BANK OF ENGLAND

Economic modelling and forecasting

2-6 February 2015

Unit roots

Ole Rummel

Adviser, CCBS at the Bank of England

ole.rummel@bankofengland.co.uk

BANK OF ENGLAND

Philosophy of my presentations

Everything should be made as simple as possible, but not simpler

Albert Einstein

Outline

- What do macroeconomists do?
- Why do we need to worry about unit roots and nonstationarity?
- Spurious regressions
- Units roots and nonstationarity
- Testing for unit roots: the (augmented) Dickey-Fuller test
- Unit-root tests, seasonality and structural change
- Moving beyond the Dickey-Fuller test (more ‘power’)
- Confirmatory analysis
- Modifications to the standard ADF test
- New developments in unit-root testing
- Summary

What do macroeconomists do?

- (Central-bank) macroeconomists do – at least – four things (Stock and Watson (2001)):
  - they describe and summarise macroeconomic data;
  - they make macroeconomic forecasts;
  - they quantify what we do and do not know about the true structure of the macroeconomy; and
  - they advise (and sometimes become) macroeconomic policymakers

Why do we need to worry about unit roots?

- Macroeconomic data often contain a **unit root**
- It turns out that models of **nonstationary data** (i.e., data that contain unit roots) require special treatment
- We must test whether series are stationary because:
  – the mean, variance and higher moments may depend on time (implications for modelling and forecasting);
  – standard statistics do not have standard distributions (implications for hypothesis testing);
  – of the problem of nonsense or spurious regressions (implications for estimation and model-building);
  – a nonstationary time series implies that the effects of shocks never die out (economic implications?); and
  – of cointegration (Granger (1981), Engle and Granger (1987))

Spurious regressions (1)

- Take two **independent** time series generated by these data-generating processes (DGPs):

$$y_t = y_{t-1} + u_t \qquad u_t \sim N(0,\sigma_u^2) \qquad (1)$$

$$x_t = x_{t-1} + v_t \qquad v_t \sim N(0,\sigma_v^2) \qquad (2)$$

- As we will see, these two processes are said to be **integrated** or to contain a **unit root**
- Consider the regression:

$$y_t = \alpha + \beta x_t + \varepsilon_t \qquad (3)$$

- This is a simple regression model but – because x_t and y_t are independent – there is **no** relationship between the variables
- If any relationship were to be detected, this would be a spurious or nonsense relationship, giving rise to a **spurious regression**

Spurious regressions (2)

- Yule (1926) pointed out that β is significant in such models even for quite large sample sizes
- If the two series also contain a time trend then the associated R^2 is often very high
- This seems to indicate that we should be cautious about regressions based on variables containing a unit root
- Granger and Newbold (1974) found that such regressions often have a high R^2 but a Durbin-Watson (*DW*) statistic close to zero (indeed, the R^2 often exceeds the *DW*)
- But if the model is estimated in **first differences**:

$$\Delta y_t = \beta \Delta x_t + \varepsilon_t \qquad (4)$$

then $R^2 \approx 0$, indicating no relationship, and $DW \approx 2$

Spurious regressions (3)

- Moreover, integrated data lead to problems with the distributions of standard *t*- and *F*-statistics
- Granger and Newbold (1974) suggested that critical values of the order of 11.2 (rather than the usual value of 1.96) should be used for inference based on *t*-statistics
- Phillips (1986) analysed spurious regressions rigorously and found expressions for the limiting distribution
- In particular, he was able to show that the estimated β converges to a **non-zero** random variable (rather than its true value of zero)...
- ...and concluded that scaling the calculated t_β by $\sqrt{T}$ is better, such that the 5 per cent significance level for a sample size of 50 becomes $\sqrt{50} \times 1.96 \approx 13.9$

Spurious regressions (4)

- But things can get very complicated, indeed
- Entorf (1997) showed that for unit root processes **with drift**, i.e.:

$$y_t = \mu + y_{t-1} + u_t \qquad u_t \sim N(0,\sigma_u^2) \tag{5}$$

$$x_t = \gamma + x_{t-1} + v_t \qquad v_t \sim N(0,\sigma_v^2) \tag{6}$$

  the estimated β from a regression involving y_t and x_t converges to a constant ($\beta \rightarrow \mu/\gamma$)
- Spurious regression results therefore depend critically on the underlying time-series properties of the data
- We need to investigate unit root processes further if we are to rescue time-series regressions

Unit roots: definition of stationarity

- A stochastic process, y_t, is **weakly** or **covariance stationary** if it satisfies the following requirements:
  - E(y_t) is constant and independent of t
  - Var(y_t) is constant and independent of t
  - Cov(y_t, y_{t-k}) = $\gamma(k)$ is constant and independent of t for all k
- 'Classical' econometric theory assumes that observed data come from a stationary process, where population means and variances are constant over time

Unit roots: nonstationarity (1)

- Take the following two processes:

$$x_t = \alpha + \rho x_{t-1} + u_t \qquad |\rho| < 1 \tag{7}$$

$$y_t = \alpha + y_{t-1} + v_t \tag{8}$$

- Then, at any time T ($T \geq 1$):

$$x_T = \rho^T x_0 + \alpha \sum_{i=0}^{T-1} \rho^i + \sum_{i=0}^{T-1} \rho^i u_{T-i} \tag{9}$$

$$y_T = y_0 + \alpha T + \sum_{i=0}^{T-1} v_{T-i} \tag{10}$$

Unit roots: nonstationarity (2)

- Even if $x_0 = y_0 = 0$, we get:

$$E(x_T) = \frac{\alpha}{1-\rho}$$

$$E(y_T) = \alpha T$$

where the expectation of y_T, E(y_T), is not independent of T

- We also note that a stationary series (such as x_t) is **mean reverting**: it has a tendency to return to its mean value
- **Nonstationary** series (such as y_t) are not like that: they can go **anywhere**

Unit roots: nonstationarity (3)

