CHARTA LEARNING · LESSON NOTES

Random walks, unit roots and seasonality

Quantitative Methods

1 A story: the wandering boat

- A boat drifts with no anchor
- Each gust of wind is a random push
- Old pushes never wear off
- Best guess for later: where it is now

IN THE LESSON

- Picture a small boat on a wide lake with no anchor.
- Each minute a gust pushes it a little left or right, and nobody can say which way the next gust will blow.
- Where will the boat be in an hour?
- The best guess is simply where it is now, because the gusts have no pattern.
- But your uncertainty keeps growing, because every old push stays in the boat's position for ever.
- Nothing drags it back to a home spot.
- A great many prices behave like this boat, including share prices and exchange rates.
- Analysts call such a series a random walk, and it breaks the usual forecasting tools in a very particular way.

2 The random walk, written down

x(t) = x(t − 1) + error
With drift: expected x(t + k) = x(t) + k × b0

x now: value this period	The value we have now
x last period	Last period's value
Error: the surprise	Average zero, no pattern
b zero: drift or intercept	Steady push each period
k: periods ahead	How far ahead we look

IN THE LESSON

- Here is the boat as a formula.
- The value this period equals the value last period, plus an error.
- The error is the unpredictable gust, with an average of zero.
- That is a random walk.
- If we add a constant, called drift, to every step, the boat has a steady current as well as gusts, and we call it a random walk with drift.
- With drift, the expected value k periods ahead is today's value plus k times the drift.
- Without drift, the expected value at every horizon is just today's value.
- Notice what has happened to the lag coefficient.
- In an autoregressive model it is called b one.
- Here b one is exactly one.

3 Example 1: forecasting a random walk

Harbour Index = 480, drift 2 a month, error s.d. 6 a month

Forecast, no drift	today's value	**480**
Forecast with drift	480 + 6 × 2	**492**
Error variance, 6 months	6×6^2	**216**
Error s.d., 6 months	√216	**14.70**

IN THE LESSON

- Meet our analyst, Marta Ilves, who works at a made-up firm called Larkfield Capital.
- She follows the Harbour Index, a fictional share index, which stands at four hundred and eighty.
- She models it as a random walk.
- First, with no drift, her best forecast six months ahead is four hundred and eighty, today's value.
- Now suppose the index drifts up two points a month on average.
- Six months of drift add twelve, so the forecast is four hundred and ninety two.
- The monthly error has a standard deviation of six.
- Errors add up, so the variance over six months is six times thirty six, which is two hundred and sixteen.
- The standard deviation is the square root, about fourteen point seven.

4 Uncertainty that never settles

Chart: Error variance vs Forecast horizon (periods)

- Random walk: no ceiling
- AR(1), b1 = 0.5
- Ceiling 1.33

IN THE LESSON

- Let's draw the idea.
- Along the bottom is the forecast horizon, how many periods ahead we look.
- Up the side is the variance of our forecast error, with a surprise of size one each period.
- First, the rising line.
- For a random walk the variance after k periods is k times the one period variance, so the line climbs straight up and never stops.
- Now the second line, for a stationary A R one model with b one of one half.
- Surprises fade there, so the line bends over and levels off at about one point three three, shown by the grey line.
- That ceiling is what a stable model looks like.
- The climbing line has no ceiling.

5 Random walk versus covariance stationary

Random walk: does it meet each requirement?

Requirement	Random walk
Constant mean	No
Constant variance	No, grows
Mean reverting level	None
Best forecast	Today

IN THE LESSON

- Now let's pin down why a random walk is not covariance stationary.
- A covariance stationary series needs three things: a constant mean, a constant variance, and a covariance between values that depends only on how far apart they are, not on the date.
- A random walk fails them.
- Its expected value can wander if there is drift, and it has no level it returns to.
- Its variance grows with time, as we just drew.
- The mean reverting level of an A R one model is b zero divided by one minus b one, and with b one equal to one that sum divides by zero, so no such level exists.
- Remember the table as a quick check on the exam.

