Non-Stationary Time Series

FRM Part I

1 A story: three kinds of movement

- Sales: a steady climb needs a trend
- Seaside shop: a yearly rhythm needs a seasonal pattern
- Share price: may wander with no fixed centre
- Each movement needs its own tool

IN THE LESSON

- Picture three charts on a desk.
- The first shows a company's sales, climbing a little every quarter.
- The second shows a seaside shop's takings, which jump every summer and sink every winter, year after year.
- The third shows the price of a share, which seems to wander wherever it likes, and today's price is the best clue to tomorrow's.
- Each chart moves in a different way, and each needs a different tool.
- The first needs a trend, the second needs a seasonal pattern, and the third may have no fixed centre at all.
- This lesson gives you the tool for each, and a way of telling which one you face.

2 What a trend model says

- Series = predictable part of t + random shock
- t counts the periods: 1, 2, 3, ...
- The shock, ε, cannot be predicted
- Key question: what shape is the predictable part?

IN THE LESSON

- A time series is a list of values recorded in order, such as monthly sales or a daily index level.
- A trend model says that the expected value of the series depends on time itself.
- We number the periods one, two, three and so on, and call that number t.
- The model splits each observation into two parts: a predictable part that depends on t, and a random shock, which we call epsilon, that nobody can predict.
- The question is the shape of the predictable part.
- It might be a straight line, a curve that grows by a steady percentage, or a bend that rises and then falls.
- Choosing the shape that fits the pattern is the first skill in this topic.

3 The linear trend

$$y_t = \beta_0 + \beta_1 \times t + \varepsilon_t$$

The series, y	The series at time t
Starting level, beta zero	Level when t is zero
Slope, beta one	Constant change per period
Time index, t	The period number
Shock, epsilon	Unpredictable shock

IN THE LESSON

- The simplest shape is a straight line.
- The series equals a starting level, beta zero, plus beta one times the time index, plus a shock.
- Beta one is the slope, and it has a plain meaning: the constant change in the level of the series each period.
- If beta one is two, the series rises by two units every period, whatever its current level.
- If beta one is negative, it falls along a straight line instead.
- This suits series that add roughly the same amount each period, such as a steady rise in a count of customers.
- It is easy to estimate with ordinary regression of the series on the time index.
- Its weakness is that the same amount is added forever, however large the series becomes.

4 The log-linear trend

$$\ln y_t = \beta_0 + \beta_1 \times t + \varepsilon_t$$
$$\text{Level} = \exp(\beta_0 + \beta_1 \times t)$$

Log of the series, ln y	Natural log of the series
Starting level, beta zero	Log of the starting level
Slope, beta one	Continuous growth rate per period
Time index, t	The period number

IN THE LESSON

- Many financial series, such as asset values, volumes or sales, grow by a roughly constant percentage, not a constant amount.
- A series that grows by three per cent a year adds more units each year as it gets bigger, so a straight line fits badly.
- The fix is to take the natural logarithm of the series and fit a straight line to that.
- The slope, beta one, is now the continuously compounded growth rate per period.
- To get back to the level, undo the logarithm with the exponential function.
- The fitted level equals the exponential of beta zero plus beta one times t.
- The common mistake is to read beta one as a unit change in the series.
- It is a growth rate.

5 Example 1: straight line against steady growth

Illustration. Linear: $y = 100 + 2t$. Log-linear: $\ln y = 4.5 + 0.03t$

Linear: change each period	constant	**+2**
Log-linear level at t = 10	exp(4.5 + 0.03×10)	**121.51**
Log-linear level at t = 11	exp(4.5 + 0.03×11)	**125.21**
Growth factor per period	exp(0.03)	**1.0305**

IN THE LESSON

- Here are two made up trends side by side.
- The linear trend is one hundred plus two t.
- Its change per period is always plus two, the same at every level.
- The log-linear trend has beta zero of four point five and beta one of point zero three.
- At time ten the log is four point five plus point three, which is four point eight, and the exponential of that is about one hundred twenty one point five.
- At time eleven the log is four point eight three, and the level is about one hundred twenty five point two.
- The growth factor between the two is the exponential of point zero three, about one point zero three zero five, so the series grows by about three per cent each period and adds more units as it grows.