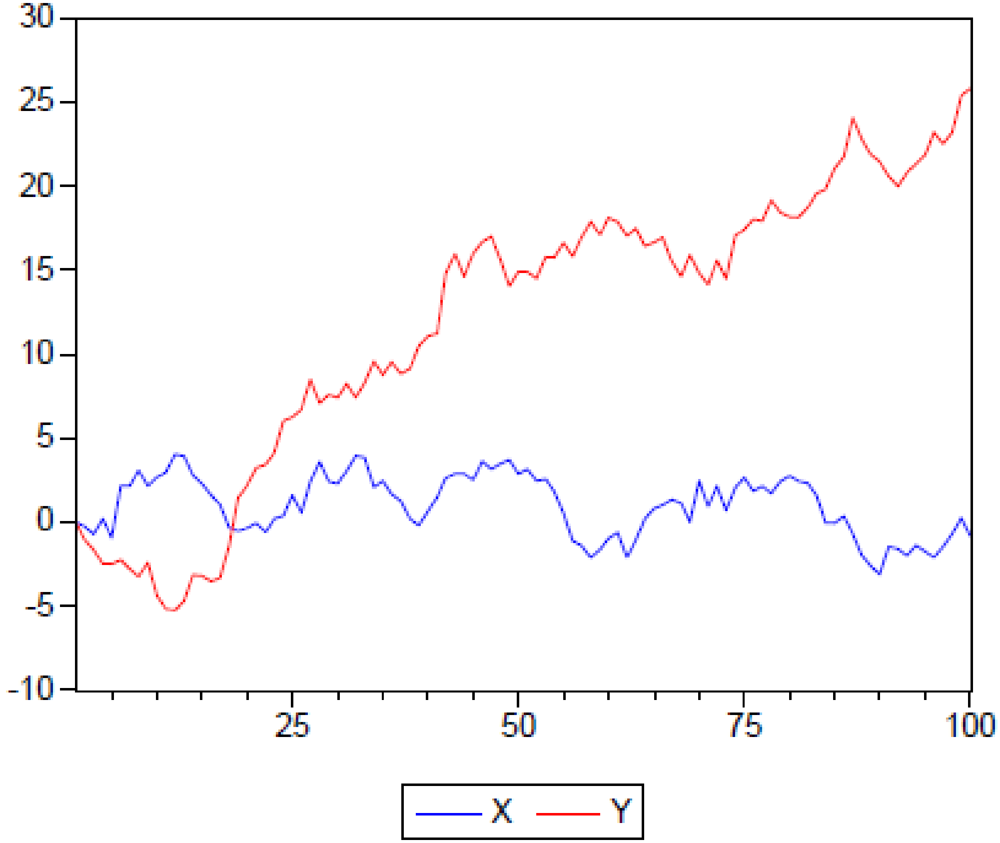

Unit roots: nonstationarity (4)

- We also get:

$$\mathrm{var}(x_T) = \frac{\sigma^2}{1-\rho^2}$$

$$\mathrm{var}(y_T) = T\sigma^2$$

where the variance of y_T, var(y_T), is not independent of T – in fact, as T gets bigger, the variance tends to infinity

- The second process, y_t, is actually a **random walk with drift**, which, if α > 0, tends to rise over time – such a series is said to contain a **stochastic trend**

Unit roots: trend-stationarity

- While the term **trend** is deceptively easy to define, trends in practice are fairly tricky to deal with
- The main reason for this is that there are two very different types of trending behaviour that are difficult to distinguish
- Much economic data grows over time so two alternative models are:

$$y_t = \alpha + y_{t-1} + u_t \qquad (11)$$

$$y_t = \alpha + \beta t + v_t \qquad (12)$$

- The second of these is a **trend-stationary** process in which t is a simple time trend taking integer values from 1 to T
- Shocks to (12) are transitory and the process reverts back to a (linear) **deterministic trend** rather than a constant

Dealing with nonstationarity (1)

- A process is said to be integrated of order *d*, denoted I(*d*), if it can be rendered stationary by differencing *d* times, that is, y_t is nonstationary, but $\Delta^d y_t$ is stationary
- If $d = 2$, then y_t is I(2) and needs to be differenced twice to achieve stationarity:

$$\begin{aligned} \Delta^2 y_t &= \Delta(\Delta y_t) = \Delta(y_t - y_{t-1}) \\ &= \Delta y_t - \Delta y_{t-1} = (y_t - y_{t-1}) - (y_{t-1} - y_{t-2}) \\ &= y_t - 2y_{t-1} + y_{t-2} \end{aligned}$$

- By analogy, a stationary process is integrated of order zero, I(0), and does not require any differencing to achieve stationarity

Dealing with nonstationarity (2)

- Dealing with trend-stationarity throws up an important point
- We have seen that differencing a series containing a stochastic trend is the correct course of action – is that also the case for a trend-stationary series?
- The strategy of differencing a trend-stationary series is **not** to be recommended, as this results in:

$$y_t = \alpha + \beta t + v_t$$

$$y_t - y_{t-1} = \alpha + \beta t + v_t - (\alpha + \beta(t-1) + v_{t-1})$$

$$\Delta y_t = \beta + v_t - v_{t-1} \qquad (13)$$

- In other words, by taking the first-difference of the trend-stationary series we have introduced a moving-average error which has a unit root – this is known as **over-differencing**

Testing for unit roots

- Suppose a series is generated by an AR(1) process:

$$y_t = \alpha + \rho y_{t-1} + \varepsilon_t \qquad (14)$$

- You might test for a **random walk**, equivalent to $\rho = 1$, by estimating:

$$y_t = \alpha + \rho y_{t-1} + \varepsilon_t$$

  by OLS and testing whether $\rho = 1$

- When the null hypothesis, H_0: $\rho = 1$, is true, (14) reduces to:

$$y_t = \alpha + y_{t-1} + \varepsilon_t$$

$$y_t - y_{t-1} = \Delta y_t = \alpha + \varepsilon_t$$

  such that Δy_t is stationary, meaning that y_t in levels is a **random walk** and thus nonstationary

The Dickey-Fuller unit-root test (1)

- In practice, the Dickey-Fuller (DF) test for a unit root, equivalent to $\rho = 1$ in (14), uses the regression:

$$(y_t - y_{t-1}) = \alpha + (\rho - 1)y_{t-1} + \varepsilon_t \qquad (15)$$

or

$$\Delta y_t = \alpha + (\rho - 1)y_{t-1} + \varepsilon_t \qquad (16)$$

- A standard test statistic would be:

$$t_\rho = \frac{\hat{\rho} - 1}{\sqrt{\mathrm{var}(\hat{\rho})}}$$

which is a ‘normally’ defined test with the null hypothesis of H_0: $(\rho - 1) = 0$ or, equivalently, H_0: $\rho = 1$

- The OLS *t*-statistic testing $(\rho - 1) = 0$ or, equivalently, $\rho = 1$ in (16), is called the **Dickey-Fuller statistic**

The Dickey-Fuller unit-root test (2)

- But this test statistic has a distribution different from the usual *t*-distribution (an alternative statistic, based on the so-called normalised bias, $T(\rho - 1)$, can also be used)
- This test uses the – nonstandard – Dickey-Fuller distribution (Dickey and Fuller (1979))
- We must use this distribution instead of a *t*-distribution because the OLS estimate of ρ is biased downwards...
- ...which means that we need to use the distribution tabulated by Dickey to find the correct critical values for our **one-tailed** test – the alternative hypothesis is H_1: $\rho < 1$
- The key point is that the estimate for ρ is **not** (usually) normally distributed

The Dickey-Fuller unit-root test (3)

- To complicate matters further, the critical values for these tests depend on the form of the underlying model and whether a drift or time trend is included
- In particular, the Dickey-Fuller test regression must be extended to deal with the possibility that under the alternative hypothesis, the series may be stationary around a deterministic trend
- Potential models are based on the general specification:

$$\Delta y_t = \mu + \delta t + (\rho - 1)y_{t-1} + \varepsilon_t \qquad (17)$$

where:

  - $\mu = \delta = 0$ is referred to as a model with no drift;
  - $\delta = 0$ is referred to as a model with drift; and
  - otherwise as a model with trend

The Dickey-Fuller unit-root test (4)

- There are therefore three forms of the Dickey-Fuller test, namely:
  - **model 1** (no drift): $\Delta y_t = (\rho - 1)y_{t\text{-}1} + \varepsilon_t$;
  - **model 2** (with drift): $\Delta y_t = \mu + (\rho - 1)y_{t\text{-}1} + \varepsilon_t$; and
  - **model 3** (with trend): $\Delta y_t = \mu + \delta t + (\rho - 1)y_{t\text{-}1} + \varepsilon_t$
- For each of the three models the form of the Dickey-Fuller test is still the same, namely the test of whether $(\rho - 1) = 0$
- For a sample size of 100 and a significance level of 5 per cent, the critical value for each of these models is:

