ARCH, cointegration and choosing a time-series model

Quantitative Methods

1 A story: stormy weeks come in bunches

- Calm days tend to follow calm days
- Stormy days tend to follow stormy days
- Meera's fund model: errors small, then large
- The spread of errors is not constant

IN THE LESSON

- Think about the weather at a seaside town.
- A calm morning is usually followed by another calm morning.
- A stormy day is usually followed by another stormy day.
- The weather is not predictable, but the size of the surprises clusters.
- Markets behave the same way.
- Meera Okafor models the monthly return of a fund, and the model's mistakes, which we call residuals, are small for months, then large for months, then small again.
- The model has an average error, but the spread of the errors is not the same every month.
- That uneven spread is what we are about to name and measure.

2 Surprises come in clusters

Chart: Size of error vs Time (months)

- Calm
- Turbulent
- Calm again

IN THE LESSON

- Here is the idea as a picture.
- The horizontal axis is time, month by month.
- The vertical axis is the size of the model's error that month, ignoring its sign.
- Follow the line from the left.
- For the first few months it stays low, which is a calm spell.
- Then it jumps, and for several months in a row the errors are big.
- That is the turbulent spell.
- Notice that the jump is not a one-off.
- Large errors are followed by large errors.
- Then the line settles back to low again.
- A model that assumes the error spread is the same in every month would miss this pattern completely.

3 Three long words, three short meanings

ARCH, word by word

Word	Meaning
Heteroskedasticity	Error variance not constant
Conditional	Depends on what we know now
Autoregressive	Depends on its own past
ARCH(1)	Depends on last period's error

IN THE LESSON

- The full name is autoregressive conditional heteroskedasticity, which is a mouthful, so everyone says arch.
- Take the words one at a time.
- Heteroskedasticity means the error variance is not constant.
- Conditional means the variance depends on something we know at the time, here what happened last period.
- Autoregressive means that something is the variable's own past, in this case the past squared errors.
- Put together, arch says this: the variance of the error in one period depends on the size of the errors before it.
- That is exactly the clustering we just drew, written as an equation.

4 The ARCH(1) test

$$\varepsilon^2(t) = a0 + a1 \times \varepsilon^2(t-1) + u(t)$$

Squared residual today	Residual from the model, squared
a0: constant	Baseline variance
a1: ARCH slope	Test: is it different from zero?
Squared residual last period	Same thing, one period earlier
Error term	Fresh error of this small model

IN THE LESSON

- Now the test.
- We take the residuals from the original model, such as Meera's return model, and we square each one.
- Squaring matters: it removes the sign and keeps only the size, and size is what clusters.
- Then we run a small regression of each squared residual on the squared residual one period earlier.
- The constant is a zero, the slope is a one, and u is a fresh error.
- The null hypothesis is that a one equals zero, which means the variance does not depend on the past.
- If a one is significantly different from zero, arch is present.
- We judge it with an ordinary t test on a one.

5 Example 1: does Meera's model have ARCH?

Squared residuals on last period's squared residuals

Constant a0	estimated	**0.5**
Slope a1	standard error 0.095	**0.4**
t-statistic	0.4 ÷ 0.095	**4.21**
Critical value, 5%	given as a table value	**1.98**
Decision	4.21 is beyond 1.98	**ARCH present**

IN THE LESSON

- Example one.
- Meera collects the residuals from a monthly return model over about one hundred twenty months, squares them, and runs the test.
- The estimated constant is zero point five.
- The slope a one is zero point four, with a standard error of zero point zero nine five.
- So the t statistic is zero point four divided by zero point zero nine five, about four point two.
- The critical value at the five per cent level is about one point nine eight, given as a table value.
- Four point two is far beyond one point nine eight, so we reject the null.
- Arch is present in this fund's residuals.

6 A significant a1 does two jobs

- Warning: the model's standard errors are unreliable
- Fix: correct them, for example robust standard errors
- Opportunity: the ARCH equation forecasts variance
- Bigger past errors mean a bigger forecast

IN THE LESSON

- A significant a one gives Meera two separate conclusions, and it helps to keep them apart.
- The first is a warning.
- The original return model assumed a constant error variance, and that assumption has failed.
- So the standard errors of that model, and the t statistics built on them, cannot be trusted.
- The fix is to correct for the heteroskedasticity, for example by using robust standard errors.
- The second is an opportunity.
- The arch equation is itself a model of the variance, so it can forecast how volatile the next period will be.
- A warning about the old model, and a forecast from the new one.