6 Curved trends: quadratic and beyond

- Quadratic: $y = \beta_0 + \beta_1 t + \beta_2 t^2 + \varepsilon$
- $\beta_2 < 0$: rises, peaks, falls. $\beta_2 > 0$: a valley
- Peak at $t = -\beta_1 \div (2 \times \beta_2)$
- Choose the form that matches the pattern of growth

IN THE LESSON

- Some series bend.
- A quadratic trend adds the square of time to the regression, so the slope itself changes as time passes.
- A polynomial trend can add even higher powers.
- Suppose the coefficient on t squared is negative.
- Then the series rises, slows, peaks and turns down, like a ball thrown into the air.
- The turning point is minus beta one divided by two times beta two, where beta two is the coefficient on t squared.
- A positive coefficient on t squared gives the opposite shape, a valley followed by an ever steeper climb.
- Do not add powers just to chase the data, because each extra term uses up a parameter.
- Choose the form that matches the pattern of growth that you see.

7 Example 2: where does the quadratic peak?

Illustration. $y = 5 + 6t - 0.1t^2$. So $\beta_1 = 6$ and $\beta_2 = -0.1$

Peak time	$-6 \div (2 \times -0.1)$	**t = 30**
Level at the peak	$5 + 180 - 90$	**95**
Level at t = 60	$5 + 360 - 360$	**5**

IN THE LESSON

- Take a made up quadratic: five, plus six t, minus point one t squared.
- Beta one is six and beta two is minus point one.
- The peak comes at minus six divided by two times minus point one, which is minus six over minus point two, equal to thirty.
- So the series peaks at period thirty.
- What is the level there?
- Six times thirty is one hundred eighty.
- Thirty squared is nine hundred, and point one of that is ninety.
- Add it up: five plus one hundred eighty minus ninety gives ninety five.
- After period thirty the fitted series falls.
- At period sixty it is back to five, since six times sixty is three hundred sixty, and point one times three thousand six hundred is three hundred sixty.

8 What seasonality is

- Seasonality: a pattern that repeats every s periods
- Quarterly data: s = 4. Monthly data: s = 12
- A series can have a trend and seasonality together
- A seasonal dummy = 1 in its own season, 0 otherwise

IN THE LESSON

- Seasonality is a pattern that repeats every s periods.
- With quarterly data s is four, and with monthly data it is twelve.
- Think of heating bills peaking in winter, retail sales spiking before the holidays, or a farm's revenue arriving at harvest.
- These repeat at the same point in each cycle, and they are quite different from a trend, which moves in one direction.
- A series can have both at once: a rising trend with a seasonal wave on top.
- A regression can capture the wave using seasonal dummy variables.
- A dummy is a variable that equals one in its own season and zero in every other.
- In this way each season gets its own adjustment to the level.

9 Trend plus seasonal dummies

$$y_t = \beta_0 + \beta_1 t + \gamma_2 D_2 + \gamma_3 D_3 + \gamma_4 D_4 + \varepsilon_t$$

Starting level, beta zero	Baseline season's level
Slope, beta one	Trend slope per period
Seasonal dummy coefficient	Season minus baseline
Shock, epsilon	Unpredictable shock

IN THE LESSON

- Here is the regression with a trend and seasonal dummies for quarterly data.
- The series equals an intercept, plus a slope times t, plus a coefficient times the second quarter dummy, plus a coefficient times the third quarter dummy, plus a coefficient times the fourth quarter dummy, plus a shock.
- Notice that there are three dummies, not four.
- The first quarter is left out and acts as the baseline.
- The intercept is the baseline season's average level, after allowing for the trend.
- Each dummy coefficient is that season's average difference from the baseline.
- So a coefficient of eight on the second quarter dummy says that quarter two runs eight above quarter one, at the same point on the trend.