Model	Coefficients	5 per cent critical value
No drift	$\mu = 0, \delta = 0$	-1.95
With drift	$\mu \neq 0, \delta = 0$	-2.89
With trend	$\mu \neq 0, \delta \neq 0$	-3.45

The Dickey-Fuller unit-root test (5)

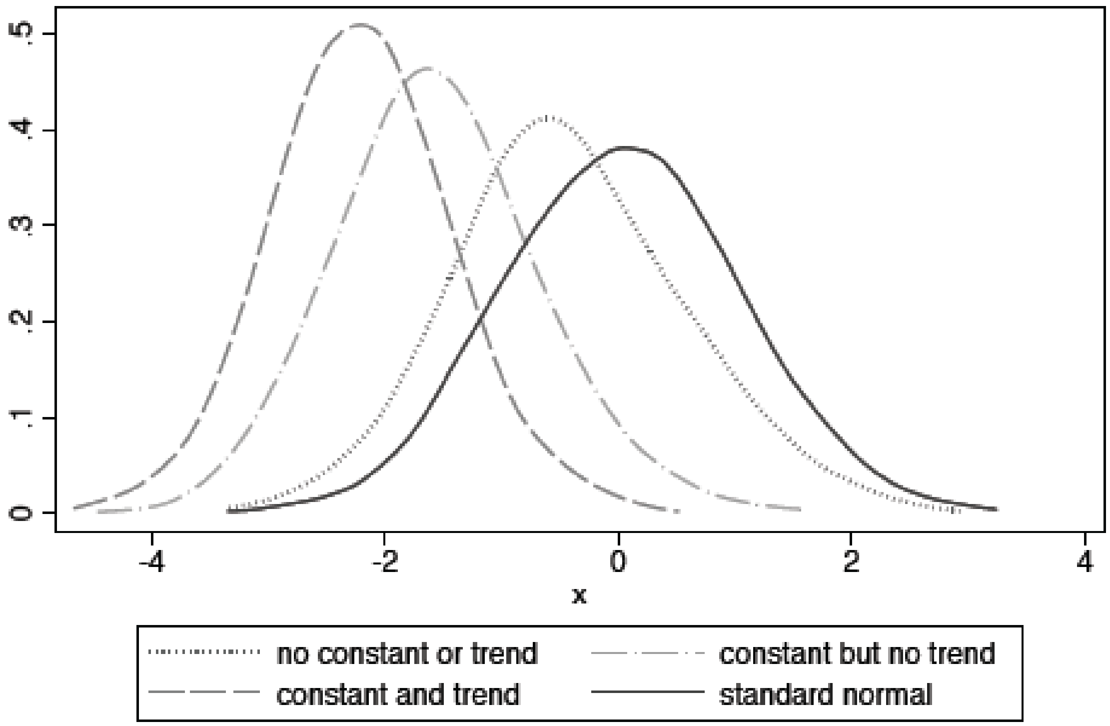

- The pertinent distribution of the different versions of the Dickey-Fuller test statistics changes depending on whether a constant and/or a time trend is included

The augmented Dickey-Fuller unit-root test (1)

- The Dickey-Fuller unit-root test is strictly only valid if the residuals are identically and independently distributed (iid)
- If there is residual autocorrelation we need to use a ‘different’ test...
- ...which can be straightforwardly generalised
- The **augmented Dickey-Fuller (ADF)** unit-root test is given by:

$$\Delta y_t = \mu + \delta t + (\rho - 1) y_{t-1} + \sum_{i=1}^{k} \beta_i \Delta y_{t-i} + \varepsilon_t \qquad (18)$$

which allows for a more complex dynamic structure that accounts for residual autocorrelation

- The test is then whether (ρ – 1) = 0, which is – again – distributed as the Dickey-Fuller test

The augmented Dickey-Fuller unit-root test (2)

- The test statistic associated with the ‘best’ time-series model is then used
- We need – again – to be careful about specifying which form of the test is correct: with or without trend and so on
- This again turns out to be quite complicated

The augmented Dickey-Fuller unit-root test (3)

- The **regression** run and the **actual** data-generating process both affect the critical values
- If they are of particular forms, then they are sometimes based on Dickey-Fuller distributions, variations on this and sometimes normal distributions
- Some caution should be used when interpreting ADF tests

Model	Data-generating process	Asymptotic distribution
No drift, no trend	$\mu = 0, \delta = 0$	Dickey-Fuller
With drift, no trend	$\mu = 0, \delta = 0$ $\mu \neq 0, \delta = 0$	‘Dickey-Fuller like’ 1 Normal
With drift, with trend	$\mu = 0, \delta = 0$ $\mu \neq 0, \delta \neq 0$	‘Dickey-Fuller like’ 2 Normal

Unit-root tests: selecting the proper lag length *k* (1)

- The choice of additional lags to purge ε_t of any residual autocorrelation may be based on a sequential general-to-specific testing procedure:
  - Hall (1994) showed than when the order *k* is selected through *t*-tests on the β_k to β_1 parameters in (18) (or *via* an application of the (modified) information criteria), the relevant DF statistics still apply
- Banerjee *et al.* (1993, p. 107) favour a generous parameterisation since:
  - *...if too may lags are present...the regression is free to set them to zero at the cost of some loss of efficiency, whereas too few lags implies some remaining autocorrelation...and hence the inapplicability of even the asymptotic distributions....*
- All of these must be checked against the relevant DF statistics - note that introducing extra variables reduces the power of the test

Unit-root tests: selecting the proper lag length *k* (2)

- The choice of additional lags to purge ε_t of any residual autocorrelation may also be based on (modified) information criteria
- Ng and Perron (1995, 2001) analysed rules for truncating long autoregressions when performing unit-root tests and developed information criteria with penalty functions adequate for integrated time series
- **Modified** information criteria (which are data-dependent) are superior to conventional information criteria in truncating long autoregressions with integrated variables when moving average errors are present
- Their preferred criterion is the MAIC, which can be interpreted as a modified form of the Akaike information criterion (AIC)

A possible unit-root testing strategy for the ADF test

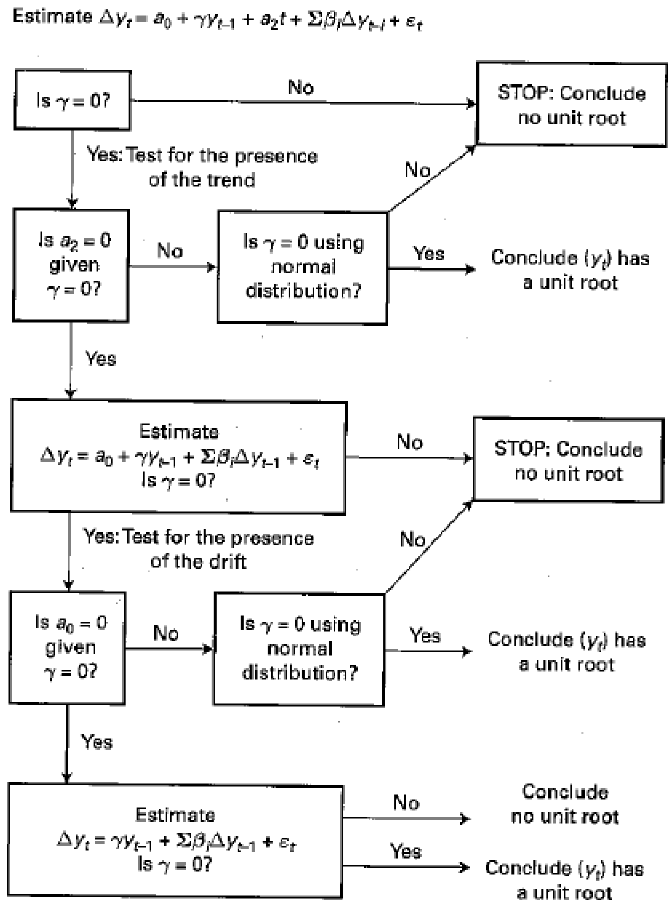

Source: Enders (1995, p. 257), based on Dolado *et al.* (1990).