6 Example 2: a stable model, then b1 = 1

Stable model: x(t) = 2 + 0.8 × x(t − 1)

Mean reverting level	2 ÷ (1 − 0.8)	**10**
Forecast from 20	2 + 0.8 × 20	**18**
Two steps ahead	2 + 0.8 × 18	**16.4**
If b1 were 1	2 ÷ (1 − 1)	**Undefined**

IN THE LESSON

- Let's see the contrast with numbers.
- Suppose Marta fits a different series, with the model x of t equals two plus zero point eight times x of t minus one.
- The slope is below one in size, so the series is covariance stationary.
- Its mean reverting level is two divided by one minus zero point eight, which is two divided by zero point two, so ten.
- If the series is now at twenty, the next forecast is two plus zero point eight times twenty, eighteen.
- The one after is two plus zero point eight times eighteen, sixteen point four.
- The forecasts glide back toward ten.
- Now set b one to one.
- The level formula becomes two divided by zero.
- It is undefined.
- That is the point of the lesson so far.

7 The trap with random walks

- b1 near 1 does not give a real long run target
- Check for a unit root before quoting b0 ÷ (1 − b1)
- Random walk forecast = latest value, not the average
- Variance grows with the horizon: no stable variance

IN THE LESSON

- A common exam slip goes like this.
- A vignette shows an A R one fit to an exchange rate, with b zero of zero point zero zero six and b one of zero point nine nine three.
- A student computes the mean reverting level, b zero over one minus b one, and announces a long run target.
- But zero point nine nine three is an estimate, with a standard error.
- It may well be statistically no different from one.
- And if the true value is one, there is no target at all.
- The tidy number you computed is meaningless.
- A second slip is saying a random walk forecast is the same as the average of past values.
- It is not.
- The best forecast is the latest value.

8 What a unit root is, and where it appears

- Unit root: b1 equals 1 in the AR(1) model
- Same thing as a random walk, possibly with drift
- Common in prices, rates and index levels
- Changes (returns) usually have no unit root

IN THE LESSON

- Now the second idea.
- A series has a unit root when the lag coefficient b one in its A R one model equals one.
- That is exactly the random walk.
- The phrase comes from the algebra of the model, where one is the root that causes the trouble.
- Unit roots are common in series that are levels with no fixed anchor: share prices, exchange rates, price indexes, and the level of an economy's output.
- People who study these see it again and again.
- Yields and spreads sometimes behave the same way.
- By contrast, changes in those levels, such as returns, usually have no unit root at all.
- The level wanders.
- The change from period to period is steady.

9 Why a unit root causes trouble

- Estimates and t-tests on b1 are not valid
- The mean reverting level does not exist
- Two wandering series can look related by chance
- Detect it, then transform the series

IN THE LESSON

- Why does it matter?
- Because the ordinary tools assume the series is covariance stationary, and a unit root series is not.
- Three things go wrong.
- First, the estimate of b one is biased, and the usual t tests on it cannot be trusted, so you will make wrong decisions with confidence.
- Second, the mean reverting level and the long horizon forecasts have no meaning.
- Third, if you regress one wandering series on another, you often get a high R squared and a very large t statistic for no reason at all, a false relationship.
- The cure is not to hope.
- It is to detect the unit root, and then change the series so that it is stationary.

10 The fix: first differencing

y(t) = x(t) − x(t − 1)
Random walk: y(t) = error

y: the change in x	The change, a new series
x now: value this period	This period's level
x last period	Last period's level
Error: the surprise	What is left: stationary

IN THE LESSON

- The fix is called first differencing.
- Instead of modelling the level x, we model the change, which we will call y.
- The change y in period t is x of t minus x of t minus one.
- If x is a random walk, then the change is just the error, which has a constant mean and variance and no pattern.
- That is covariance stationary.
- So the changes can be modelled with an ordinary A R model, even when the levels cannot.
- This works when the series has a single unit root.
- If the changes still have a unit root, you would difference once more, but one round is usually enough.

11 From a level to a forecast and back

Four boxes, in order

Test the level → Take changes →
Fit AR to changes → Add to latest level

IN THE LESSON

- Here is the whole routine as a flow, read from left to right.
- First you test the level for a unit root with the Dickey-Fuller test, which we meet shortly.
- If it has one, you move to the second box and take the change, this period minus last period.
- Third, you fit an A R model to those changes.
- Everything that model knows is about changes, so its answer is a forecast change.
- Fourth, and this is where people slip, you rebuild the level.
- You add the forecast change to the latest observed level.
- If you skip that last box, you will report a forecast of about one point two as if it were the index.