7 Forecasting next period's variance

Forecast variance = a0 + a1 × (latest residual)²

Forecast variance	Variance expected next period
a0: constant	Baseline variance
a1: ARCH slope	Weight on the latest surprise
Latest residual, squared	Square it: the sign drops out

IN THE LESSON

- Here is the forecast.
- We take the fitted arch equation and look one period ahead.
- The forecast variance is the constant a zero, plus a one times the square of the latest residual.
- Notice what is not in the formula: the sign.
- A big surprise either way lifts the forecast in the same way.
- The forecast is a variance, so it is in squared units.
- To get a volatility, in the same units as the returns, take the square root.
- Because a one is positive here, a larger latest residual always gives a larger forecast, which is the clustering written as arithmetic.

8 Example 2: next month's volatility

a0 = 0.5, a1 = 0.4, last residual −2.5

Square the residual	$(-2.5)^2$	**6.25**
Weight it	0.4 × 6.25	**2.50**
Forecast variance	0.5 + 2.50	**3.00**
Forecast volatility	√3.00	**1.73%**
After a quiet month	√0.5	**0.71%**

IN THE LESSON

- Example two.
- Meera's fitted equation has a zero of zero point five and a one of zero point four.
- Last month's residual was minus two point five percentage points.
- Square it, and the minus sign vanishes: it is six point two five.
- Now multiply by zero point four to get two point five, and add the constant zero point five.
- The forecast variance is three point zero zero.
- Take the square root, and the forecast volatility is about one point seven three per cent.
- Compare that with a quiet month.
- After a residual of zero, the forecast variance would be only zero point five, a volatility of about zero point seven one per cent.

9 Bigger surprises, bigger forecasts

Forecast variance

Residual 0	**0.5**
Residual ±1	**0.9**
Residual ±2	**2.1**
Residual ±2.5	**3.0**

Illustration with a0 = 0.5 and a1 = 0.4, not real data

IN THE LESSON

- Let us see the forecast as bars.
- The vertical axis is the forecast variance for next month.
- Each bar is a different latest residual, in percentage points, either sign.
- After a residual of zero, the forecast variance is just the constant, zero point five.
- After a residual of one, it is zero point nine.
- After two, it is two point one.
- After two and a half, it is three point zero, six times the quiet level.
- Notice the bars grow faster than the residuals do, because the residual is squared before it counts.
- That is why one big shock moves the forecast so much, and why the next few months are expected to stay stormy.

10 Mistakes to avoid with ARCH

- 1. Forgetting to square the residual
- 2. No autocorrelation does not rule out ARCH
- 3. The forecast is a variance: take the root
- 4. A significant a1 is not only a standard-error fix

IN THE LESSON

- Before we move on, the common slips.
- One.
- Using the residual without squaring it.
- Try it on Meera's numbers: zero point five plus zero point four times minus two point five gives minus zero point five, a negative variance, which cannot exist.
- That is a quick way to catch the error.
- Two.
- Assuming that no autocorrelation in the residuals rules out arch.
- Autocorrelation is about the residuals themselves, while arch is about their squares, so you can have one without the other.
- Three.
- Forgetting that the forecast is a variance, so you must take the square root before quoting a volatility.
- Four.
- Treating a significant a one as only a standard error problem.
- It also means the variance can be forecast, which is the useful part.

11 A story: the walker and the dog

- A walker wanders at random, so does their dog
- Alone, neither path can be predicted
- The lead keeps the gap between them steady
- Two drifting series on a lead: cointegrated

IN THE LESSON

- A walker leaves a pub in the middle of the night.
- They wander in every direction, with no plan at all.
- Their dog is on a long lead, and the dog also wanders.
- Look at either one alone and you cannot predict where it will be in an hour.
- Yet you can say something sure about the two together: they will never be very far apart, because the lead keeps pulling them together.
- Each path is a random drift, but the distance between them is steady.
- Hold on to that picture.
- Two series that drift but are tied by a lead are cointegrated, and the lead is a long-term economic link.