The Phillips-Perron unit-root test

- Rather than using (long) autoregressions to approximate general (serially) dependent processes, the unit-root tests due to Phillips and Perron (1988) adjust for serial correlation nonparametrically by directly modifying the test statistics
- Phillips-Perron tests have the same limiting null distribution as the DF distribution and therefore the same critical values
- But these tests rely on asymptotic theory, which means that large samples (i.e., long data series) are required for them to work well
- Perron and Ng (1996) consider modified Phillips-Perron tests (referred to as *M*-tests)
- These *M*-tests appear to be robust to, e.g., measurement errors and additive outliers in the observed series

Unit-root tests: seasonality and breaks

- Also note that:
  - seasonally adjusted data tend to look more integrated than they really are (there is a tendency for the OLS estimate of ρ in the ADF test (18) to be biased toward one when y_t is a seasonally adjusted series, thus rejecting the null hypothesis of nonstationarity substantially less often that it should according to the appropriate critical values); and
  - a break (shift) in the mean or trend also make the data appear non-stationary
- In other words, whenever possible, seasonally **unadjusted** data should be used when testing for unit roots, since the filters used to adjust for seasonal patterns often distort the underlying properties of the data (Davidson and MacKinnon (1993), Section 19.6)

Unit-root tests and structural change (1)

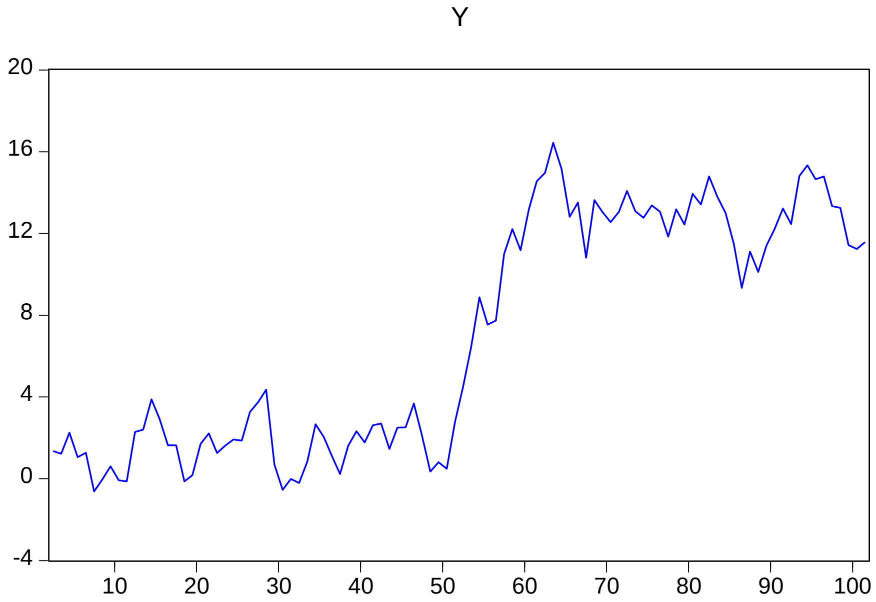

- Consider the following data-generating process:

$$y_t = 0.2 + 0.75y_{t\text{-}1} + dum_t + \varepsilon_t \qquad (19)$$

with a representative plot

Unit-root tests and structural change (2)

Null Hypothesis: Y has a unit root
Exogenous: Constant, Linear Trend
Lag Length: 0 (Automatic - based on Modified AIC, maxlag=12)

		t-Statistic	Prob.*
Augmented Dickey-Fuller test statistic		-2.114385	0.5313
Test critical values:	1% level	-4.052411	
	5% level	-3.455376	
	10% level	-3.153438	

*MacKinnon (1996) one-sided p-values.

- Results of an ADF test on the generated data in EViews

Unit-root tests and structural change (3)

Null Hypothesis: Y has a unit root
Exogenous: Constant, Linear Trend
Lag length: 0 (Spectral GLS-detrended AR based on Modified AIC, maxlag=12)
Sample: 2 101
Included observations: 100

		MZa	MZt	MSB	MPT
Ng-Perron test statistics		-8.17614	-1.98613	0.24292	11.2578
Asymptotic critical values*:	1%	-23.8000	-3.42000	0.14300	4.03000
	5%	-17.3000	-2.91000	0.16800	5.48000
	10%	-14.2000	-2.62000	0.18500	6.67000

*Ng-Perron (2001, Table 1)

HAC corrected variance (Spectral GLS-detrended AR)	1.601272

- M-tests on the generated data in EViews

Unit-root tests and structural change (4)

- Dealing with the structural break is straightforward when the timing of the (single) structural break is known
- We include a dummy variable, $BREAK_t$, in (18) to capture the structural break according to:

$$\Delta y_t = \mu + \delta t + (\rho - 1) y_{t-1} + \sum_{i=1}^{k} \beta_i \Delta y_{t-i} + \gamma BREAK_t + \varepsilon_t \qquad (20)$$

where the structural break dummy variables is defined as:

$$BREAK_t = \begin{cases} 0 & : \quad t \leq \tau \\ 1 & : \quad t > \tau \end{cases} \qquad (21)$$

and τ is the observation where the break occurs

Unit-root tests and structural change (5)

- The unit-root test is still based on testing $(\rho - 1) = 0$, but the *p*-values of the test statistic are now also a function of the timing of the structural break, τ
- The correct *p*-values for a unit root with a (single) structural break are available in Perron (1989), who argues that most time series are trend-stationary if one allows for a one-time change in the intercept or in the slope (or both) of the trend function
- Certain ‘big shocks’ do not represent a realisation of the underlying data-generating process and the null should be tested against the trend-stationarity alternative by allowing, both under the null and the alternative hypotheses, for the presence of a one-time break (at a known point in time) in the intercept or in the slope (or both) of the trend function

Unit-root tests and structural change (6)

Zivot-Andrews Unit Root Test
Sample: 2 101
Included observations: 100
Null Hypothesis: Y has a unit root with a structural break in the intercept
Chosen lag length: 0 (maximum lags: 4)
Chosen break point: 52.00000

	t-Statistic	Prob. *
Zivot-Andrews test statistic	-5.137049	2.03E-05
1% critical value:	-5.34	
5% critical value:	-4.93	
10% critical value:	-4.58	

* Probability values are calculated from a standard t-distribution and do not take into account the breakpoint selection process