10 The dummy variable trap

Quarterly data, s = 4. Two valid set-ups and one trap

Set-up	What the coefficients mean
Intercept + 3 dummies	Dummy = gap from baseline
No intercept + 4 dummies	Dummy = season's own mean
Intercept + 4 dummies	Trap: multicollinearity

IN THE LESSON

- How many dummies should you use?
- With an intercept in the model, use one fewer than the number of seasons.
- For quarterly data, that is three.
- If you include all four, the four dummies add up to one in every row, which is exactly the intercept column.
- The regression then cannot separate them.
- This is perfect multicollinearity, called the dummy variable trap, and the software will fail or drop one for you.
- The omitted season is the baseline.
- There is a second route: drop the intercept and use all four dummies.
- Then each coefficient is simply that season's own mean level.
- The two routes describe the same fit, but the meaning of the coefficients changes, so always check which one you have.

11 Example 3: reading seasonal coefficients

Illustration. y = 100 + 2t + 8·Q2 − 5·Q3 + 15·Q4. Q1 is the baseline

t = 5 (Q1)	100 + 10 + 0	**110**
t = 6 (Q2)	100 + 12 + 8	**120**
t = 7 (Q3)	100 + 14 − 5	**109**
t = 8 (Q4)	100 + 16 + 15	**131**

IN THE LESSON

- Here is a made up quarterly model: one hundred, plus two t, plus eight for quarter two, minus five for quarter three, plus fifteen for quarter four.
- Quarter one is the baseline.
- Take time five, which is a first quarter.
- The fitted value is one hundred plus ten, which is one hundred ten, with no dummy.
- Time six is a second quarter: one hundred plus twelve plus eight gives one hundred twenty.
- Time seven is a third quarter: one hundred plus fourteen minus five gives one hundred nine.
- Time eight is a fourth quarter: one hundred plus sixteen plus fifteen gives one hundred thirty one.
- Notice that each dummy coefficient is a difference from the baseline, not a level.
- Fifteen does not mean the fourth quarter equals fifteen.

12 Testing whether seasonality is there

- Null: all dummy coefficients are zero (no seasonality)
- Alternative: some season differs from baseline
- Use a joint F test, not just single t statistics
- Large F: reject the null and keep the seasonal dummies

IN THE LESSON

- Before building a seasonal model, ask whether the seasons matter at all.
- Each dummy coefficient has its own t statistic, but the real question is joint: are all the seasonal differences zero together?
- The tool is the joint F test.
- The null hypothesis says every dummy coefficient equals zero, meaning no seasonality.
- The alternative says at least one season differs from the baseline.
- If the F statistic is larger than its critical value, reject the null and keep the dummies.
- A small, unreliable single coefficient does not decide the matter.
- A plot of the series by season, or of the residuals from a trend-only fit, usually shows the pattern long before you run the test.

13 Forecasting a seasonal series

- 1. Time index of the target date is T + h
- 2. Count h steps on, wrapping after s: the season
- 3. Forecast = b0 + b1 × (T + h) + season term
- Omitted season: its seasonal term is zero

IN THE LESSON

- Now we forecast.
- Suppose the model is a trend plus seasonal dummies and we have data up to time T.
- To forecast h steps ahead, do three things in order.
- First, find the target date's time index, which is T plus h.
- Second, work out which season that date falls in, by counting h steps forward from the last observation and wrapping round after s periods.
- For monthly data, if the last observation is December, then one step ahead is January.
- Third, evaluate the fitted equation at T plus h, and add the dummy coefficient for that season.
- If it is the omitted season, add nothing.
- Always write the season and the time index down before you calculate anything.

