Testing for Stationarity

Qiao Zhou

qiao_zhou@mfe.berkeley.edu

2016-04-13

Test for Stationarity

Stationary time series have many desired properties. If one regresses one nonstationary time series on its lagged series or another nonstationary time series, the regression results will be spurious. On such situations, no matter how significant the test result is, the regression will not make much sense. To test for stationarity of a time series, Dickey Fuller Test or Augmented Dickey Fuller Test can be used.

OLS for non-stationary time series

OLS is well-behaved as long as time series is stationary and ergodic, but all properties break down for nonstationary time series. For an $AR(1)$ process ,

$$x_t = \phi x_{t-1} + e_t$$

When $\phi < 1$, we have

$$\sqrt{T}\left(\hat{\phi} - \phi\right) \xrightarrow{d} N(0, Q_{xx}^{-1} \Omega Q_{xx}^{-1})$$

Where $\Omega = \mathrm{E}\left(x_i x_i' e_i^2\right)$. Therefore, the distribution of $\hat{\phi}$ will look normal when sample size T becomes sufficiently large.

However, when $\phi = 1$, the model becomes essentially a random walk model, and $\sqrt{T}\left(\hat{\phi} - \phi\right)$ no longer converges to a well-defined distribution. In other word, $\sqrt{T}\left(\hat{\phi} - 1\right)$ is not well defined if $\phi = 1$. $T\left(\hat{\phi} - 1\right)$ will converge to a well-defined distribution at a rate proportional to the sample size T (a.k.a. superconsistent). Since it converges at a rate of T, its distribution looks similar regardless of sample size. Also, the larger the ϕ, the faster the convergence. The distribution of $T\left(\hat{\phi} - \phi\right)$ when $\phi = 1$ is called the Dickey Fuller distribution, which will be covered in the next section.

Dickey Fuller Test

For an $AR(1)$ process,

$$x_t = \phi x_{t-1} + e_t$$

$$x_t - x_{t-1} = \phi x_{t-1} - x_{t-1} + e_t$$

$$\Delta x_t = (\phi - 1) x_{t-1} + e_t$$

$$= \rho x_{t-1} + e_t$$

Dickey Fuller Test (DF test) test for the null hypothesis that the time series x_t has a unit root, which is equivalent to test for that $\phi = 1$ or $\rho = \ \phi - 1 = 0$:

$$H_0 : \rho = 0$$

Where the DF statistic is

$$DF = \frac{\hat{\rho}}{s.e.(\hat{\rho})}$$

Usually, the t-test has a normal distribution in the limiting case, but here $T(\hat{\phi} - 1)$ converges to a well-defined distribution that is not normal but is skewed to the left. This distribution is called the Dickey Fuller distribution. The standard t-tests of $H_0 : \rho = 0$ tends to reject too soon.

In practice, the preferred test for stationarity is the Augmented Dickey Fuller Test (ADF Test), which has better small sample properties. The ADF test uses the null hypothesis that $H_0 : \rho = 0$, where the difference equation used is similar to DF but with lagged difference terms up to the pth order augmented:

$$\Delta x_t = \rho x_{t-1} + \delta_1 \Delta x_{t-1} + \delta_2 \Delta x_{t-2} + \cdots + \delta_p \Delta x_{t-p} + e_t$$

Where the DF statistic is

$$DF = \frac{\hat{\rho}}{s.e.(\hat{\rho})}$$

For example, we take the 10-Year Treasury Constant Maturity yield data. It visually appears nonstationary.

