

# Econometrics 2b: Handout #6

Matthew J. Lindquist\*

April 2009

## 1 Tests for Stationarity

### 1.1 The Dickey-Fuller Test

- Model  $Y_t$  as an  $AR(1)$  process

$$Y_t = \theta Y_{t-1} + \varepsilon_t$$

$$\varepsilon_t \sim IID(0, \sigma^2).$$

If  $|\theta| < 1$ , then the process is stationary. If  $|\theta| > 1$ , then the process explodes.  
If  $|\theta| = 1$ , then the process exhibits a unit root and the process is nonstationary.

- The unit root test examines whether or not  $\theta = 1$ .
- In practice, we first make the following transformation

$$\begin{aligned} Y_t - Y_{t-1} &= \theta Y_{t-1} - Y_{t-1} + \varepsilon_t \rightarrow \\ \Delta Y_t &= (\theta - 1) Y_{t-1} + \varepsilon_t \\ &= \pi Y_{t-1} + \varepsilon_t \end{aligned}$$

and then test whether or not  $\pi = 0$ .

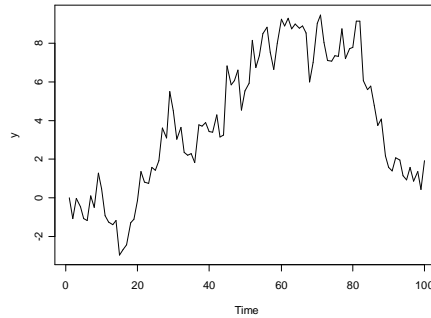
- If  $\pi = 0$ , then  $\theta = 1$  and  $Y_t$  has a unit root.
- $H_0 : \pi = 0$
- $H_1 : \pi < 0$ .
- Under  $H_0$  the  $t$ -value of the estimated coefficient of  $Y_{t-1}$  does NOT follow the standard  $t$ -distribution. You must compare them with the DF-test statistics like those reported in Table A on p. 439 in Enders (2004) [look at the  $\tau$ -statistic].

---

\*Department of Economics, Stockholm University, SE-106 91 Stockholm, Sweden. Tel. 46+8+163831. E-mail: Matthew.Lindquist@ne.su.se.

**Example 1** *Determining the order of integration - the DF-test.*

**step 1:** Graph the time series,  $y$ : `plot.ts(y)`



Does this time series look stationary?

**step 2:** Implement the DF-test in  $\mathbb{R}$  by writing:

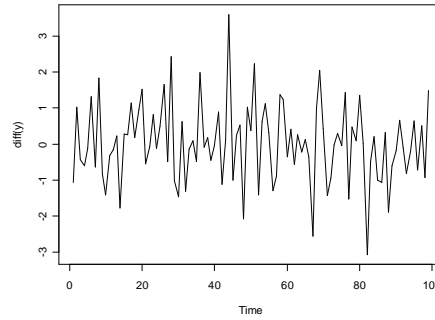
```
dfuller.reg <- lm(diff(y) ~ lag(y,-1)[1:99] -1)
summary(dfuller.reg)
```

Examine the regression summary. The  $t$ -statistic associated with the coefficient in front of the lagged value of  $y$  is equal to -0.997. Compare this to the appropriate (non-standard) critical values from Table A in Enders (2004). The 1% and 5% critical values of  $\tau$  (the  $t$ -statistic in a regression with no constant and no deterministic trend) for a sample size = 100 are -2.60 and -1.95, respectively. Since the  $t$ -value from the regression is **greater than** both of these critical values (i.e., closer to zero), we can not reject that  $\pi = 0$  and that there is a unit root present in this series.

**Conclusion:** Do NOT reject  $y$  is  $I(1)$ .

**step 3:** Examine whether  $y$  is  $I(1)$  or  $I(2)$ .

**step 4:** Graph the time series  $\Delta y$ : `plot.ts(diff(y))`



Does this series look stationary? What is our definition of stationarity? What should I be looking for?