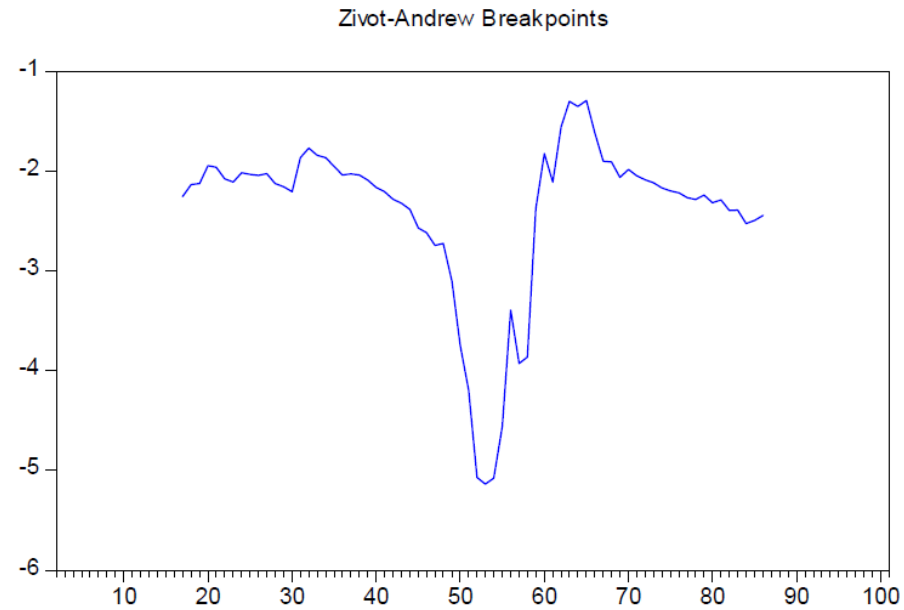

- Zivot and Andrews (1992) unit-root tests with an **unknown** break date on the generated data in EViews

Art vs. science: telling time series apart (1)

- It can be difficult to tell unit root and trend stationary processes apart
- We generated three time series *w*, *x* and *z* in EViews:
  - one series is trend stationary;
  - another series is difference stationary with a drift (i.e., it should not be constant around a time trend); and
  - the third series has a *near unit root*, i.e., the AR term, ρ, is close to one

Art vs. science: telling time series apart (2)

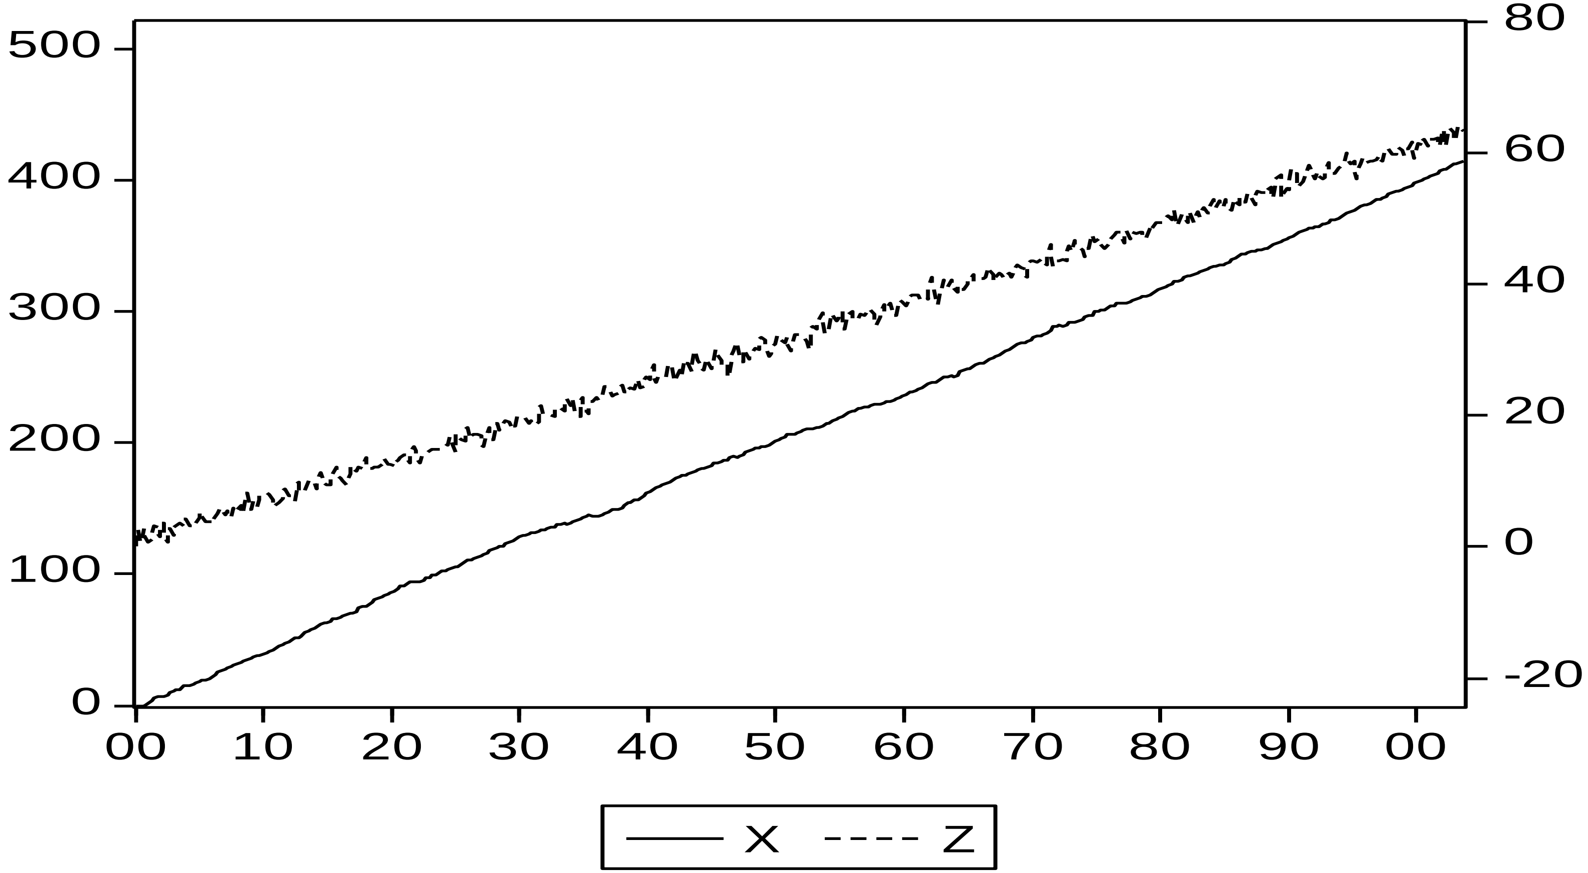

Art vs. science: telling time series apart (3)

