Fuller** test is used to detect nonstationarity:

Method 1: Examining Autocorrelations of the AR model

Stationary Time Series:

- Autocorrelations at all lags equals to zero, or
- Autocorrelations decrease rapidly to zero as the number of lags increases in the model.

Nonstationary time series:

- Autocorrelations at all lags are not equal to zero, or
- Autocorrelations do not decrease rapidly to zero as the number of lags increases in the model.

Method 2: Using Dickey-Fuller Test

Subtracting x_{t-1} from both sides of AR (1) equation we have:

$$x_t - x_{t-1} = b_0 + (b_1 - 1)\, x_{t-1} + \varepsilon_t$$

(or)

$$x_t - x_{t-1} = b_0 + g_1 x_{t-1} + \varepsilon_t$$

where,

$g_1 = (b_1 - 1)$.

- If $b_1 = 1$, then $g_1 = 0$. This implies that there is a unit root in AR (1) model.

Null Hypothesis: H_0: $g_1 = 0 \rightarrow$ time series has a unit root and is Nonstationary

Alternative Hypothesis: H_1: $g_1 < 0 \rightarrow$ time series does not have a unit root and is Stationary

- t-statistic is calculated for predicted value of g_1 and critical values of t-test are computed from Dickey-Fuller test table (these critical t-values in absolute value > than typical critical t-values).

Practice: Example 12
Volume 1, Reading 11.

6. MOVING-AVERAGE TIME SERIES MODELS

Moving average (MA) is different from AR model. MA is an average of successive observations in a time series. It has lagged values of residuals instead of lagged values of dependent variable.

6.1 Smoothing Past Values with an n-Period Moving Average

n-period moving average is used to smooth out the fluctuations in the value of a time series across different time periods.

$$\frac{x_t + x_{t-1} + x_{t-2} + \ldots.. + x_{t-(n-1)}}{n}$$

Drawbacks of Moving Average:

- It is biased towards large movements in the actual data.
- It is not the best predictor of the future.
- It gives equal weights to all the periods in the moving average.

Distinguishing AR time series from a MA time series:

- Autocorrelations of most AR (p) time series start large and decline gradually.
- Autocorrelations of MA (q) time series suddenly drop to 0 after the first q autocorrelations.

7. SEASONALITY IN TIME-SERIES MODELS

When a time series variable exhibit a repeating patterns at regular intervals over time, it is known as seasonality e.g. sales in Dec. > sales in Jan. A time series with seasonality also has a non-constant mean and thus is not covariance stationary.

Detecting seasonality: In case of seasonality in the data, autocorrelation in the model differs by season. For example, in case of quarterly sales data of a company, if the fourth autocorrelation of the error term differs significantly from 0 → it is a sign of seasonality in the model.

Decision Rule: When t-statistic of the fourth lag of autocorrelations of the error > critical t-value → reject null hypothesis that fourth autocorrelations is 0. Thus, there is seasonality problem.

Correcting Seasonality: This problem can be solved by adding seasonal lags in an AR model i.e. after including a seasonal lag in case of quarterly sales data, the AR model becomes:

$$x_t = b_0 + b_1 x_{t-1} + b_2 x_{t-4} + \varepsilon_t$$

In case of monthly sales data, the AR model becomes:

$$x_t = b_0 + b_1 x_{t-1} + b_2 x_{t-12} + \varepsilon_t$$

NOTE:

R^2 of the model without seasonal lag will be less than the R^2 of the model with seasonal lag. This implies that when time series exhibit seasonality, including a seasonal lag in the model improves the accuracy of the model.

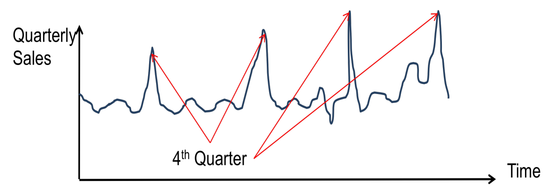

Practice: Example 15
Volume 1, Reading 9.

8. AUTOREGRESSIVE MOVING-AVERAGE MODELS (ARMA)

An ARMA model combines both autoregressive lags of the dependent variable and moving-average errors.

Drawbacks of ARMA model:

- Parameters of ARMA models are usually very unstable.
- ARMA models depend on the sample used.
- Choosing the right ARMA model is a difficult task because it is more of an art than a science.

9. AUTOREGRESSIVE CONDITIONAL HETEROSKEDASTICITY MODELS (ARCH)

When regression model has (conditional) heteroskedasticity i.e. variance of the error in a particular time-series model in one period depends on the variance of the error in previous periods, standard errors of the regression coefficients in AR, MA or ARMA models will be incorrect, and hypothesis tests would be invalid.

ARCH model:

ARCH model must be used to test the existence of conditional heteroskedasticity. An ARCH (1) time series is the one in which the variance of the error in one period depends on size of the squared error in the previous period i.e. if a large error occurs in one period, the variance of the error in the next period will be even larger.

To test whether time series is ARCH (1), the squared residuals from a previously estimated time-series model are regressed on the constant and first lag of the squared residuals i.e.

$$\hat{\varepsilon}_t = \alpha_0 + \alpha_1 \hat{\varepsilon}_{t-1}^2 + \mu_t$$

where,

μ_t is an error term

Decision Rule: If the estimate of α_1 is statistically significantly different from zero, the time series is ARCH (1). If a time-series model has ARCH (1) errors, then the variance of the errors in period t+1 can be predicted in period t using the formula:

$$\hat{\sigma}_{t+1}^2 = \hat{\alpha}_0 + \alpha_1 \hat{\varepsilon}_t^2$$

Consequences of ARCH:

- Standard errors for the regression parameters will not be correct.
- When ARCH exists, we can predict the variance of the error terms.

