CHARTA LEARNING · LESSON NOTES

Empirical Properties of Correlation

FRM Part II

1 A story: diversification that vanished

- Kestrel: $200m book of 20 stocks, many industries
- Calm years: stocks move fairly independently
- Sell-off: almost all fall on the same days
- Losses far beyond what the risk model said
- Cause: correlation was treated as a fixed number

IN THE LESSON

- Kestrel runs a two hundred million dollar book of twenty stocks spread across many industries.
- In calm years the stocks move fairly independently.
- One firm's bad news is offset by another's good news, and the book's risk looks modest.
- Then a sell-off arrives.
- Almost every stock falls on the same days, and the book loses far more than the risk model said it could.
- Nothing was wrong with the stocks one by one.
- What changed was correlation, the tendency of two returns to move together.
- The model had treated correlation as a fixed number, estimated in good times.
- This lesson is about how correlations really behave.

2 What we measure

- Correlation: from −1 to +1, how returns move together
- 20 stocks: 20 × 19 ÷ 2 = 190 pairs
- Average pairwise correlation: the mean over all pairs
- Correlation volatility: SD of changes in correlation
- Level = how strongly; volatility = how unstable

IN THE LESSON

- First, some terms.
- The correlation between two stocks is a number between minus one and plus one that measures how closely their returns move together.
- With twenty stocks there are twenty times nineteen divided by two, which is one hundred and ninety pairs.
- So Kestrel summarises them with the average pairwise correlation, the simple average over all the pairs.
- It estimates this each month from recent daily returns, which gives it a time series of correlations.
- Correlation volatility is the standard deviation of the changes in that series over time.
- Keep the two apart.
- The level tells you how strongly stocks move together now.
- The volatility tells you how unstable that level is.

3 Correlation and the state of the economy

Equity correlations depend on the state of the economy

	Expansion	Recession
Average correlation	Lowest	Highest
Correlation volatility	Lower	Higher
Diversification	Works well	Weakens

IN THE LESSON

- Equity correlations depend on the state of the economy.
- That is the central empirical finding, and the table sets it out.
- In an expansion, when output and profits are growing, average correlations between stocks are at their lowest.
- In a recession or a sharp sell-off, they are at their highest.
- Correlation volatility moves the same way: it is lower in expansions and higher in recessions.
- So in stress, correlation is both higher and less stable.
- The consequence for diversification is in the last row.
- Diversification works well in an expansion, and it weakens in a recession, which is exactly when it is needed most.