12 Example 3: taking changes of the index

Index: 300.0, 304.5, 307.1, 312.8, 315.0

Change, quarter 2	304.5 − 300.0	**4.5**
Change, quarter 3	307.1 − 304.5	**2.6**
Change, quarter 4	312.8 − 307.1	**5.7**
Change, quarter 5	315.0 − 312.8	**2.2**
Average change	15.0 ÷ 4	**3.75**

IN THE LESSON

- Let's do it with data.
- Marta has five quarterly readings of the Harbour Index: three hundred, three hundred and four point five, three hundred and seven point one, three hundred and twelve point eight, and three hundred and fifteen.
- Differencing is subtraction.
- The first change is four point five.
- The second is two point six.
- The third is five point seven.
- The fourth is two point two.
- Four numbers come out of five, because the first level has no earlier level to subtract.
- The average change is the sum, fifteen, divided by four, which is three point seven five.
- That equals the total rise of fifteen over four steps.
- The levels climb without end, but the changes are a steady, well behaved series.

13 Example 4: forecast the changes, then the level

Model of changes: y(t) = 0.60 + 0.30 × y(t − 1)

Next change	0.60 + 0.30 × 2.2	**1.26**
Next level	315.0 + 1.26	**316.26**
Change after that	0.60 + 0.30 × 1.26	**0.978**
Level after that	316.26 + 0.978	**317.238**
The slip	0.60 + 0.30 × 315.0	**95.1**

IN THE LESSON

- Now suppose a Dickey-Fuller test on the index level cannot reject a unit root, so Marta fits an A R one model to a long history of quarterly changes.
- It reads: change equals zero point six plus zero point three times last quarter's change.
- The latest level is three hundred and fifteen and the latest change is two point two.
- The next change is zero point six plus zero point three times two point two, which is one point two six.
- The next level is three hundred and fifteen plus one point two six, three hundred and sixteen point two six.
- The change after that is zero point six plus zero point three times one point two six, zero point nine seven eight, giving a level of three hundred and seventeen point two three eight.
- The slip: feeding the level three hundred and fifteen into the model gives ninety five, nonsense.

14 How do we test for a unit root?

- Idea: fit AR(1), test whether b1 equals 1
- But the usual t-distribution is wrong under a unit root
- Software's t on b1 tests zero, not one
- Solution: the Dickey-Fuller test

IN THE LESSON

- So how do we find out whether a series has a unit root?
- The obvious idea is to fit the A R one model and test whether b one differs from one.
- It does not work.
- When the true b one is one, the usual t statistic does not follow the t distribution at all, so the usual critical values give the wrong answer.
- A second trap is the t statistic that software prints beside b one.
- That tests whether b one is different from zero, which is the wrong question, and with b one near one it is gigantic, even for a random walk.
- We need a test with its own table.
- That test is named after two statisticians, Dickey and Fuller.

15 The Dickey-Fuller rewrite

x(t) − x(t − 1) = b0 + g1 × x(t − 1) + error
g1 = b1 − 1

Change in x	Change in x this period
b zero: drift or intercept	Intercept
g one: b one minus one	Zero if there is a unit root
x last period	Last period's level
b one: lag coefficient	The AR(1) lag coefficient

IN THE LESSON

- The trick is a small piece of algebra.
- Start with the A R one model, x of t equals b zero plus b one times x of t minus one plus error.
- Subtract x of t minus one from both sides.
- On the left you get the change in x.
- On the right, the lagged value now has the coefficient b one minus one, which we call g one.
- So the change in x equals b zero plus g one times last period's x, plus error.
- A unit root, b one equal to one, becomes g one equal to zero.
- A stationary series, b one below one, becomes g one below zero.
- We have turned the question into whether one coefficient is zero.

16 The four stages of the test

Stages of the Dickey-Fuller test

Rewrite as a change → Estimate g1 and its t → Compare with DF value → Act on the verdict

IN THE LESSON

- Here are the stages in order.
- Stage one, rewrite: subtract last period's value to get the change on the right, as we just did.
- Stage two, estimate: run the regression of the change on last period's level, and read off g one and its t statistic.
- Stage three, compare: set that t statistic against the Dickey-Fuller critical value, never against the ordinary t table.
- Stage four, act.
- If you reject, the series is stationary and you can model it directly.
- If you cannot reject, there is a unit root, and you difference the series and then fit the A R model to the changes.
- The test is the gate in front of the model.