14 Example 4: three seasonal forecasts

Illustration. y = 100 + 2t + 8·Q2 − 5·Q3 + 15·Q4. Last data: T = 20 (Q4)

h = 1: t = 21, Q1	100 + 42 + 0	**142**
h = 2: t = 22, Q2	100 + 44 + 8	**152**
h = 6: t = 26, Q2	100 + 52 + 8	**160**
Same season, a year apart	160 − 152 = 4 periods × 2	**8**

IN THE LESSON

- We use the same made up model, and the last observation is time twenty, which is a fourth quarter.
- One step ahead, the target is time twenty one.
- Counting on from the fourth quarter, that is a first quarter, the baseline, so the forecast is one hundred plus forty two, which is one hundred forty two.
- Two steps ahead is time twenty two, a second quarter: one hundred plus forty four plus eight gives one hundred fifty two.
- Six steps ahead is time twenty six.
- Counting six steps from quarter four lands in quarter two again, so the forecast is one hundred plus fifty two plus eight, which is one hundred sixty.
- Compare the two second quarter forecasts: they differ by eight, which is four periods of trend at two each.

15 Example 5: the two usual mistakes

Illustration, same model, forecast for h = 2 from T = 20

Correct: t = 22, Q2	100 + 44 + 8	**152**
Wrong: h in place of T + h	100 + 4 + 8	**112**
Wrong: last season (Q4)	100 + 44 + 15	**159**

IN THE LESSON

- Let's see what the common errors cost, for the two step ahead forecast.
- The correct answer is time twenty two, a second quarter, so the forecast is one hundred plus forty four plus eight, one hundred fifty two.
- The first mistake is to plug in h in place of T plus h for the trend term.
- With h equal to two, the trend adds only four, and the forecast comes out as one hundred plus four plus eight, which is one hundred twelve.
- That is forty below the truth.
- The second mistake is to use the season of the last observation, quarter four, instead of the season of the target date.
- That adds fifteen where eight belongs, giving one hundred plus forty four plus fifteen, which is one hundred fifty nine.
- Find the season of the target date, and use T plus h.

16 Trend values and interval forecasts

- Trend value: insert t into the fitted equation
- Linear: b0 + b1 × t. Log-linear: exp(b0 + b1 × t)
- SE = √(RSS ÷ (n − k)), the residual std deviation
- Multiples: 1.645 (90%), 1.96 (95%), 2.576 (99%)

IN THE LESSON

- Once a trend model is estimated, the estimated trend value for any date comes from inserting that date's time index.
- For the linear model, it is the intercept plus the slope times t.
- For the log-linear model, it is the exponential of the intercept plus the slope times t.
- A point forecast h steps ahead just uses t equal to T plus h.
- A point forecast is a single number, so we add an interval around it.
- The interval is the point forecast plus and minus a multiple of the standard error of the regression.
- That standard error is the residual standard deviation: the square root of the residual sum of squares divided by n minus k.
- Here k counts the estimated parameters.
- The multiples are about one point six four five for ninety per cent, one point nine six for ninety five and two point five seven six for ninety nine.

17 The interval forecast

Interval = Point forecast ± multiple × SE
SE = √(RSS ÷ (n − k))

Point forecast	Trend value at t = T + h
Multiple (1.96 for 95%)	1.645, 1.96 or 2.576
Standard error of regression	Residual SD, not the series' SD
Limit of the interval	Ignores coefficient error; no h
Log model	Logs first, then exp: asymmetric

IN THE LESSON

- Here is the interval forecast in one line.
- The interval is the point forecast, plus or minus the multiple times the standard error of the regression.
- Be careful about what the standard error means.
- It is the standard deviation of the residuals, not the standard deviation of the series itself.
- The series has a trend, so its own spread is far too large.
- Also note two limits of this simple interval.
- It ignores uncertainty in the estimated coefficients, and its width does not grow with the horizon h.
- For a log-linear model, build the interval on the log scale first, using the log standard error, and then exponentiate both ends.
- The result is asymmetric around the point forecast.