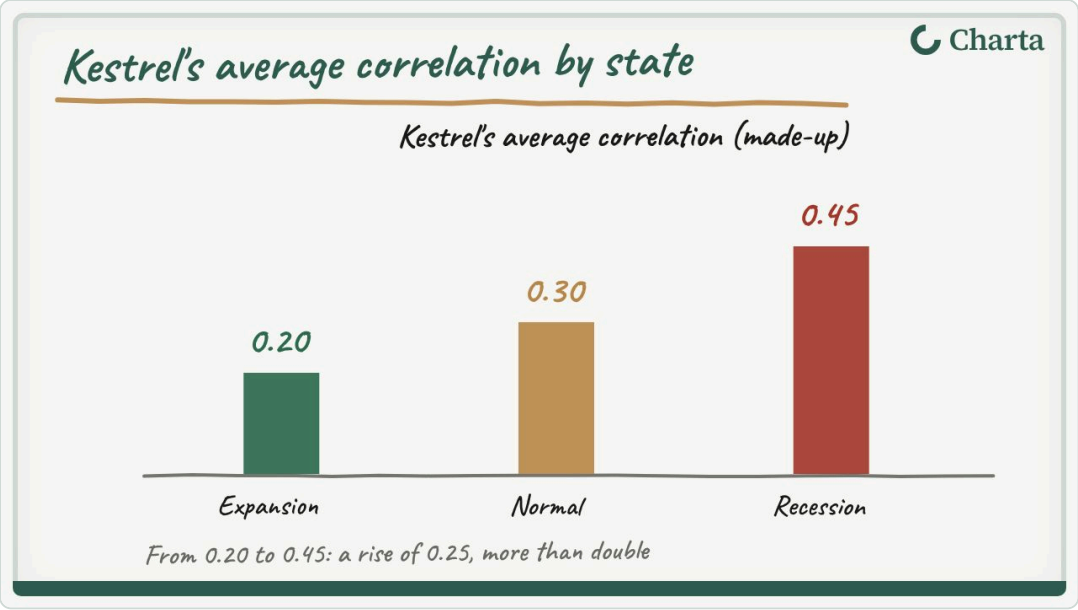
4 Why correlations rise in bad times

- Return = systematic part + firm-specific part
- Systematic: economy, rates, risk appetite hit all firms
- Expansion: firm-specific news dominates
- Recession: common factors swamp firm news
- Big shocks come and go, so correlation jumps more

IN THE LESSON

- Why does this happen?
- Think of each stock's return as two parts.
- One part comes from systematic factors, which are forces that hit every firm at once, such as the economy, interest rates and investors' appetite for risk.
- The other part is firm specific news, such as a new product or a lost contract.
- In an expansion, firm specific news dominates, so stocks go their own way.
- In a recession or a sell-off, the systematic factors become large and swamp the firm specific news.
- Every stock is pushed the same way, so correlations rise.
- And because those big common shocks come and go abruptly, the correlation itself jumps around more.

5 Kestrel's average correlation by state



IN THE LESSON

- Here is what Kestrel's risk team finds in its own history when it sorts the months into three states of the economy.
- These are made up numbers for our example.
- In expansion months, the average pairwise correlation of its stocks is point two zero.
- In normal months, it is point three zero.
- In recession months, it is point four five.
- The level rises steadily as the economy weakens.
- From point two zero to point four five is a rise of point two five, and the recession level is more than double the expansion level.
- Keep the normal level of point three zero in mind.
- It will reappear later as the long run level that Kestrel's correlation reverts to.

6 Example 1: measuring correlation volatility

Five month-end readings of Kestrel's average correlation in a calm year

Calm readings	0.19, 0.22, 0.21, 0.22, 0.19	**4 changes**
Calm changes	+0.03, −0.01, +0.01, −0.03	**mean 0**
Calm: SD of changes	√((9+1+1+9)÷4) hundredths	**0.022**

IN THE LESSON

- In a calm year, five month end readings of Kestrel's average correlation are point one nine, point two two, point two one, point two two and point one nine.
- That gives four changes: plus point zero three, minus point zero one, plus point zero one and minus point zero three.
- They average zero, so to find their standard deviation, square each change, average the squares and take the root.
- Working in hundredths, nine plus one plus one plus nine is twenty, and divided by four that is five.
- The root of five is about two point two four, so the correlation volatility in the calm year is about point zero two two.

7 Example 1: the same measure in stress

Five month-end readings of Kestrel's average correlation in a stress year

Stress readings	0.37, 0.52, 0.47, 0.52, 0.37	**±0.15, ±0.05**
Stress: SD of changes	√((225+25+25+225)÷4) hundredths	**0.112**
Stress versus calm	√125 ÷ √5 = √25	**5 times**

Stress: a higher level and 5 times the correlation volatility

IN THE LESSON

- Now a stress year.
- The readings are point three seven, point five two, point four seven, point five two and point three seven, with changes of plus or minus point one five and plus or minus point zero five.
- In hundredths, two hundred and twenty five plus twenty five plus twenty five plus two hundred and twenty five is five hundred.
- Over four that is one hundred and twenty five, with a root of about eleven point one eight, so the volatility is about point one one two.
- Compared with the calm year, that is the root of one hundred and twenty five over the root of five, the root of twenty five, so five times as large.
- Here we divide by the number of changes, four.
- Dividing by three instead gives slightly larger figures but the same five to one ratio.
- Stress brings a higher level and five times the correlation volatility.