12 A warning: a lovely regression that means nothing

Ice cream sales on a stamp price, 100 months (made up)

Output	Value
R^2	0.91
Slope b1 (SE 0.028)	0.62
t-statistic on slope	22.1
Dickey-Fuller t on residuals	−1.20

IN THE LESSON

- Why do we care?
- Because of a famous trap.
- Meera regresses a town's monthly ice cream sales on the price of a rare stamp.
- Both have drifted upward for years, but they have nothing to do with each other.
- Look at the output.
- The fit is very good, an R squared of zero point nine one.
- The slope is zero point six two, with a standard error of zero point zero two eight, so its t statistic is about twenty two.
- Everything looks wonderful.
- But the residuals themselves wander, and the Dickey-Fuller test on them, minus one point two zero, shows no stationarity.
- The regression is spurious: strong numbers, no real link.

13 Two series that drift apart

Chart: Level of series vs Time

- Series Y
- Series X
- Gap keeps widening

IN THE LESSON

- Here is the trap as a picture.
- The horizontal axis is time.
- The vertical axis is the level of each series.
- Follow the two lines from the left.
- Both start together and both wander upward, so a regression sees them move together and gives a high R squared.
- But look at the gap.
- It does not stay steady.
- As time passes the red line pulls away from the green line, and the gap keeps widening.
- Nothing holds the two together.
- If the gap can grow without limit, the regression errors are not stationary, and the usual tests and standard errors cannot be trusted.

14 Four cases before you regress levels

Test each series for a unit root first

Unit roots found	Regression in levels
Neither series	Fine: ordinary regression
Only one series	Unreliable
Both, not cointegrated	Spurious
Both, cointegrated	Valid long-run link

IN THE LESSON

- This leads to a decision rule.
- Before regressing one series on another, test each one for a unit root.
- Recall from Level One that a series with a unit root is non-stationary, a random walk.
- That gives four cases.
- If neither series has a unit root, ordinary regression is fine.
- If only one has a unit root, the estimates and tests are unreliable, and there is nothing to cointegrate.
- If both have unit roots and they are not cointegrated, the regression is spurious.
- If both have unit roots and they are cointegrated, the regression captures a valid long-run relationship.
- Only that last case rescues a pair of unit-root series.

15 What cointegration really means

- Cointegrated: a long-run link keeps them together
- The regression errors are stationary
- Each series alone may still wander
- Needs a unit root in both series

IN THE LESSON

- Now the definition.
- Two series are cointegrated if a long-term economic relationship keeps them from drifting apart, so that the regression errors, the gap between them, are stationary.
- The word to focus on is stationary.
- Each series alone wanders, but a particular combination of them does not.
- That is the lead in our story.
- Revenue and national output, or a futures price and its spot price, are natural candidates, because economics ties them together.
- Notice also that cointegration needs both series to have unit roots.
- If one of them is already stationary, there is nothing to cointegrate, and the pair fails the test for a different reason.

16 Two series on a lead

Chart: Level of series vs Time

- Series Y
- Series X
- Gap stays steady

IN THE LESSON

- Now the good case, as a picture.
- The axes are the same as before: time along the bottom, the level of each series up the side.
- Follow the green and the red lines.
- Both climb and wobble, and neither could be predicted alone.
- But this time the gap between them stays about the same, however far they travel.
- Whenever one wanders ahead, the long-run link pulls it back.
- The gap is the regression error, so the errors are stationary.
- This is what the Engle-Granger test is looking for: a steady gap between two series that each have a unit root.

17 The Engle-Granger method in four steps

Two series, Y and X

Unit-root test on each →
Regress Y on X in levels →
Dickey-Fuller on residuals →
Beyond critical value?

IN THE LESSON

- The test has four steps.
- Step one: test each series for a unit root.
- If neither has one, stop, because ordinary regression is fine, and if only one has one, stop too.
- Step two, when both have unit roots: regress one series on the other in levels, and keep the residuals, which are the gap between them.
- Step three: apply a Dickey-Fuller test to those residuals, but use Engle-Granger critical values, not the ordinary table.
- Step four: decide.
- If the test rejects a unit root in the residuals, the residuals are stationary and the series are cointegrated.
- If it does not, the regression may be spurious.