18 Example 6: a linear trend interval

Illustration. y = 100 + 2t, n = 40, k = 2, RSS = 1520. Forecast T + h = 40 + 3

Standard error	√(1520 ÷ 38) = √40	**6.32**
Point forecast at t = 43	100 + 2 × 43	**186**
Margin for 95%	1.96 × 6.32	**12.40**
95% interval	186 ± 12.40	**173.6 to 198.4**

IN THE LESSON

- Let's build an interval from made up estimates.
- The linear trend is one hundred plus two t.
- We have forty observations and two parameters, with a residual sum of squares of fifteen hundred twenty.
- Divide by thirty eight, which is forty minus two, and you get forty.
- The square root is six point three two, the standard error.
- We forecast three steps ahead from time forty, so t is forty three.
- The point forecast is one hundred plus eighty six, which is one hundred eighty six.
- For a ninety five per cent interval, multiply one point nine six by six point three two to get about twelve point four.
- The interval runs from one hundred seventy three point six to one hundred ninety eight point four.

19 Example 7: a log-linear interval

Illustration. ln y = 4.5 + 0.03t, SE on the log scale 0.05, t = 40

Log forecast	4.5 + 0.03 × 40	**5.70**
Log interval	5.70 ± 1.96 × 0.05	**5.602 to 5.798**
Point forecast	exp(5.70)	**298.9**
95% interval for the level	exp of each end	**271.0 to 329.6**

IN THE LESSON

- Now the log-linear model, with ln y equal to four point five plus point zero three t.
- Say the standard error on the log scale is point zero five, and we want the level at time forty.
- On the log scale the point forecast is four point five plus one point two, which is five point seven.
- Ninety five per cent margin: one point nine six times point zero five is point zero nine eight.
- So the log interval runs from five point six zero two to five point seven nine eight.
- Now exponentiate each end.
- The point forecast becomes about two hundred ninety eight point nine.
- The interval runs from about two hundred seventy one to about three hundred twenty nine point six.
- The lower gap is about twenty eight, the upper gap about thirty one, so the interval is asymmetric.

20 The random walk

- Random walk: today = yesterday + an unpredictable shock
- The shock is white noise: mean zero, constant variance
- Nothing pulls the series back to an average
- Add a constant tilt each step and you have drift

IN THE LESSON

- Now we leave deterministic patterns behind.
- Imagine a person who takes a step each second, and the direction is decided by a coin toss, with no memory.
- Where will they be in a minute?
- The best guess is where they stand now, but the uncertainty widens with every step.
- This is a random walk.
- Today's value equals yesterday's value plus a shock that cannot be predicted.
- Nothing pulls the series back to an average.
- The shocks are white noise, meaning they have mean zero, constant variance and no correlation over time.
- Many asset prices are modelled this way as a first approximation.
- If the person also has a slight tilt in one direction, each step adds a constant, called drift.

21 Random walk with drift

$y_t = d + y_{t-1} + \varepsilon_t$
$\text{Var}(y_{T+h}) = h \times \sigma^2$, $\text{SD} = \sigma \times \sqrt{h}$

Drift, d	Constant drift each step
Shock, epsilon	White-noise shock
Shock volatility, sigma	Volatility of the shock
Horizon, h	Steps ahead

IN THE LESSON

- Here is the random walk with drift.
- The value today equals the drift, plus yesterday's value, plus a shock.
- Without drift, set d to zero.
- The best forecast of any future value is simply the latest value, plus d for each step ahead, so h steps ahead gives the latest value plus h times d.
- Shocks never fade.
- A shock this period is still fully present in every later value.
- The variance of the value h steps ahead, around the current value, is h times sigma squared, which grows in a straight line with the horizon.
- The standard deviation is sigma times the square root of h.
- The usual mistake is to multiply by h itself.