**step 5:** Implement the DF-test in  $\mathbb{R}$  by writing:

```
dfuller.reg <- lm(diff(y,1,2) ~ lag(diff(y),-1)[1:98] -1)
summary(dfuller.reg)
```

Examine the regression summary. The  $t$ -statistic associated with the coefficient in front of the lagged value of  $y$  is equal to -11.48. Compare this to the appropriate (non-standard) critical values from Table A in Enders (2004). The 1% and 5% critical values of  $t_{nc}^*$  (the  $t$ -statistic in a regression with no constant and no deterministic trend) for a sample size = 100 are -2.60 and -1.95, respectively. Since the  $t$ -value from the regression is **less than** both of these critical values, we can reject the hypothesis that  $\pi = 0$  and that there is a unit root present in this series.

**Conclusion:**  $x$  is  $I(1)$ .

**Note:** The time series was, in fact, a random walk. It was simulated in  $\mathbb{R}$  by writing:

```
y <- arima.sim(list(order=c(0,1,0)), n=100)
```

- Different models have different  $\tau$  distributions. These values are also reported in Table A in Enders (2004).

$$\Delta Y_t = \delta Y_{t-1} + \varepsilon_t$$

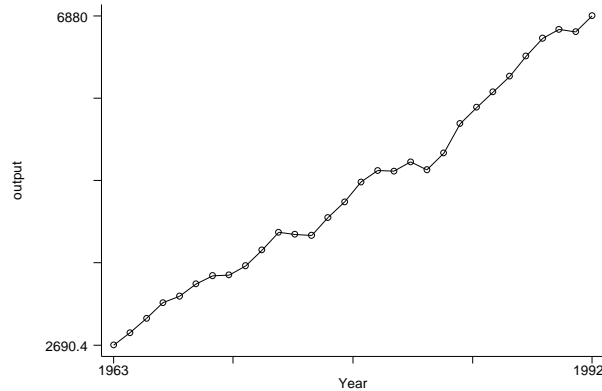
$$\Delta Y_t = \alpha + \delta Y_{t-1} + \varepsilon_t$$

$$\Delta Y_t = \alpha + \beta t + \delta Y_{t-1} + \varepsilon_t$$

- Note that under  $H_0$  the distribution of the  $F$ -test statistic also changes. So, to test the hypothesis that  $\alpha = 0 = \beta$  (in the equation above) we have to use the critical  $\phi$ -values calculated by Dickey and Fuller. These values can be found in Table B in Enders (2004).

## Example 2 *Testing Annual US GDP for a Unit Root*

**step 1:** Graph US GDP. Does the series look stationary?



Annual US GDP, 1963-1992.

**step 2:** Implement the DF-test with no constant and no trend:

```
dfuller.reg <- lm(diff(output) ~ lag(output,-1)[1:29] -1)
summary(dfuller.reg)
```

The  $t$ -statistic on lagged output is 6.91. The critical value,  $\tau$ , from Enders' Table A is -1.95 (5%,  $n = 25$ ). Since  $t > \tau$ , we can not reject  $H_0 : \pi = 0$ . We can, therefore, not reject the presence of a unit root in annual US GDP.

**step 3:** Implement the DF-test with a constant, but no trend:

```
dfuller.reg <- lm(diff(output) ~ lag(output,-1)[1:29])
summary(dfuller.reg)
```

The  $t$ -statistic on lagged output is 0.391. The critical value,  $\tau_\mu$ , from Enders' Table A is -3.00 (5%,  $n = 25$ ). Since  $t > \tau_\mu$ , we can not reject  $H_0 : \pi = 0$ . We can, therefore, not reject the presence of a unit root in annual US GDP.