18 The test on the residuals

Δε(t) = a + γ × ε(t − 1) + u(t) H0: γ = 0

Change in the residual	This period's residual minus last
Gamma: unit-root slope	0: unit root; below 0: stationary
Last residual	Residual one period earlier
Engle-Granger value	Stricter than the ordinary table

IN THE LESSON

- Here is the Dickey-Fuller regression applied to the residuals.
- We take the change in the residual from one period to the next, and regress it on last period's residual.
- Gamma is the slope.
- If gamma is zero, the residual does not pull back toward zero, which means it has a unit root.
- If gamma is negative, the residual is pulled back, which means it is stationary.
- So the null hypothesis is that gamma equals zero.
- We compute the t statistic on gamma and compare it with the Engle-Granger critical value, which is more negative than the ordinary Dickey-Fuller one.
- A t statistic beyond that value means we reject the unit root.

19 Example 3: revenue and national output

60 quarters, levels; critical values given

Revenue, unit root?	−1.30 vs −2.91	**Not rejected**
GDP, unit root?	−0.85 vs −2.91	**Not rejected**
Residuals, Engle-Granger	−4.10 vs −3.34	**Rejected**
Verdict	gap is stationary	**Cointegrated**

The regression captures a long-run link

IN THE LESSON

- Example three.
- Meera regresses the quarterly revenue of a fictional company, Brightline Telecom, on nominal G D P, over sixty quarters, in levels.
- Step one.
- The Dickey-Fuller statistic for revenue is minus one point three, and for G D P it is minus zero point eight five.
- The ordinary five per cent critical value, given as about minus two point nine one, is not beaten by either, so both have unit roots.
- Step two and three.
- The Dickey-Fuller statistic on the regression residuals is minus four point one zero, against an Engle-Granger critical value, given as about minus three point three four.
- Minus four point one zero is further left, so we reject.
- The two are cointegrated.

20 Example 4: three pairs, three verdicts

Residual test uses the Engle-Granger value, −3.34

Pair A: both unit roots	−4.10 beyond −3.34	**Cointegrated**
Pair B: both unit roots	−3.00 not beyond −3.34	**Spurious risk**
Pair B, wrong table	−3.00 beyond −2.91	**False comfort**
Pair C: one unit root	no test needed	**Unreliable**

IN THE LESSON

- Example four, three more pairs that Meera meets.
- Pair A is the one we just solved: both unit roots, residual statistic minus four point one zero beyond minus three point three four, so cointegrated.
- Pair B has both unit roots, but the residual statistic is minus three point zero zero.
- It would pass the ordinary table value of about minus two point nine one, which is the trap.
- But the Engle-Granger value is minus three point three four, and minus three is not beyond it.
- Not cointegrated, so the regression may be spurious.
- Pair C has a unit root in only one series.
- We stop at step one: the regression is unreliable and cointegration does not apply.

21 Mistakes to avoid with cointegration

- 1. Ordinary table used on the residuals
- 2. Testing the series, not the residuals
- 3. Trusting R^2 on two trending series
- 4. Only one unit root: no cointegration

IN THE LESSON

- The common slips here.
- One.
- Using the ordinary Dickey-Fuller critical value on the residuals.
- The residuals come from a regression fitted to make them as small as possible, so they already look more stationary than they should, and the stricter Engle-Granger value allows for that.
- Two.
- Testing the original series for cointegration, instead of testing the regression residuals.
- Three.
- Believing a high R squared and big t statistics on trending series.
- Four.
- Concluding that cointegration is possible when only one series has a unit root.
- Both must have one.
- And when the pair is cointegrated, remember the regression describes the long-run relationship.

22 A story: the doctor's checklist

- Do not start with the cleverest model
- Plot first, then try the simplest model
- Test the residuals after every step
- Each test result chooses the next step

IN THE LESSON

- A good doctor does not order every test at once.
- They ask what is wrong, look at the patient, try the simplest treatment, and only move on when a result says the simple answer is not enough.
- Choosing a time-series model works the same way.
- Meera has housing starts, a monthly count of new homes begun, and she must forecast it.
- She does not jump to the cleverest model.
- She plots, tries the simplest idea, tests the residuals, and lets each test result decide the next step.
- The tools are the ones we know: trend models, autoregressive models, differencing, seasonal lags, arch and cointegration.

23 Step one: plot housing starts

Chart: Housing starts vs Time (months)

- December dip
- Rising trend

IN THE LESSON

- Always begin with the plot.
- The horizontal axis is time in months, and the vertical axis is housing starts.
- Follow the line from the left.
- It climbs over time, which is a trend.
- But it also dips sharply, and the dips come at the same point each year, every December.
- That is seasonality.
- You also look for a structural break, a sudden permanent shift in the level, and for growth that gets steeper, which would suggest logs.
- Meera sees a trend and a seasonal dip.
- A plain trend model would leave the December pattern in its residuals, so she already expects to need more than a trend.