8 Correlation and the risk of a stock book

$\sigma_p = \sigma \times \sqrt{\left(\frac{1}{n} + \left(1 - \frac{1}{n}\right) \times \bar{\rho}\right)}$
VaR = z × σp × book value

Volatility of each stock, σ	Same for every stock
Number of stocks, n	Held in equal amounts
Average correlation, $\bar{\rho}$	Mean over all pairs
Confidence multiple, z	2.33 at 99% confidence

IN THE LESSON

- To see what correlation does to risk, use a standard result.
- Suppose every stock has the same volatility, sigma, and the book holds n stocks in equal amounts.
- Then the portfolio volatility, sigma p, is sigma times the square root of one over n, plus one minus one over n, times the average correlation, rho bar.
- The first term shrinks as you add stocks.
- That is diversification at work.
- The second term does not shrink, because it is driven by correlation.
- Assuming normal returns, value at risk, V a R, is then z times sigma p times the value of the book, where z is two point three three at ninety nine per cent confidence.
- So a higher average correlation feeds straight into a higher V a R.

9 Example 2: Kestrel's VaR, calm versus stressed

$200m, 20 equal stocks, σ = 30% a year each. 1-year 99% VaR, z = 2.33. $\bar{\rho}$ = 0.20 or 0.45

Calm: σp	30% × √(0.05 + 0.95 × 0.20)	**14.70%**
Calm: VaR	2.33 × 14.70% × $200m	**$68.5m**
Stress: σp	30% × √(0.05 + 0.95 × 0.45)	**20.73%**
Stress: VaR	2.33 × 20.73% × $200m	**$96.6m**
Increase	96.6 − 68.5 = 28.1	**+41%**

IN THE LESSON

- For Kestrel: two hundred million dollars, twenty stocks, each with thirty per cent volatility a year, and one year V a R at ninety nine per cent, with z of two point three three.
- At the calm correlation of point two zero, the inside term is point zero five plus point nine five times point two zero, point two four.
- So sigma p is thirty per cent times the root of point two four, fourteen point seven per cent.
- Then V a R is two point three three times fourteen point seven per cent times two hundred million, or sixty eight point five million.
- With the recession correlation of point four five, the inside term is point zero five plus point nine five times point four five, point four seven seven five.
- Sigma p becomes twenty point seven three per cent, and V a R, two point three three times that times two hundred million, is ninety six point six million.
- That is twenty eight point one million, or forty one per cent, more.

10 Why it matters: stressed correlation inputs

- Calm-period matrix understates downturn risk
- Kestrel: $68.5m on calm inputs, $96.6m on stressed
- VaR needs stressed correlation inputs
- Hedge ratios too: they shift when correlation jumps
- Correlations mean-revert after a spike

IN THE LESSON

- The practical lesson follows directly.
- A correlation matrix estimated in calm markets understates portfolio risk in a downturn.
- Kestrel's V a R just showed it: sixty eight point five million dollars on calm inputs, but ninety six point six million on stressed ones.
- So V a R needs stressed correlation inputs, such as correlations measured in past recessions and sell-offs, alongside the calm estimates.
- Hedge ratios need them too, because a hedge sized on a calm correlation can behave very differently when correlations jump.
- One more property matters.
- Correlations mean revert: after a spike, they drift back towards a long run level.
- That is our next topic.

11 Exam traps: level versus volatility

- Recession: higher correlation AND higher volatility
- Level = the correlation; volatility = how much it moves
- High but steady differs from low but jumpy
- Diversification is weakest exactly when needed
- Calm estimates are optimistic, not conservative

IN THE LESSON

- Before moving on, here are the exam points for this part.
- First, compared with an expansion, a recession means higher correlation and higher correlation volatility.
- Second, do not confuse the two ideas.
- The level is the correlation itself, say point four five.
- The volatility is how much it moves from one period to the next.
- A question may describe a correlation that is high but steady, or one that is low but jumpy, and ask which property has changed.
- Third, diversification is weakest exactly when it is needed.
- Fourth, a calm period estimate is not a conservative estimate.
- It is an optimistic one, because it understates downturn risk.