17 Null, alternative and the decision rule

- Null: unit root, g1 = 0
- Alternative: stationary, g1 < 0
- Reject only if t is below the DF critical value
- A negative g1 estimate alone proves nothing

IN THE LESSON

- Let's be exact about the hypotheses.
- The null hypothesis is that there is a unit root: g one equals zero.
- The alternative is that the series is stationary: g one is below zero.
- Because the alternative is one sided, we only care about a t statistic that is very negative.
- The decision rule: if the t statistic is more negative than the Dickey-Fuller critical value, reject the null and treat the series as stationary.
- If not, you fail to reject, and you treat it as having a unit root.
- Notice the tone of failing to reject.
- The test starts by assuming a unit root, and only a strong result moves you off it.
- A negative g one estimate, by itself, proves nothing.

18 Dickey-Fuller values are further out

5% cut-off, size (all negative)

Ordinary t	**1.65**
DF, intercept	**2.86**
DF, intercept + trend	**3.41**

Bigger bar means a harder test to pass

IN THE LESSON

- A picture of the critical values makes the point.
- Each bar is the size of the five per cent cut off, though every real value is negative.
- The first bar is the ordinary one sided t value, about one point six five.
- The second is the Dickey-Fuller value for a test with an intercept, about two point eight six.
- The third is the value when a time trend is also included, about three point four one.
- These are example table values for this lesson, so use the figures a question gives you.
- See how much further out the Dickey-Fuller bars are.
- A t statistic of minus two would clear the ordinary cut off but fall short of the Dickey-Fuller one, so the ordinary table rejects a unit root too often.

19 Example 5: testing the index level

Regression of the change in index on lagged level

Estimate of g1	from the regression	**−0.018**
Standard error	from the regression	**0.010**
t-statistic	−0.018 ÷ 0.010	**−1.80**
Dickey-Fuller 5% value	intercept version	**−2.86**
Decision	−1.80 is not below −2.86	**Unit root**

IN THE LESSON

- Now run the test.
- Marta regresses the quarterly change in the Harbour Index on the previous level, with an intercept, using one hundred and twenty quarters.
- The estimate of g one is minus zero point zero one eight, with a standard error of zero point zero one zero.
- The t statistic is minus zero point zero one eight divided by zero point zero one zero, which is minus one point eight.
- The Dickey-Fuller five per cent critical value for this version is minus two point eight six.
- Minus one point eight is not below minus two point eight six, so she cannot reject the unit root.
- Notice that g one is negative, and b one is zero point nine eight two, below one.
- That alone changes nothing.
- Using the ordinary value of minus one point six five would have rejected wrongly.

20 Example 6: test the changes as well

Same cut-off in both tests: minus 2.86

Series tested	t-statistic
Index level	−1.80
Change in the index	−5.30
Cut-off to beat	−2.86

IN THE LESSON

- Having failed to reject for the level, Marta differences the index and runs the same test on the changes.
- This time the t statistic on g one is minus five point three.
- Compare it with minus two point eight six.
- It is far below the cut off, so she rejects the unit root for the changes.
- The changes are covariance stationary, and the original level has exactly one unit root.
- That is the pattern to expect: the level fails the test, the first difference passes it.
- If the changes had also failed, she would difference a second time and test again.
- Only then is it safe to fit an A R model and read its coefficients and its tests.

21 Versions of the test, and the link to AR models

- Versions: intercept, trend, extra lags
- Each version has its own critical values
- The test is the AR(1) model, rewritten
- Difference, fit AR, then check the residuals

IN THE LESSON

- A few practical points tie this back to autoregressive models.
- First, the test comes in versions.
- One has just an intercept.
- One adds a time trend, for series that climb steadily, and its critical value is further out.
- A third adds extra lagged changes on the right to soak up autocorrelation, and is called the augmented test.
- You pick the version that fits the series, and use its own table.
- Second, the test is an A R one model in disguise.
- It is the same regression, only rewritten.
- Third, once you have differenced, you build the A R model on the changes, then check its residual autocorrelations, as you would for any model, before you forecast with it.