Generalized least squares or other methods that correct for heteroskedasticity must be used to estimate the correct standard error of the parameters in the time-series model.

Autoregressive model versus ARCH model:

- Using AR (1) model implies that model is correctly specified.
- Using ARCH (1) implies that model can not be correctly specified due to existence of conditional heteroskedasticity in the residuals; therefore, ARCH (1) model is used to forecast variance/volatility of residuals.

10. REGRESSIONS WITH MORE THAN ONE TIME SERIES

1. When neither of the time series (dependent & independent) has a unit root, linear regression can be used.
2. One of the two time series (i.e. either dependent or independent but not both) has a unit root, we should not use linear regression because error term in the regression would not be covariance stationary.
3. If both time series have a unit root, and the time series are not cointegrated, we cannot use linear regression.
4. If both time series have a unit root, and the time series is cointegrated, linear regression can be used. Because, when two time series are cointegrated, the error term of the regression is covariance stationary and the t-tests are reliable.

Cointegration: Two time series are cointegrated if

- A long term financial or economic relationship exists between them.
- They share a common trend i.e. two or more variables move together through time.

Two Cointegrated Time Series

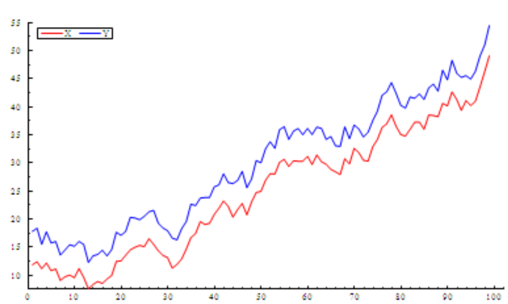

NOTE:

Cointegrated regression estimates the long-term relation between the two series. Therefore, it is not the best model of the short-term relation between the two series.

Detecting Cointegration: The Engle-Granger Dickey-Fuller test can be used to determine if time series are cointegrated.

Engle and Granger Test:

1. Estimate the regression $y_t = b_0 + b_1 x_t + \varepsilon_t$
2. Unit root in the error term is tested using Dickey-fuller test but the critical values of the Engle-Granger are used.
3. If test fails to reject the null hypothesis that the error term has a unit root, then error term in the regression is not covariance stationary. This implies that two time series are not cointegrated and regression relation is spurious.
4. If test rejects the null hypothesis that the error term has a unit root, then error term in the regression is covariance stationary. This implies that two time series are cointegrated and regression results and parameters will be consistent.

NOTE:

- When the first difference is stationary, series has a single unit root. When further differences are required to make series stationary, series is referred to have multiple unit roots.
- For multiple regression model, rules and procedures for unit root and stationarity are the same as that of single regression.

12. SUGGESTED STEPS IN TIME-SERIES FORECASTING

Following is a guideline to determine an accurate model to predict a time series.

1. Select the model on the basis of objective i.e. if the objective is to predict the future behavior of a variable based on the past behavior of the same variable, use Time series model and if the objective is to predict the future behavior of a variable based on assumed casual relationship with other variables Cross sectional model should be used.

2. When time-series model is used, plot the series to detect Covariance Stationarity in the data. Trends in the time series data include:

- A linear trend
- An exponential trend
- Seasonality
- Structural change i.e. a significant shift in mean or variance of the time series during the sample period

3. When there is no seasonality or structural change found in the data, linear trend or exponential trend is appropriate to use i.e.

i. Use linear trend model when the data plot on a straight line with an upward or downward slope.
ii. Use log-linear trend model when the plot of the data exhibits a curve.
iii. Estimate the regression model.
iv. Compute the residuals
v. Use Durbin-Watson statistic to test serial correlation in the residual.

4. When serial correlation is detected in the model, AR model should be used. However, before using AR model, time series must be tested for Covariance Stationarity.

- If time series has a linear trend and covariance nonstationary; it can be transformed into covariance stationary by taking the first difference of the data.
- If time series has exponential trend and covariance nonstationary; it can be transformed into covariance stationary by taking natural log of the time series and then taking the first difference.
- If the time series exhibits structural change, two different time-series model (i.e. before & after the shift) must be estimated.
- When time series exhibits seasonality, seasonal lags must be included in the AR model.

5. When time series is converted into Covariance Stationarity, AR model can be used i.e.

- Estimate AR (1) model;
- Test serial correlation in the regression errors; if no serial correlation is found only then AR (1) model can be used. When serial correlation is detected in AR (1), then AR (2) should be used and tested for serial correlation. When no serial correlation is found, AR (2) can be used. If serial correlation is still present, order of AR is gradually increasing until all serial correlation is removed.

6. Plot the data and detect any seasonality. When seasonality is present, add seasonal lags in the model.

7. Test the presence of autoregressive conditional heteroskedasticity in the residuals of the model i.e. by using ARCH (1) model.

8. In order to determine the better forecasting model, calculate out-of-sample RMSE of each model and select the model with the lowest out-of-sample RMSE.

Practice: End of Chapter Practice Problems for Reading 9 & FinQuiz Item-set ID# 11585.