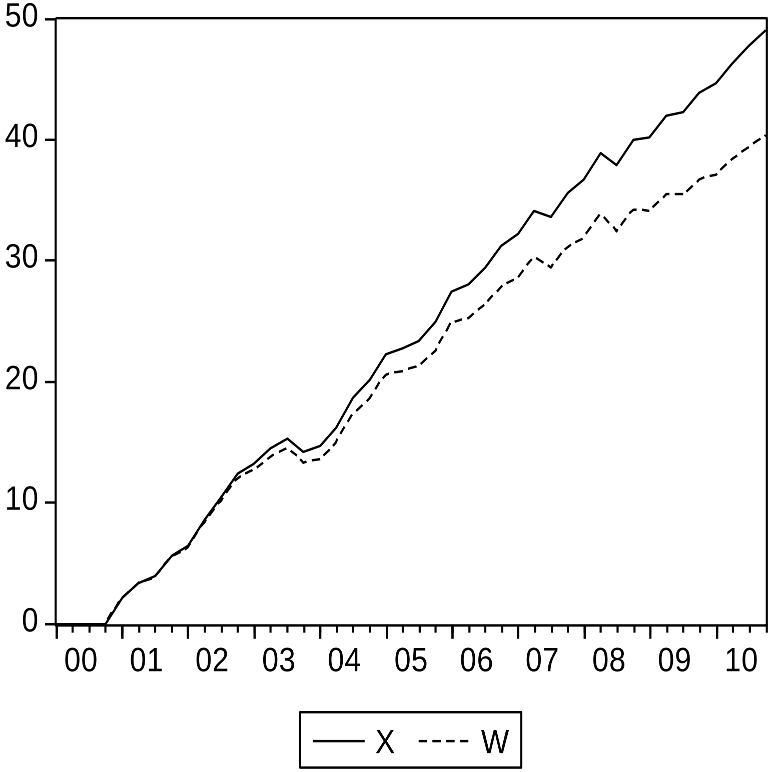

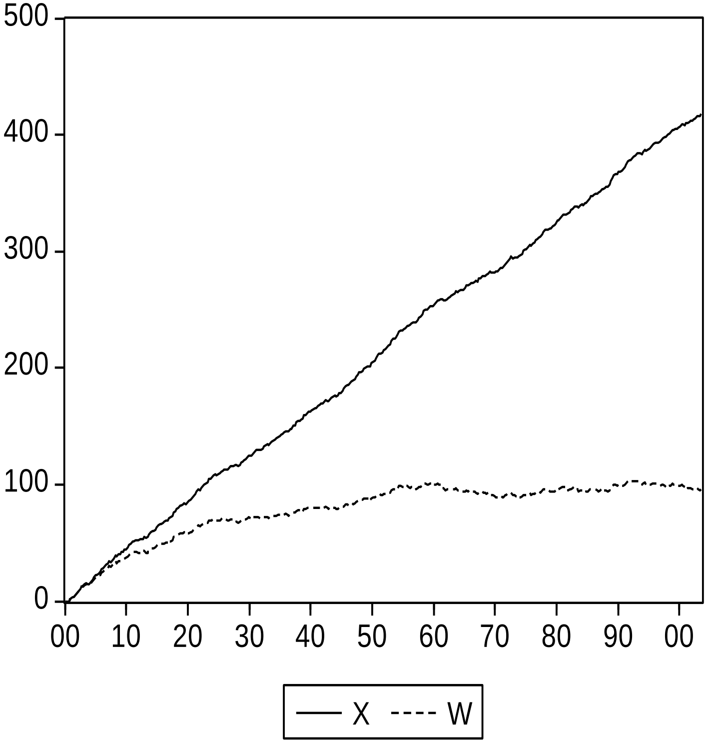

Art vs. science: telling time series apart (4)

- The processes for x_t and z_t are given by:

$$x_t = 1 + x_{t-1} + 0.6u_t \qquad (22)$$

$$z_t = 1 + 0.15trend + 0.8e_t \qquad (23)$$

  where both u_t and e_t are white noise
- Could you tell the difference with a near-unit root process?

$$w_t = 1 + 0.97w_{t-1} + 0.6u_t \qquad (24)$$

Art vs. science: ADF tests (1)

- It is difficult to tell series apart over small samples
- Note the low power of the ADF unit-root test (the null hypothesis is that the series in question has a unit root):

(small sample, 40 observations)

x: ADF statistic	0.3932	*p*-value	0.9804
w: ADF statistic	-0.4921	*p*-value	0.8828
z: ADF statistic	0.4520	*p*-value	0.8076
z: ADF trend	-8.1189	*p*-value	0.0000

Art vs. science: ADF tests (2)

- It is difficult to tell series apart over small samples
- Note the low power of the ADF unit-root test (the null hypothesis is that the series in question has a unit root):

(big sample, 400 observations)

x: ADF statistic	-0.7704	*p*-value 0.8258
w: ADF statistic	-6.9012	*p*-value 0.0000
z: ADF statistic	3.5121	*p*-value 0.9999
z: ADF trend	-20.648	*p*-value 0.0000

- 400 observations – this is a long time series (most central banks do not have that many observations!)

The power of ADF tests (1)

- Tests for unit roots are not especially good at distinguishing between a slowly mean-reverting series (ρ close to 1) and a true unit root process ($\rho = 1$)
- Part of the problem concerns the **power** of the test and the presence of the deterministic regressors in the estimating equation
- Formally, the power of a test is equal to the probability of rejecting a false null hypothesis given a sample of finite size
- A test with **higher** power than the ADF is more likely to reject the null hypothesis of a unit root against the stationary alternative when the alternative is true; thus, a more powerful test is better able to distinguish between $\rho = 1$ and a root that is large but less than one

The power of ADF tests (2)

- Note that, unlike many other hypotheses testing situations, the power of tests of the unit root hypothesis against stationary alternatives depends very little on the number of observations *per se* and very much on the **span** of the data

The power of ADF tests (3)

- It is generally alleged that unit-root tests have 'low' power – should practitioners abandon the use of augmented Dickey-Fuller tests? Can we do better?
- There are several different alternatives available to us (Haldrup and Jansson (2006)):
  - **confirmatory analysis** using alternative unit-root tests;
  - **modifications** to the standard Dickey-Fuller class of tests (which can achieve improved power);
  - the use of (stationary) **covariates**;
  - assumptions regarding the **initial condition**; and
  - **non-Gaussian** (Rothenberg and Stock (1997)) and **GARCH errors** (Seo (1999))

Testing the null of stationarity: the KPSS test (1)

- Kwiatkowski *et al.* (KPSS) (1992) proposed a test for testing the null hypothesis of stationarity or I(0) (i.e., the **absence** of a unit root) rather than testing for the null of nonstationarity or I(1) (i.e., the **presence** of a unit root)
- Consider the regression:

$$y_t = \alpha + \delta t + z_t \qquad (25)$$

with:

$$z_t = z_{t-1} + \varepsilon_t \qquad \varepsilon_t \sim N(0,\sigma_\varepsilon^2) \qquad (26)$$

- The test is to see if $z_t = z_{t-1}$ or that it is constant…
- …which is done by testing the null hypothesis H_0: $\sigma_\varepsilon^2 = 0$
- Although not directly apparent, this null hypothesis also implies a unit root (over-differencing) in the representation of Δy_t

Testing the null of stationarity: the KPSS test (2)

- The test is one of **parameter constancy** rather than a direct test of the null hypothesis of stationarity
- Set up as in (25), this is a test of trend stationarity: for simple stationarity, drop the time trend from the regression
- The critical values for the KPSS test are:

Model	5 per cent critical value	1 per cent critical value
Drift only	0.463	0.739
With drift and trend	0.146	0.216

The tests due to Elliott *et al.* (1996) (1)