22 Example 8: uncertainty grows with the root of h

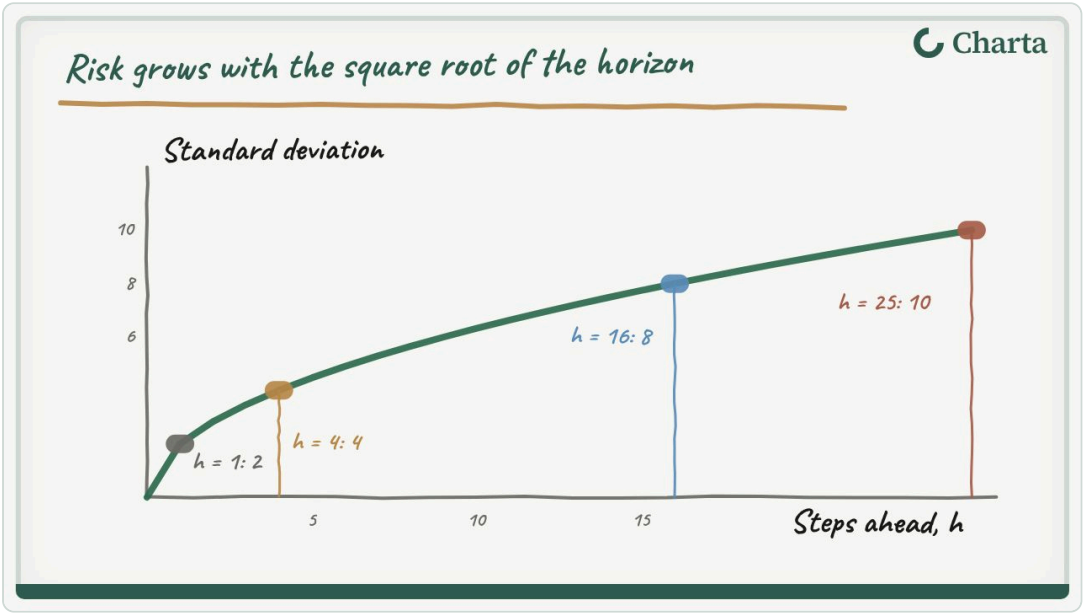
Illustration. d = 0.5, σ = 2, latest value 100

SD at h = 4	2 × √4	**4**
SD at h = 16	2 × √16	**8**
Forecast at h = 16	100 + 16 × 0.5	**108**
95% interval at h = 16	108 ± 1.96 × 8	**92.3 to 123.7**
Wrong: scale by h, not √h	2 × 16	**32**

IN THE LESSON

- Suppose a series follows a random walk with drift point five, a shock volatility of two, and a latest value of one hundred.
- These are made up numbers.
- Four steps ahead the standard deviation is two times the square root of four, which is four.
- Sixteen steps ahead it is two times the square root of sixteen, which is eight.
- Four times the horizon only doubles the risk.
- Now the forecast sixteen steps ahead is one hundred plus sixteen times point five, which is one hundred eight.
- A ninety five per cent interval is one point nine six times eight, about fifteen point seven, either side, so from about ninety two point three to about one hundred twenty three point seven.
- Scaling by h instead would give thirty two, and the interval would be far too wide.

23 Risk grows with the square root of the horizon



IN THE LESSON

- This picture shows the standard deviation of a random walk forecast, with a shock volatility of two.
- Along the bottom is the horizon, the number of steps ahead.
- Up the side is the standard deviation of the forecast error.
- The curve rises quickly at first and then flattens, because it follows the square root of h.
- At one step, the standard deviation is two.
- At four steps it has doubled to four.
- At sixteen steps it is eight, and at twenty five steps it is ten.
- A straight line would climb far faster.
- The random walk's risk never stops growing, but it grows more slowly the further out we look.

24 Unit roots and the AR(1) model

- AR(1): $y_t = \varphi \times y_{t-1} + \varepsilon_t$
- $|\varphi| < 1$: stationary, shocks fade, returns to its mean
- $\varphi = 1$: a random walk. The series has a unit root
- AR coefficients that sum to 1 also mean a unit root
- First difference of a unit-root series is stationary

IN THE LESSON

- A random walk is a special case of a more general model, the first order autoregression, or A R one.
- Here today equals phi times yesterday plus a shock.
- If phi is below one in size, the series is stationary, which means it keeps coming back to a stable average and shocks fade away.
- If phi equals exactly one, we get the random walk.
- We say the series has a unit root, because the characteristic equation of the autoregression has a root equal to one.
- Equivalently, the autoregressive coefficients add up to one.
- Such a series is non stationary, but its first difference, today minus yesterday, is stationary.
- For a random walk that difference is just the shock.