**step 4:** Implement the DF-test with a constant and a trend:

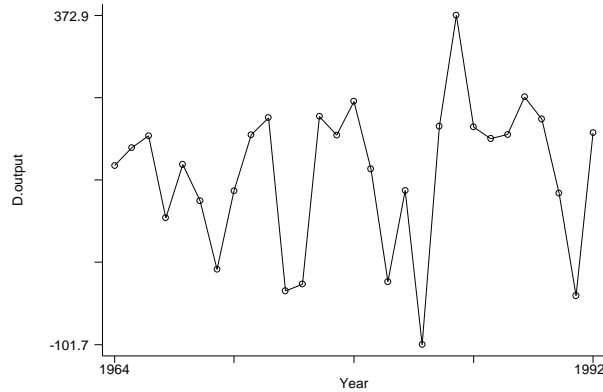
```
dfuller.reg <- lm(diff(output) ~ lag(output,-1)[1:29] + Year[1:29])
summary(dfuller.reg)
```

The  $t$ -statistic on lagged output is -1.674. The critical value,  $\tau_\tau$ , from Enders' Table A is -3.60 (5%,  $n = 25$ ). Since  $t > \tau_\tau$ , we can not reject  $H_0 : \pi = 0$ . We can, therefore, not reject the presence of a unit root in annual US GDP.

**Conclusion:** All three tests tell us that US GDP has at least one unit root.

**step 5:** Test if GDP is  $I(2)$ .

**step 6:** Graph  $\Delta\text{GDP}$ . Does this series look stationary?



**step 7:** Implement the DF-test on  $\Delta\text{GDP}$  with no constant and no trend:

```
dfuller.reg <- lm(diff(diff(output)) ~ lag(diff(output),-1)[1:28] -1)
summary(dfuller.reg)
```

The  $t$ -statistic on lagged  $\Delta\text{GDP}$  is -2.082. The critical value,  $\tau$ , from Enders' Table A is -1.95 (5%,  $n = 25$ ). Since  $t < \tau$ , we can reject  $H_0 : \pi = 0$ . We can, therefore, reject the presence of a unit root in annual  $\Delta\text{GDP}$ .

**step 8:** Implement the DF-test on  $\Delta\text{GDP}$  with a constant, but no trend:

```
dfuller.reg <- lm(diff(diff(output)) ~ lag(diff(output),-1)[1:28])
summary(dfuller.reg)
```

The  $t$ -statistic on lagged  $\Delta\text{GDP}$  is -4.223. The critical value,  $\tau_\mu$ , from Enders' Table A is -3.00 (5%,  $n = 25$ ). Since  $t < \tau_\mu$ , we can reject  $H_0 : \pi = 0$ . We can, therefore, reject the presence of a unit root in  $\Delta\text{GDP}$ .

**step 9:** Implement the DF-test on  $\Delta\text{GDP}$  with a constant and a trend:

```
dfuller.reg <- lm(diff(diff(output)) ~ lag(diff(output),-1)[1:28] + Year[1:28])
summary(dfuller.reg)
```

The  $t$ -statistic on lagged  $\Delta\text{GDP}$  is -4.206. The critical value,  $\tau_\tau$ , from Enders' Table A is -3.60 (5%,  $n = 25$ ). Since  $t < \tau_\tau$ , we not reject  $H_0 : \pi = 0$ . We can, therefore, reject the presence of a unit root in  $\Delta\text{GDP}$ .

**Conclusion:** All three tests reject the presence of a unit root (at the 5% level) in  $\Delta\text{GDP}$ . Thus, we can conclude that annual US GDP is  $I(1)$ .

## 1.2 The Augmented Dickey-Fuller (ADF) Test

- The DF-test we assume  $\varepsilon_t$  and  $\varepsilon_{t-k}$  are uncorrelated for all  $k$ .
- You can check the correlogram (using the **acf** and/or **pacf** functions in  $\mathbb{R}$ ) of the residuals from the DF regression to see if this holds or not.
- If this does not hold, we can augment the DF-test by adding additional lagged differences to the test equation. Fortunately, this does not change the distribution of the test statistic. So, we can continue using the same table for the DF-test statistics as before

$$\Delta Y_t = \alpha + \beta t + \delta Y_{t-1} + \gamma_i \sum_{i=1}^m \Delta Y_{t-i} + \varepsilon_t.$$