- Elliott *et al.* (1996) (ERS) showed that it is possible to enhance the power of a unit-root test by estimating the model using something close to first differences, which is called **local GLS detrending**
- The so-called DF-GLS test, which is a more efficient version of the ADF test, is computed in two steps
- Instead of creating the first difference of y_t, ERS preselect a constant close to unity, α (= 1 + $\underline{c}/T$), and subtract αy_{t-1} from y_t, resulting in a **quasi-differenced** or **local detrended** series denoted y_t^d
- The value of α that seems to provide the best power is $\underline{\alpha} = (1 - 7/T)$ for the case of an intercept ($\underline{c}$ = -7) and $\underline{\alpha} = (1 - 13.5/T)$ if there is an intercept and trend ($\underline{c}$ = -13.5)

The tests due to Elliott *et al.* (1996) (2)

- ERS then estimate the basic ADF regression using the local detrended data, y_t^d:

$$\Delta y_t^d = \varphi y_{t-1}^d + \sum_{j=1}^{m} \beta_j \Delta y_{t-j}^d + v_t \qquad (27)$$

- The lag length *m* is selected using information criteria (ERS suggest the SBC/BIC) and the null hypothesis of a unit root can be rejected if $\varphi \neq 0$
- Because the quasi-differenced or local detrended data depend on the value of $\underline{c}$, the critical values are different to the DF critical values which rely on simple detrending
- The critical values of the test also depend on whether or not a trend is included in the test

The tests due to Elliott *et al.* (1996) (3)

- The second unit-root test defined in the paper, known as the ERS (feasible) point optimal or P_T test is also based on the quasi-differenced of local detrended data, y_t^d
- Rather than use an ADF-type test on y_t^d, the ERS (feasible) point optimal test assesses the null hypothesis that $\alpha = 1$ against the alternative hypothesis that $\alpha = \underline{\alpha}$ directly

Unit-root tests with (stationary) covariates (1)

- We have seen that the inclusion of **deterministic** regressors affects the distribution of the (augmented) Dickey-Fuller tests
- In most applications of unit-root tests, the series y_t being tested for a unit root is not observed in isolation
- Instead, one typically observes at least one time series, say z_t, in addition to the time series y_t of interest
- What happens when we include other **stochastic** regressors (or covariates) into the ADF regression?
- Hansen (1995) argued that standard univariate unit-root tests that ignore useful information from correlated (stationary) covariates lead to less powerful unit-root tests

Unit-root tests with (stationary) covariates (2)

- Hansen (1995) proposed constructing a more powerful unit-root test – the covariate-augmented DF or CADF test – using the following model:

$$\Delta y_t^d = \varphi y_{t-1}^d + \sum_{j=1}^{m} \beta_j \Delta y_{t-j}^d + \sum_{j=-m}^{m} \mu_j z_{t-j}^d + v_t \qquad (28)$$

  where Δ is the difference operator, y_t^d is the GLS-demeaned and detrended variable, y_t, to be tested, z_t^d is the OLS-demeaned and detrended vector z_t of stationary covariates and v_t is a disturbance error

- Hansen (1995), Elliott and Jansson (2003) and Pesavento (2006) have shown that the null hypothesis, H_0: $\varphi = 0$, in the presence of stochastic covariates does not have a standard Dickey-Fuller distribution

Unit-root tests with (stationary) covariates (3)

- The use of unit-root tests using a single variable does not make economic sense in many situations and can sometimes give misleading results
- The idea of using covariates in unit-root testing is a good one, but the fact that only I(0) covariates can be considered is very restrictive:
  - these variables are judged to be I(0) by some earlier tests that are based on single-variable unit-root tests
- Why not jointly test for unit roots in y_t and z_t instead of a test for unit roots in y_t conditional on z_t being I(0)?

There are many unit-root tests available

- In testing for a unit root, one faces a large array of possible methods
- Monte Carlo studies do not point to any dominant test
- There are two reasons for this:
  - a uniformly most powerful test does not exist (as proven theoretically in Elliott *et al.* (1996)); and
  - an ordering of unit-root tests depends on the treatment of the initial condition of the process

Unit-root tests and the initial condition

- We generally make the (apparently innocuous) assumption that the (unobserved) initial observation y_0 is equal to zero
- But the power of any unit-root test will depend on the **deviation** of the initial observation y_0 from its modelled deterministic part
- In fact, the initial condition turns out to be a crucial aspect of the unit-root testing problem, as it profoundly affects both the power and the form of optimal tests
- Müller and Elliott (2003) find that all popular unit-root tests are close to optimal, but that their idiosyncratic power characteristics can be explained by different implied treatments of the initial condition
- The choice of unit-root tests in practice comes down to what types of initial conditions are likely for the application at hand

Things may not be as bad as they seem

- Recent research, examining the role of the initial condition, y_0, has become less critical of ADF tests
- While Müller and Elliott (2003) found that the ADF test has near optimality properties, tests based on Dickey-Fuller statistics are attractive only when the deviation of the initial observation, y_0, from its modelled deterministic part is 'large' – but can we come up with a compelling reason why the potentially mean-reverting series should start off far from its equilibrium value?
- A useful choice of unit-root test when the initial condition is relatively 'small' is the ERS test described above

Unit roots and nonlinearity (1)

- Established tests for a unit root, such as the ADF, assume a linear adjustment process for the series back to its long-run equilibrium value (called the **attractor**)
- Although the Dickey-Fuller test can be augmented with deterministic regressors and lagged changes of y_t, the dynamic adjustment process is still assumed to be linear…and hence symmetric
- Pippenger and Goering (1993) as well as Balke and Fomby (1996) have shown that tests for unit roots have low power in the presence of **asymmetric** adjustment

Unit roots and nonlinearity (2)

- There is a large and growing literature designed to test for the presence of an attractor in the presence of nonlinear adjustment
- For example, Enders and Granger (1998) generalise the Dickey-Fuller methodology to consider the null hypothesis of a unit root against the alternative hypothesis of a momentum threshold autoregressive (MTAR) model

Unit roots and nonlinearity (3)

- The TAR model in this case is given by:

$$\Delta y_t = I_t \rho_1(y_{t-1} - \tau) + (1 - I_t)\rho_2(y_{t-1} - \tau) + \varepsilon_t$$

$$I_t = \begin{cases} 1 & : \quad \Delta y_{t-1} \geq \tau \\ 0 & : \quad \Delta y_{t-1} < \tau \end{cases} \qquad (29)$$

where I_t is a dummy variable that divides the observations on y_t into two subsets

- The attractor is τ since Δy_t has an expected value of zero when $\Delta y_{t-1} = \tau$
- Using (29) above, it is possible to test for an attractor even though the adjustment process is nonlinear

Unit roots and nonlinearity (4)

- To begin with, note that if $\rho_1 = \rho_2 = 0$, the process is a random walk
- A sufficient condition for the $\{y_t\}$ sequence to be stationary is $-2 < (\rho_1, \rho_2) < 0$
- The Dickey-Fuller test emerges as a special case in which $\rho_1 = \rho_2$
- If it is not possible to reject the null hypothesis $\rho_1 = \rho_2 = 0$, it can be concluded that there is an attractor – but again, as in the Dickey-Fuller tests, it is not possible to use a classical *F*-statistic to test the null hypothesis

Unit roots and nonlinearity (5)