25 Why unit roots cause trouble

- Unrelated random walks: high R^2, big t, low DW
- Spurious regression; it worsens as n grows
- AR estimate biased down; usual t values mislead
- Shocks permanent; forecast error variance unbounded

IN THE LESSON

- Why do we care?
- Unit roots make ordinary regression unreliable.
- Take two independent random walks, with no connection at all, and regress one on the other.
- You will typically find a high R squared, a large t statistic and a very low Durbin Watson statistic, as if there were a strong relationship.
- This is a spurious regression, and it gets worse, not better, as the sample grows, because the t statistic no longer follows its usual distribution.
- The estimated autoregressive coefficient is also biased downward in finite samples, so ordinary critical values reject too often.
- Shocks are permanent, there is no mean reversion, and the forecast error variance grows without bound.

26 Stationary against unit root

Stationary / unit root, side by side

Feature	Stationary / Unit root
Mean reversion	Yes / None
Effect of a shock	Fades / Permanent
Forecast error variance	Levels off / Grows with h
Remedy for a unit root	Difference or cointegrate

IN THE LESSON

- Let's put the two cases side by side.
- A stationary series has a stable mean to which it returns.
- A shock fades away over time, so the effect eventually dies.
- Long horizon forecasts settle at the mean and the forecast error variance levels off.
- A unit root series has no such anchor.
- A shock stays in all future values, forecasts follow the latest value, and the error variance grows with the horizon.
- The remedies are two.
- If the series share a common stochastic trend, model their relationship using cointegration.
- Otherwise, take first differences, which turn the unit root series into a stationary one, and regress the differences.
- Trusting levels regressions with a unit root is the trap.

27 The Dickey-Fuller test

$\Delta y_t = \gamma \times y_{t-1} + \varepsilon_t$, $\gamma = \varphi - 1$
H_0: $\gamma = 0$ (unit root). H_1: $\gamma < 0$ (stationary)

Gamma = phi minus 1	Zero under a unit root
Change in y	Change = $y_t - y_{t-1}$
Critical values	5% DF: −1.94, −2.86, −3.41
One-sided test	Reject if t is below the value
Shock, epsilon	Error term

IN THE LESSON

- How do we check for a unit root?
- We use the Dickey Fuller test.
- Subtract yesterday's value from both sides of the autoregression.
- The change in the series equals gamma times yesterday's value, plus a shock, where gamma is phi minus one.
- The null hypothesis is a unit root, gamma equals zero.
- The alternative is stationarity, gamma below zero, so the test is one sided.
- We take the t ratio of the estimate of gamma, but we do not compare it with normal or ordinary t critical values.
- We use Dickey Fuller critical values, which are more negative.
- At five per cent they are about minus one point nine four with no constant, minus two point eight six with a constant, and minus three point four one with a constant and a trend.

28 Example 9: running the test

Illustration. Constant, no trend. 5% Dickey-Fuller critical value: −2.86

Case 1: $\hat{\gamma} = -0.05$, SE 0.02	$t = -0.05 \div 0.02$	**−2.5**
Decision, Case 1	−2.5 is not < −2.86	**Do not reject**
Case 2: $\hat{\gamma} = -0.12$, SE 0.03	$t = -0.12 \div 0.03$	**−4.0**
Decision, Case 2	−4.0 < −2.86	**Reject**

IN THE LESSON

- Let's run the test twice, using a regression with a constant, where the five per cent critical value is minus two point eight six.
- These are made up estimates.
- In the first case, gamma hat is minus point zero five, with a standard error of point zero two.
- The t ratio is minus point zero five over point zero two, which is minus two point five.
- That is not below minus two point eight six, so we fail to reject the unit root.
- Note that normal critical values would have rejected it, since minus two point five is below minus one point six four five.
- In the second case gamma hat is minus point one two with a standard error of point zero three.
- The t ratio is minus four, below minus two point eight six, so we reject the unit root.