24 What the plot tells you to try first

Plot first, then choose the first move

What the plot shows	First move
Steady rise	Linear trend model
Steady percentage growth	Log-linear trend model
Residuals serially correlated	Move to an AR model
Unit root in the level	Difference it first

IN THE LESSON

- Here is how the plot steers the first move.
- If the series grows by a steady amount each period, a linear trend model suits.
- If it grows by a steady percentage, so it curves upward, take logs first and fit a log-linear trend.
- If the plot shows no seasonality and no structural break, a trend model is a fair first try, and then you check its residuals for serial correlation.
- If they are correlated, move to an autoregressive model.
- If the series has a unit root, difference it first.
- If there is a structural break, the series needs separate treatment before and after the break.

25 The order of the checks

Start

Plot the data → Trend model if no season or break → Residuals correlated? → Make it stationary

Then

Fit AR(1), test residuals → Add lags until clean → Test for ARCH → Compare out-of-sample RMSE

IN THE LESSON

- Now the whole sequence on one board.
- Top row.
- Understand the question, plot the data, and if there is no seasonality or break, try a trend model.
- Check its residuals for serial correlation.
- If they are correlated, make the series covariance stationary, by differencing, and by taking logs first if it grows exponentially.
- Bottom row.
- Fit an A R one model and test the residual autocorrelations.
- Add lags, or a seasonal lag, until none is significant.
- Then test for arch, and correct if needed.
- Finally compare the candidate models on out-of-sample root mean squared error.
- If two series are involved, check unit roots and cointegration first.

26 Is a residual autocorrelation significant?

t = autocorrelation at lag k ÷ (1 ÷ √T)

t-statistic	Compare with about 1.98
Residual autocorrelation	Residuals now vs k periods ago
Standard error: one over root T	Standard error, T = observations

IN THE LESSON

- One test appears at every stage: is a residual autocorrelation really different from zero?
- If the model is right, its residuals should have no autocorrelation at any lag.
- We take the autocorrelation at lag k, and divide by its standard error, which is one divided by the square root of T, the number of observations.
- The result is a t statistic.
- Compare it with the critical t value, about one point nine eight for a long sample at the five per cent level, given in the example.
- If the t statistic is beyond it, that lag still has something to teach us, and we add that lag to the model.

27 Example 5: housing starts, step by step

120 months; log changes; critical t about 1.98

Plot	trend plus December dip	**Seasonal**
Unit root	Dickey-Fuller not rejected	**Difference logs**
AR(1) residual, lag 12	0.38 ÷ 0.091	**t = 4.2**
Add lag 12	4.2 beyond 1.98	**AR(1), lag 12**
Then ARCH test	a1 not significant	**No correction**

IN THE LESSON

- Example five.
- Meera has one hundred twenty months of housing starts, with a trend and a December dip.
- A Dickey-Fuller test cannot reject a unit root in the level, so she takes logs and differences them, which gives monthly growth rates.
- She fits an A R one model on the growth rates, and checks the residual autocorrelations.
- At lag twelve it is zero point three eight.
- The standard error is one over the square root of one hundred twenty, about zero point zero nine one, so the t statistic is about four point two.
- That is beyond one point nine eight, so she adds lag twelve, the seasonal lag.
- After that no residual autocorrelation is significant, and the arch test on the residuals shows no significant slope.

28 Comparing models: out-of-sample RMSE

RMSE = √(sum of squared errors ÷ n)

RMSE	Typical miss, in the data's units
Forecast errors, squared	Actual minus forecast, squared
n: out-of-sample forecasts	Forecasts not used in fitting

IN THE LESSON

- The last step is to choose between the models that survive the tests.
- The measure is the root mean squared error.
- Take each forecast error, the actual value minus the forecast.
- Square the errors so positives and negatives do not cancel and big misses count heavily.
- Average the squares, then take the square root, which brings the answer back to the original units.
- Crucially, the errors must be out-of-sample, meaning forecasts for periods that were not used to fit the model.
- A model can fit its own past very well and still forecast badly, and out-of-sample errors catch that.
- The model with the lower out-of-sample RMSE is preferred.