- If the null hypothesis of nonstationary is rejected, it is possible to test for symmetric versus asymmetric adjustment
- In particular, if the null is rejected (meaning the process has an attractor), then we can perform the tests for symmetric adjustment ($\rho_1 = \rho_2$) using a standard F-distribution
- An alternative to the MTAR model is the TAR model, in which the adjustment depends on $y_{t\text{-}1}$ rather than on $\Delta y_{t\text{-}1}$
- While the MTAR is favoured as having more power as a unit-root test over the TAR model, the recommendation is to estimate both and select the adjustment mechanism by a model selection criterion such as the AIC or SBC/BIC

Summary (1)

- ‘Classical’ econometric theory generally assumes stationary data (particularly constant means and variances across time periods), but empirical evidence is strongly against the validity of this assumption
- Moreover, stationarity is an important basis for empirical modelling (‘spurious regressions’), and statistical inference when the stationarity assumption is incorrect can induce serious mistakes
- Due to the effects of the assumption of a unit root in a variable on both the econometric method used and the economic interpretation of the model examined, it is quite common to pre-test the data for unit roots

Summary (2)

- When using augmented Dickey-Fuller unit-root tests, confidence in the empirical outcomes is obtained when the test results appear robust to changes in the sample size, outliers, additional lags and the inclusion or exclusion of deterministic components
- But be aware of the low power of unit-root tests – and the near-observational equivalence of unit roots and stationary processes with roots close to one in **finite** samples (although the distinction between a unit-root process and a near unit-root process need not be crucial for practical modelling)
- Unit-root tests can be constructed with both excellent size and local asymptotic power properties, but to achieve these dual objectives it is necessary to use GLS detrended data

References and further reading (1)

Balke, N S and Fomby, T B (1997), ‘Threshold cointegration’, *International Economic Review*, Vol. 38, No. 3, pages 627-43.

Banerjee, A, Dolado, J J, Galbraith, J W and Hendry, D F (1993), *Co-integration, error correction and the econometric analysis of non-stationary data*, Oxford, Oxford University Press.

Campbell, J Y and Perron, P (1991), ‘Pitfalls and opportunities: what macroeconomists should know about unit roots’, in Blanchard, O J and Fischer, S (eds.), *NBER Macroeconomics Annual*, Vol. 6, pages 141-201.

Dickey, D A and Fuller, W A (1979), ‘Distribution of the estimators for autoregressive time series with a unit root’, *Journal of the American Statistical Association*, Vol. 74, No. 366, pages 427-31.

References and further reading (2)

Davidson, R and MacKinnon, J G (1993), *Estimation and inference in econometrics*, Oxford, Oxford University Press.

Dickey, D A and Pantula, S G (2002), 'Determining the order of differencing in autoregressive processes', *Journal of Business & Economic Statistics*, Vol. 20, No. 1, pages 18-24.

Elliott, G and Jansson, M (2003), 'Testing for unit roots with stationary covariates', *Journal of Econometrics*, Vol. 115, No. 1, pages 75-89.

Elliott, G, Rothenberg, T J and Stock, J H (1996), 'Efficient tests for an autoregressive unit root', *Econometrica*, Vol. 64, No. 4, pages 813-36.

References and further reading (3)

Enders, W and Granger, C W J (1998), 'Unit-root tests and asymmetric adjustment with an example using the term structure of interest rates', *Journal of Business & Economic Statistics*, Vol. 16, No. 3, pages 304-11.

Engle, R F and Granger, C W J (1987), 'Co-integration and error correction: representation, estimation and testing', *Econometrica*, Vol. 55, No. 2, pages 251-76.

Granger, C W J (1981), 'Some properties of time series data and their use in econometric model specification', *Journal of Econometrics*, Vol. 16, No. 1, pages 121-30.

Granger, C W J and Newbold, P (1974), 'Spurious regressions in econometrics', *Journal of Econometrics*, Vol. 2, No. 2, pages 111-20.

References and further reading (4)

Haldrup, N and Jansson, M (2006), ‘Improving size and power in unit root testing’, *Palgrave handbook of econometrics, vol. 1: econometric theory*, Basingstoke, Palgrave Macmillan, pages 252-77.

Hall, A (1994), ‘Testing for a unit root in time series with pretest data-based model selection’, *Journal of Business & Economics Statistics*, Vol. 12, No. 4, pages 461-70.

Hansen, B E (1995), ‘Rethinking the univariate approach to unit root testing: using covariates to increase power’, *Economic Theory*, Vol. 11, No. 5, pages 1148-71.

References and further reading (5)

Kwiatkowksi, D, Phillips, P C B, Schmidt, P and Shin, Y (1992), ‘Testing the null hypothesis of stationarity against the alternative of a unit root: how sure are we that economic time series have a unit root?’, *Journal of Econometrics*, Vol. 54, No. 1-3, pages 159-78.

Maddala, G and Kim, I-M (1998), *Unit roots, cointegration, and structural change*, Cambridge, Cambridge University Press.

Müller, U K and Elliott, G (2003), ‘Tests for unit roots and the initial condition’, *Econometrica*, Vol. 71, No. 4, pages 1269-86.

Ng, S and Perron, P (1995), ‘Unit root tests in ARMA models with data-dependent methods for the selection of the truncation lag’, *Journal of the American Statistical Association*, Vol. 90, No. 429, pages 268-81.

References and further reading (6)

Ng, S and Perron, P (2001), 'Lag length selection and the construction of unit root tests with good size and power', *Econometrica*, Vol. 69, No. 6, pages 1519-54.

Perron, P (1989), 'The great crash, the oil price shock and the unit root hypothesis', *Econometrica*, Vol. 57, No. 6, pages 1361-1401.

Perron, P and Ng, S (1996), 'Useful modifications to some unit root tests with dependent errors and their local asymptotic properties', *Review of Economic Studies*, Vol. 63, No. 3, pages 435-65.

Pesavento, E (2007), 'Near-optimal unit root tests with stationary covariates with better finite sample size', mimeo, Emory University, Department of Economics.

References and further reading (7)

Phillips, P C B and Perron, P (1988), ‘Testing for a unit root in time series regression’, *Biometrika*, Vol. 75, No. 2, pages 333-46.

Pippenger, M K and Goering, G E (1993), ‘A note on the empirical power of unit root tests under threshold processes’, *Oxford Bulletin of Economics and Statistics*, Vol. 55, No. 4, pages 473-81.

Rothenberg, T J and Stock, J H (1997), ‘Inference in a nearly integrated autoregressive model with nonnormal innovations’, *Journal of Econometrics*, Vol. 80, No. 2, pages 269-86.

Seo, B (1999), ’Distribution theory for unit root tests with conditional heteroskedasticity’, *Journal of Econometrics*, Vol. 91, No. 1, pages 113-44.

References and further reading (8)

Stock, J H (1987), 'Asymptotic properties of least squares estimators of cointegrating vectors', *Econometrica*, Vol. 55, No. 5, pages 1035-56.

Stock, J H and Watson, M W (2001), 'Vector autoregressions', *Journal of Economic Perspectives*, Vol. 15, No. 4, pages 101-15.

Zivot, E and Andrews, D W K (1992), 'Further evidence on the great crash, the oil price shock and the unit-root hypothesis', *Journal of Business & Economic Statistics*, Vol. 10, No. 3, pages 251-70.