- How does one choose the appropriate lag length  $m$ ? The main goal is to eliminate autocorrelations in the regression residuals.
  1. Start with a general model (i.e., start with too many lags) and then reduce it through a series of  $F$ - and/or  $t$ -tests.
  2. Alternatively, one could start by adding lags until the newly added lag is insignificant. The risk with this method is that significant lags can come afterwards, especially if the data has a periodicity less than one year.
    - (a) Or, just add lags until there is no autocorrelation!
  3. AIC
  4. BIC
  5. Regardless of the method you use, you have to check the regression residuals for autocorrelation!
- My preferred method is to have as few lags as is necessary to eliminate autocorrelation in the residuals. Why? increasing the number of lags lowers the power of the test. Keep in mind, however, that with quarterly data (for example) you have to look at 4 lags, 8 lags, etc...not just lags 1, 2, 3, etc.
- In practice, I always do both. I wouldn't want to publish a paper with results that were overly sensitive to the number of lags included or excluded from this test.

## 1.3 The Power of the DF Unit Root Test

**Type II Error:** A type II error is that the null hypothesis is not rejected when the alternative is true.

- When the "truth" deviates a lot from the null hypothesis (e.g., if, in an AR(1) process,  $\theta = 0.01$ , while our  $H_0$  is  $\theta = 1$ ) the probability of making a type II error is small.
- Conversely, when the "truth" deviates only slightly from the null hypothesis (e.g. if, in an AR(1) process,  $\theta = 0.98$ , while our  $H_0$  is  $\theta = 1$ ) the probability of making a type II error is large.

**The Power of a Test:** the power of a test is equal to 1 - the probability of making a type II error, i.e. the probability of rejecting the null hypothesis when it is, indeed, false.<sup>1</sup>

- The power of our DF and ADF tests is small in many important cases, due to the high level of persistency and/or trends in macroeconomic variables.
- The result is that we tend to "over accept" the existence of a unit root.

## 1.4 The Phillips-Perron Test

- In the DF-test we assume that  $\varepsilon_t$  and  $\varepsilon_{t-k}$  are uncorrelated for all  $k$ .
- The ADF-test dealt with this potential problem by added lagged values of the dependent variable to the left hand side of our regression, i.e., we added  $\gamma_i \sum_{i=1}^m \Delta Y_{t-i}$ .
- The PP-test deals with this potential problem using nonparametric statistical methods which takes care of serial correlation without added lagged differences.
- The **PP.test** in  $\mathbb{R}$  uses the Newey-West estimator of the variance-covariance matrix.
- The **PP.test** in  $\mathbb{R}$  estimates the DF style equation with a constant and a time trend.
- The **PP.test** in  $\mathbb{R}$  maintains  $H_0 : \theta = 1$ , while the alternative is  $H_1 : \theta < 1$ .

**Example 3** *Simulate a random walk process and use the PP-test for a unit root:*

```
x<-arima.sim(list(order=c(0,1,0)),n=200)
PP.test(x)
```

This delivers the following results:

*data: x*

*Dickey-Fuller = -2.2946*

*Truncation lag parameter = 4*

*p-value = 0.4527*

A p-value greater than 0.10 means we can not reject  $H_0 : \theta = 1$ .

---

<sup>1</sup>The "size" of a test is the percentage of times we falsely reject the null hypothesis.

**Example 4** *Simulate a white noise process and use the PP-test for a unit root:*

```
x<-rnorm(200)
```

```
PP.test(x)
```

This delivers the following results:

*data: x*

*Dickey-Fuller = -31.37*

*Truncation lag parameter = 7*

*p-value = 0.01*

Thus, we can reject the unit root hypothesis at the 1% level.

- **IMPORTANT! A super helpful tip!** Due to the low power of our unit root tests. A clear rejection by the **PP.test** means that you can conclude that the series has no unit root and you don't have to do any more testing.

## 1.5 Sequential Unit Root Test

The sequential method of testing for unit roots outlined in Enders (2004) and in Pfaff (2006) is the method we will use to organize our search for unit roots. It is essential that you follow a specified methodology if you are to avoid data mining.