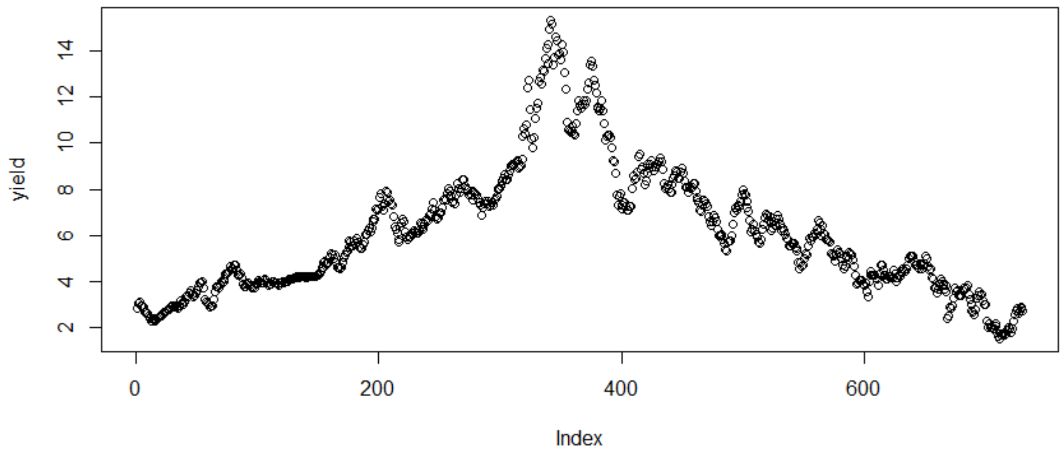

To confirm our suspicion, we perform the standard ADF test using multiple lags. In summary, we listed the DF t-statistics, p-values and their corresponding lags:

```
        Lags DF t-Statistics    P-value
0   0.0000000      -1.3199406  0.8661717
1   1.0000000      -1.7793070  0.6717261
2   2.0000000      -1.4226577  0.8226925
3   3.0000000      -1.5751185  0.7581572
4   4.0000000      -1.5049036  0.7878785
```

It is clear that the data series is not stationary, since the P-values are uniformly high across all lags. Next, we apply the differencing to the data series and plot the series again. The plot appears to be stationary by inspection.

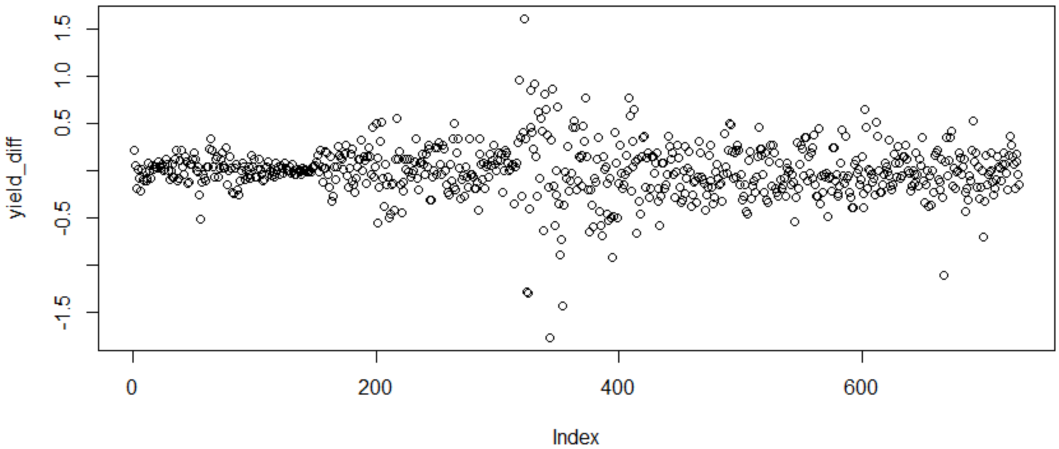

Now, we perform ADF test again, and we find that the P-values are close to zero across all lags and the t-statistics are very negative compared with the critical values. Recall that at 1% significance level, the critical value is around -3.44 only. In this case, we can reject the unit root null in favor of the alternative that the series is stationary.

```
       Lags DF t-Statistics   P-value
0    0.00000        -19.67363   0.01000
1    1.00000        -19.83454   0.01000
2    2.00000        -14.74380   0.01000
3    3.00000        -13.41081   0.01000
4    4.00000        -11.09539   0.01000
```

Conclusion

Many financial time series are persistent and nonstationary in nature. Therefore, it’s better to test for stationarity before performing further analysis. Price series, yields, inflation, etc., are usually nonstationary. To achieve stationarity, differencing can be used. Cointegration may also be used to produce stationary time series.

Appendix: R code

```
data_file = "C:/dev/empirical_methods/data/Tbill10yr.csv"
df = read.csv(file = data_file,header = TRUE)
yield = df$VALUE
yield_diff = diff(yield)

lags = 0:4
yield_dfs= rep.int(0, length(lags))
yield_pvals= rep.int(0, length(lags))

yield_diff_dfs= rep.int(0, length(lags))
yield_diff_pvals= rep.int(0, length(lags))

library(tseries)

for(i in 1:length(lags)){
  m_lag = lags[i];
  f = adf.test(yield,k =m_lag)
  yield_dfs[i] = f$statistic
  yield_pvals[i] = f$p.value

  f = adf.test(yield_diff,k =m_lag)
  yield_diff_dfs[i] = f$statistic
  yield_diff_pvals[i] = f$p.value
}


cbind(lags,yield_dfs,yield_pvals)
cbind(lags,yield_diff_dfs,yield_diff_pvals)


trial <- cbind(lags,yield_dfs,yield_pvals)
colnames(trial) <- c('Lags','DF t-Statistics', 'P-value')
rownames(trial) <- paste(lags)
tab_yield <- as.table(trial)
print(tab_yield)

trial <- cbind(lags,yield_diff_dfs,yield_diff_pvals)
colnames(trial) <- c('Lags','DF t-Statistics', 'P-value')
rownames(trial) <- paste(lags)
tab_yield_diff <- as.table(trial)
print(tab_yield_diff)
```