29 The augmented test and its limits

- ADF adds lagged differences: errors uncorrelated
- Lag count chosen by an information criterion
- Same rule: compare t ratio with DF critical values
- Failing to reject proves nothing: power is low

IN THE LESSON

- The simple test assumes the errors are not correlated over time.
- Real data often break that rule.
- The augmented Dickey Fuller test, or A D F, fixes this by adding lagged differences of the series to the regression, so that the errors become serially uncorrelated.
- The number of lags is chosen by an information criterion.
- The decision rule is unchanged: compare the t ratio of gamma with the Dickey Fuller critical values.
- Remember the caution that comes with every such test.
- Failing to reject does not prove a unit root.
- The test has low power, especially for series with a coefficient close to one, so it often fails to reject a series that is in fact stationary.

30 Four mistakes to avoid

- Log-linear β_1 is a growth rate; use exp for level
- Seasonal forecast: target season and T + h, not h
- Random walk SD = σ × √h, not σ × h
- Use Dickey-Fuller values; do not trust R^2 in levels

IN THE LESSON

- Mistake one: reading the log-linear slope as a unit change.
- It is a continuously compounded growth rate, and the level needs the exponential.
- Mistake two: in a seasonal forecast, using the season of the last observation, or plugging in h instead of T plus h.
- Find the target season and the time index first.
- Mistake three: scaling random walk risk by h.
- The standard deviation grows with the square root of h.
- Mistake four: using normal critical values in a unit root test, or trusting a high R squared and a large t statistic from a regression in levels.
- Use Dickey Fuller values, and difference the data.

31 Recap: the main ideas

- Trends: linear, log-linear (growth), quadratic (peak)
- Seasonality: s − 1 dummies; gaps from baseline
- Forecast at T + h, target season; ± multiple × SE
- Random walk: SD = σ√h. Unit root: difference the data
- Dickey-Fuller: H_0 unit root; use DF critical values

IN THE LESSON

- Let's pull it together.
- A linear trend adds a constant amount each period, the log-linear trend has a slope that is the continuous growth rate, and a quadratic trend bends and peaks at minus beta one over two beta two.
- Seasonality is modelled with one fewer dummy than the number of seasons, each coefficient being a difference from the baseline.
- To forecast, find T plus h and the target season.
- An interval forecast uses the regression standard error, and log-linear intervals are exponentiated.
- A random walk has shocks that never fade, with risk growing with the square root of h.
- A unit root causes spurious regressions, and the Dickey Fuller test checks for it.

32 Try three quick questions

1. Random walk, σ = 3. SD of forecast error at h = 9?
   Answer: 3 × √9 = 9, not 3 × 9 = 27.
2. Monthly: y = 100 + 2t + 5·March (other months 0). T = 12 (Dec). March?
   Answer: t = 15, so 100 + 30 + 5 = 135.
3. DF t ratio −3.1 with a constant. Reject unit root?
   Answer: Yes: −3.1 < −2.86 at the 5% level.

IN THE LESSON

- Now try three quick questions yourself.
- Pause after each one.
- Question one.
- A series follows a random walk with a shock volatility of three.
- What is the standard deviation of the forecast error nine steps ahead?
- Three times the square root of nine, which is nine.
- Question two.
- A monthly model has a trend of one hundred plus two t and an extra five in March.
- The last observation is time twelve, in December.
- What is the forecast for March of the next year?
- That is h equal to three, so t is fifteen, giving one hundred thirty plus five, one hundred thirty five.
- Question three.
- A Dickey Fuller t ratio is minus three point one with a constant.
- Do you reject?
- Yes, it is below minus two point eight six.