29 Example 6: which model forecasts better?

Five held-back months; errors in percentage points

Model A squared errors	36+25+16+36+25	**138**
Model A RMSE	√(138 ÷ 5)	**5.25**
Model B squared errors	16+16+9+25+16	**82**
Model B RMSE	√(82 ÷ 5)	**4.05**
Choice	lower RMSE wins	**Model B**

IN THE LESSON

- Example six.
- Meera keeps back five months and forecasts them with two models, without using those months in the fitting.
- Model A is the A R one on growth rates.
- Model B adds lag twelve.
- Model A's errors, in percentage points, are six, minus five, four, minus six and five.
- Squared, they are thirty six, twenty five, sixteen, thirty six and twenty five, a total of one hundred thirty eight.
- Divide by five to get twenty seven point six, and the square root is about five point two five.
- Model B's errors are smaller, and its squared errors total eighty two.
- Divide by five to get sixteen point four, the root is about four point zero five.
- Model B has the lower RMSE, so Meera keeps it.

30 Match the problem to the model

Situation to model

Situation	Choice
Residuals serially correlated	AR model, stationary data
Squared residuals depend on last	ARCH correction
Two trending series	Unit roots, then cointegration
Several models survive	Lowest out-of-sample RMSE

IN THE LESSON

- To finish, practise matching a situation to a model, because cases are written that way.
- Residuals from a trend model still show serial correlation: use an autoregressive model, after making the series stationary.
- Residuals clean, but squared residuals depend on last period's squared residuals: that is arch, so correct the standard errors and forecast variance with it.
- Two trending series: test each for a unit root, and if both have one, test for cointegration.
- Several candidates survive: pick the lowest out-of-sample RMSE.

31 Mistakes to avoid when choosing a model

- 1. Modelling levels that are not stationary
- 2. Judging by in-sample fit, not RMSE
- 3. Forgetting the ARCH test
- 4. Two series, no unit-root check
- 5. Ignoring a seasonal pattern

IN THE LESSON

- The common slips in model choice.
- One.
- Modelling levels that are not stationary.
- This is the main trap: a trend or an autoregressive model fitted to a series with a unit root gives unreliable results, so difference first, and take logs first if the growth is exponential.
- Two.
- Judging models by how well they fit the past, rather than by out-of-sample RMSE.
- Three.
- Stopping after the mean model, forgetting to test for arch.
- Four.
- Skipping the unit-root and cointegration checks when two series are involved.
- Five.
- Ignoring a seasonal pattern, which shows up as significant residual autocorrelation at the seasonal lag.

32 Recap: three ideas

- ARCH: squared residuals on their own lag
- Forecast variance = a0 + a1 × latest residual²
- Both unit roots: cointegrated, or spurious
- Plot, stationary, AR, ARCH, then RMSE

IN THE LESSON

- Let us pull the lesson together.
- First, arch means the variance of the error depends on the size of past errors.
- Test it by regressing squared residuals on their lag, and a significant slope means arch is present.
- Then fix the standard errors, and forecast next period's variance as the constant plus the slope times the square of the latest residual.
- Second, before regressing one series on another, test each for a unit root.
- If both have one, the regression is valid only if they are cointegrated, tested by Dickey-Fuller on the residuals with Engle-Granger values.
- Third, model choice is a sequence, ending with out-of-sample RMSE.

33 Try three quick questions

1. ARCH(1): a0 = 0.4, a1 = 0.5, latest residual −3. Forecast variance?
   Answer: 0.4 + 0.5 × 9 = 4.9
2. Both unit roots; residual t = −2.80; Engle-Granger −3.34. Verdict?
   Answer: Not beyond −3.34: not cointegrated
3. Exponential growth; trend residuals correlated. Next step?
   Answer: Take logs, difference, fit an AR model

IN THE LESSON

- Now try three quick questions yourself.
- Pause after each one.
- Question one.
- An arch one model has a zero of zero point four and a one of zero point five.
- The latest residual is minus three.
- What is the forecast variance?
- Square the residual to get nine, times zero point five is four point five, plus zero point four is four point nine.
- Question two.
- Both series have unit roots, and the Dickey-Fuller statistic on the residuals is minus two point eight, against an Engle-Granger value of minus three point three four.
- Not beyond it, so not cointegrated, and the regression may be spurious.
- Question three.
- Sales grow exponentially and a trend model leaves serially correlated residuals.
- What next?
- Take logs, difference, then fit an A R model.
- Well done.