See pp. 213-214 in Enders (2004) *Applied Econometric Time Series (2nd edition)*.

See pp. 21-29 and 55-58 in Pfaff (2006).

**Example 5** *Test Gujarati's (2003) PDI (personal disposable income) series for a unit root:*

**step 1:** Run DF regression on the full model.

```
trend<-1:87
```

```
df<-lm(diff(PDI)~PDI[1:87]+trend)
```

$t\text{-stat} = -2.588$

$\tau_\tau = -3.45$  (5%, n=100)

Conclusion: We cannot reject  $H_0: \pi = 0$  (we cannot reject a unit root).

**step 1a:** Check the residuals for autocorrelation.

```
acf(residuals(df), ci.type="ma")
```

Conclusion: no autocorrelation.

**step 2:**  $F$ -test for  $trend = \pi = 0$ .

```
df1<-lm(diff(PDI)~1)
anova(df1,df)
Analysis of Variance Table
Model 1: diff(PDI) ~1
Model 2: diff(PDI) ~PDI[1:87] + trend
Res.Df RSS Df Sum of Sq F Pr(>F)
1 86 67080
2 84 61998 2 5081 3.4422 0.03657 *
Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
F-statistic = 3.44
 $\phi_3$ -statistic = 6.49 (5%, n=100)
 $F < \phi_3 \rightarrow$  Conclusion: We cannot reject that  $trend = \pi = 0$ .
```

**step 3:** Run DF regression on the model without the trend.

```
df2<-lm(diff(PDI)~PDI[1:87])
t-stat = -0.672
 $\tau_\mu = -2.89$  (5%, n=100)
Conclusion: We cannot reject  $\pi = 0$  (i.e., a unit root).
```

**step 3a:** Check the residuals for autocorrelation.

```
acf(residuals(df2), ci.type="ma")
Conclusion: no autocorrelation.
```

**step 3b:**  $F$ -test for  $intercept = \pi = 0$

$$F = \frac{(RSS_r - RSS_u)/q}{RSS_r/(T - k)}$$

where  $T$  = the number of observations,  $k$  the number of independent variables estimated (including the intercept), and  $q$  is the number of restrictions tested.

```
RSSu<-sum(residuals(df2)^2)
RSSr<-sum(diff(PDI)^2)
q<-2
T=length(diff(PDI))
k<-3
F=((RSSr-RSSu)/q)/(RSSu/(T-k))
F
F=17.76
```

**Step 3c:** Re-do the  $F$ -test using the anova command.

```
df2<-lm(diff(PDI)~PDI[1:87])
```

```
df3<-lm(diff(PDI)~-1)
```

```
anova(df3,df2)
```

```
F=17.971
```

```
 $\phi_1 = 4.71$  (5%, n=100)
```

Conclusion: We can reject the hypothesis that  $intercept = \pi = 0$ .

**step 3d:** Use standard  $t$ -statistic to test whether or not  $\pi = 0$ . **WHY!?**

```
 $t$ -statistic = -0.672
```

Conclusion: We cannot reject  $\pi = 0$ .

**Conclusion:** PDI is a random walk with drift.

## 2 Testing for Unit Roots in the Presence of Structural Breaks

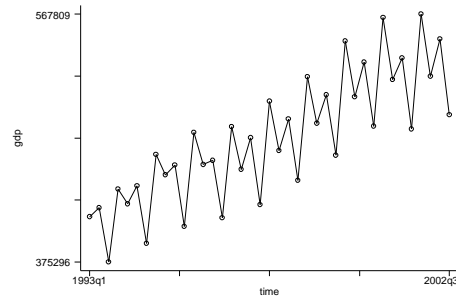
- Structural breaks (in otherwise stationary series) introduce bias in standard unit root towards accepting the unit root hypothesis. Why?
- Pfaff (2006, pp. 73-78) demonstrates the implementation of the Zivot-Andrews unit root test in  $\mathbb{R}$ . The name of the test is `ur.za()`. It can be found in a package named `urca` along with a large number of other unit root tests.
- Enders (2004, pp. 200-207) discusses the problem of structural change and discusses Perron's (1989) test for unit roots in the presence of structural breaks.
- We will return to the topic of structural breaks in a later lecture.