22 Mistakes to avoid so far

- 1. b1 of 0.99 treated as clearly below 1
- 2. Negative g1 read as proof of stationarity
- 3. Ordinary t table used for the DF test
- 4. Level modelled despite an unrejected unit root
- 5. Change forecast not added back to the level

IN THE LESSON

- Time to collect the slips from the first half.
- One.
- Treating a b one of zero point nine nine as safely below one.
- Only the test can say if it differs from one, and the usual t table is the wrong one.
- Two.
- Looking at a negative g one and calling the series stationary.
- Three.
- Comparing the Dickey-Fuller t statistic with the ordinary critical value of about minus one point six five.
- Four.
- Modelling the level with an A R model once the unit root has not been rejected, or computing a mean reverting level for a series that may be a random walk.
- And five.
- Forecasting the change and then forgetting to add it back to the latest level.

23 A story: the garden centre

- Seasonality: a pattern that repeats each year
- Fernlea sales jump every spring
- The AR(1) uses the last quarter, not the last year
- The leftover pattern shows in the residuals

IN THE LESSON

- Now the third idea.
- Imagine a chain of garden centres.
- Every spring, sales leap, as people buy seeds, pots and tools.
- Every winter, sales sag.
- That pattern repeats every year, whatever else is going on.
- We call it seasonality: a regular movement tied to a time of year.
- Our analyst, Marta, studies a made up chain called Fernlea Garden Centres, and models quarterly sales growth with a simple A R one model.
- It uses last quarter's growth to predict this quarter's.
- But last quarter is the wrong guide to the spring jump.
- The spring jump looks like the one a year ago.
- The A R one model cannot see that, and its misses show it.

24 Spotting it: the residual autocorrelations

Residual autocorrelation

Lag 1	**0.07**
Lag 2	**0.03**
Lag 3	**0.09**
Lag 4	**0.45**

Only the seasonal lag stands out

IN THE LESSON

- Here is how you detect it.
- After fitting the A R one model, you compute the autocorrelations of the residuals, at lag one, lag two, and so on.
- If the model has caught everything, they are all close to zero.
- Here are one hundred quarters of Fernlea residuals.
- At lag one the autocorrelation is only zero point zero seven.
- At lag two, zero point zero three.
- At lag three, zero point zero nine.
- These are tiny.
- Then look at lag four.
- It is zero point four five, far taller than the rest.
- With quarterly data, lag four is the same quarter in the previous year.
- A spike at the seasonal lag is the fingerprint of seasonality.
- For monthly data, we would look for the spike at lag twelve.

25 Example 7: is the spike real?

Fernlea: 100 quarters, residuals of the AR(1)

Standard error	1 ÷ √100	**0.10**
Lag 1 t-statistic	0.07 ÷ 0.10	**0.7**
Lag 3 t-statistic	0.09 ÷ 0.10	**0.9**
Lag 4 t-statistic	0.45 ÷ 0.10	**4.5**
Cut-off at five per cent	about 2	**Lag 4 real**

IN THE LESSON

- A tall bar is only worth acting on if it is big compared with noise.
- The standard error of each residual autocorrelation is one divided by the square root of the number of observations.
- With one hundred quarters, that is one over ten, zero point one.
- Divide each autocorrelation by it to get a t statistic.
- Lag one is zero point seven.
- Lag three is zero point nine.
- Both far inside plus or minus two.
- Lag four is zero point four five divided by zero point one, four point five.
- That is well beyond two, so it is significant.
- The model is misspecified in the way seasonality predicts, and the fix goes at lag four.
- Only the seasonal lag fails, so we add one lag, not a pile of them.

26 The seasonal fix: add the lag

x(t) = b0 + b1 × x(t − 1) + b2 × x(t − 4) + error

x now: value this period	Growth this quarter
x last period	Growth last quarter
x four periods ago	Same quarter a year ago
b two: seasonal weight	Weight on the seasonal lag
Error: the surprise	Check: autocorrelations near 0

IN THE LESSON

- The cure is to add the seasonal lag as a second regressor.
- For quarterly data the model becomes x of t equals b zero plus b one times x of t minus one, plus b two times x of t minus four, plus error.
- The first lag is still there, catching the short run momentum.
- The new term brings in the same quarter a year earlier, with its own weight, b two.
- For monthly data you would use lag twelve instead of four.
- After you fit it, check the residual autocorrelations again.
- If the spike at the seasonal lag has gone and no other lag is significant, the model has done its job.