## 3 Seasonal Unit Roots

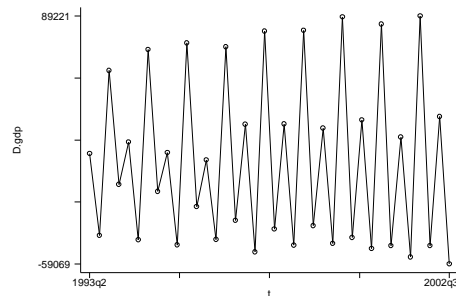
- A seasonal quarterly process has four possible roots; "1", "-1", "+i" and "-i".
  - "1" corresponds to the unit root, which can be removed by first differencing.
  - "-1" implies a two-period (biannual) cycle.
  - " $\pm i$ " produce seasonal unit roots, which can be removed by taking fourth differences.

**Example 6** *What would a stochastic seasonal trend look like?*

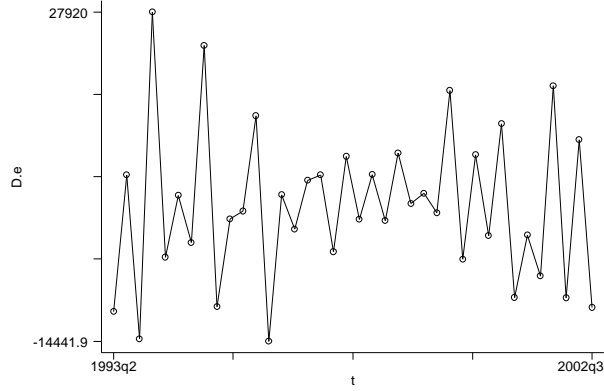
- Graph Quarterly Swedish GDP. Does the series look stationary? Does this series look like it has a seasonal component? Is this seasonal component deterministic or stochastic?



- Assuming that this series is  $I(1)$ , we can take first differences to make it stationary. This gives us the following graph. Does this series have a constant mean and/or a constant variance? If so, the seasonal component is deterministic. If not, it is stochastic.



- Alternatively, we could remove a deterministic seasonal component by using the variable dummy method. This results in a nonstationary, seasonally detrended series for Swedish GDP. What do the first differences of this series look like? Is the mean and/or variance of this series constant? If so, then it was correct to remove a deterministic seasonal component. If not, then we haven't dealt correctly with the seasonal component.



- Let us now perform a formal test for the presence of a stochastic seasonal component.

### 3.1 Hyllberg, Engel, Granger & Yoo (HEGY) Test

- The description of this test follows Pfaff (2006, p. 82)
- The HEGY test regression for quarterly data is

$$\Delta_4 y_t = \sum_{i=1}^4 \pi_i Y_{i,t-1} + \varepsilon_t$$

where the regressors  $Y_{i,t-1}$  for  $i = 1, \dots, 4$  are constructed as

$$\begin{aligned} Y_{1,t} &= y_t + y_{t-1} + y_{t-2} + y_{t-3} \\ Y_{2,t} &= -y_t + y_{t-1} - y_{t-2} + y_{t-3} \\ Y_{3,t} &= -y_t + y_{t-2} \\ Y_{4,t} &= Y_{3,t-1} = -y_{t-1} + y_{t-3}. \end{aligned}$$

Then we test for  $\pi_1 = 0$ ,  $\pi_2 = 0$  and  $\pi_3 = \pi_4 = 0$ .

- The test regression can be generalized to include deterministic components such as an intercept, trend, seasonal dummy variables, as well as lagged seasonal differences.
- The HEGY-test can be performed in  $\mathbb{R}$  using the **HEGY.test** function in the **uroot** package.