27 Example 8: forecasting with a seasonal lag

x(t) = 0.6 + 0.3 × x(t − 1) + 0.5 × x(t − 4), growth in %

Last quarter: Q4 2026	the first lag	**2.0**
Same quarter: Q1 2026	the seasonal lag	**7.0**
Forecast, Q1 2027	0.6 + 0.3 × 2.0 + 0.5 × 7.0	**4.7**
Forecast, Q2 2027	0.6 + 0.3 × 4.7 + 0.5 × 3.0	**3.51**

IN THE LESSON

- Marta refits the model with the lag included: growth equals zero point six, plus zero point three times last quarter's growth, plus zero point five times the growth four quarters ago.
- Fernlea's growth in twenty twenty six was seven per cent in the first quarter, three in the second, minus one in the third and two in the fourth.
- To forecast the first quarter of twenty twenty seven, the latest quarter is the fourth of twenty twenty six, two.
- The same quarter a year earlier is the first of twenty twenty six, seven.
- The forecast is zero point six, plus zero point six, plus three point five, which is four point seven per cent.
- For the second quarter, last quarter becomes the forecast four point seven, and the seasonal lag is three, giving three point five one.

28 The counting trap, and checking the fix

- Quarterly: seasonal lag is 4. Monthly: 12
- Q1 2027 pairs with Q1 2026, not Q2 or Q4
- Miscount gives 2.7 instead of 4.7
- Keep the first lag, then recheck the residuals

IN THE LESSON

- The usual slip is miscounting the seasonal lag.
- Four quarters before the first quarter of twenty twenty seven is the first quarter of twenty twenty six.
- Not the second, not the fourth.
- Had Marta used the second quarter of twenty twenty six, the forecast would have been zero point six plus zero point six plus one point five, two point seven, which is two points too low.
- Lag four is for quarterly data and lag twelve for monthly data, because they are one year back.
- And adding the seasonal lag does not replace the first lag.
- It sits beside it.
- Finally, re-run the residual check after the change.
- A fix you have not tested is only a hope.

29 Recap: the formulas to remember

Random walk: x(t) = x(t − 1) + error (b1 = 1)
Change: y(t) = x(t) − x(t − 1) DF: g1 = b1 − 1
Seasonal: x(t) = b0 + b1x(t − 1) + b2x(t − 4)

b one: lag coefficient	b1 = 1 means a unit root
y: the change in x	Difference to remove it
g one: b one minus one	Null: g1 = 0, unit root
x four periods ago	Seasonal lag: 4 or 12
Rebuilt level	Latest level + forecast change

IN THE LESSON

- Let's pull the lesson together.
- A random walk is this period's value equals last period's plus an error, so b one is one, the variance grows with time, and the best forecast is today's value plus any drift times the horizon.
- A unit root means b one is one, and the fix is the change, this period minus last.
- The Dickey-Fuller test rewrites the model with g one equal to b one minus one.
- Null: g one is zero, a unit root.
- Reject only if the t statistic beats the Dickey-Fuller value.
- Last, seasonality shows up as a spike at lag four or twelve, and the cure is adding that lag as a regressor.

30 Try three quick questions

1. Random walk, drift 0.5, now 40. Expected value 8 periods on?
   Answer: 40 + 8 × 0.5 = 44
2. DF t = −2.1, critical value −2.86. Conclusion?
   Answer: Cannot reject: unit root. Difference it.
3. x = 1 + 0.4·x(t−1) + 0.5·x(t−4); lags are 2 and 6. Forecast?
   Answer: 1 + 0.8 + 3.0 = 4.8

IN THE LESSON

- Now three quick questions.
- Pause after each one.
- Question one.
- A random walk with drift of zero point five a period stands at forty.
- What is the expected value eight periods ahead?
- Eight times zero point five is four, so forty four.
- Question two.
- A Dickey-Fuller t statistic is minus two point one, and the five per cent critical value is minus two point eight six.
- What do you conclude?
- Minus two point one is not below minus two point eight six, so you cannot reject the unit root, and you difference the series.
- Question three.
- The model is x equals one, plus zero point four times x lag one, plus zero point five times x lag four.
- The latest quarter is two and the same quarter last year is six.
- The forecast is one, plus zero point eight, plus three, which is four point eight.