12 Mean reversion: a rubber band

- Mean reversion: pulled back towards a long-run level
- Bigger gap from that level, stronger pull
- Speed a: share of the gap closed each period
- Long-run mean μ; time step dt, such as one month
- dt = 1: a = 1 closes the gap at once; a = 0, no pull

IN THE LESSON

- Kestrel's team notices something else in its data.
- When correlation spikes, it does not stay high for ever.
- It drifts back.
- This is mean reversion, a tendency to be pulled towards a long run level, like a rubber band that stretches and pulls back.
- The further correlation is from that level, the stronger the pull.
- Three quantities describe it.
- The speed of reversion, written ay, says what share of the gap closes each period.
- The long run mean, written mu, is the level it is pulled towards.
- And the time step, d t, is the length of one period, such as one month.
- With a time step of one, a speed of one closes the whole gap at once, and a speed of zero means no pull at all.

13 The mean reversion model

C(t) − C(t−1) = a × (μ − C(t−1)) × dt + noise
Change = speed × gap × time step + noise

Reversion speed, a	Share of the gap closed
Long-run mean, μ	Level it is pulled towards
Time step, dt	Length of one period
Last value, C(t−1)	Last period's value

IN THE LESSON

- Here is the model.
- The change in correlation from one period to the next, C of t minus C of t minus one, equals the speed ay, times the gap between the long run mean mu and last period's correlation, times the time step d t, plus noise, meaning random news.
- In words: change equals speed, times gap, times time step, plus noise.
- If last period's correlation is above mu, the gap is negative, so the expected change is down.
- If it is below mu, the expected change is up.
- Either way, correlation is pulled towards mu, and the size of the pull grows with the gap.

14 Estimating it with a regression

Fit = α + β × C(t−1)
β = −a × dt a = −β (dt = 1) μ = α ÷ a

Slope, β	Equals minus a × dt
Intercept, α	Equals a × μ × dt
Reversion speed, a	Minus slope, if dt = 1
Long-run mean, μ	Intercept ÷ speed

IN THE LESSON

- How do we estimate ay and mu?
- With a standard regression.
- Regress the change in correlation on the lagged level, meaning last period's correlation.
- The fitted line is alpha plus beta times C of t minus one.
- Now match it to the model.
- Multiply out the model and you get ay times mu times d t, minus ay times d t times C of t minus one.
- So the slope beta equals minus ay times d t, and the intercept alpha equals ay times mu times d t.
- With a time step of one, the speed is minus beta, and the long run mean is alpha divided by ay.
- Notice the minus sign: a mean reverting series has a negative slope.

15 Example 3: Kestrel's correlation data

Kestrel's average correlation at six month-ends: 0.34, 0.41, 0.53, 0.45, 0.27, 0.19

Month 1 to 2	level 0.34, then 0.41	**+0.07**
Month 2 to 3	level 0.41, then 0.53	**+0.12**
Month 3 to 4	level 0.53, then 0.45	**−0.08**
Month 4 to 5	level 0.45, then 0.27	**−0.18**
Month 5 to 6	level 0.27, then 0.19	**−0.08**

Means: lagged level 2.00 ÷ 5 = 0.40; change −0.15 ÷ 5 = −0.03

IN THE LESSON

- Let us run the regression by hand on a tiny sample, so you see every step.
- Kestrel's average correlation at six month ends was point three four, point four one, point five three, point four five, point two seven and point one nine.
- Each pair of months gives one data point: the lagged level and the change.
- From point three four to point four one is a change of plus point zero seven.
- From point four one to point five three, plus point one two.
- From point five three to point four five, minus point zero eight.
- From point four five to point two seven, minus point one eight.
- From point two seven to point one nine, minus point zero eight.
- The five lagged levels add to two, so their mean is point four zero.
- The changes add to minus point one five, so their mean is minus point zero three.

16 Example 3: gaps from the means

Gap = value − mean. Lagged-level mean 0.40; change mean −0.03

Month 1 to 2	−0.06 × +0.10	**−0.0060**
Month 2 to 3	+0.01 × +0.15	**+0.0015**
Month 3 to 4	+0.13 × −0.05	**−0.0065**
Month 4 to 5	+0.05 × −0.15	**−0.0075**
Month 5 to 6	−0.13 × −0.05	**+0.0065**

Sum of the products = −0.0120

IN THE LESSON

- Next, subtract the means.
- For each month, take the lagged level minus point four zero, and the change minus its mean of minus point zero three, which means adding point zero three.
- Then multiply the two gaps.
- Month one to two: minus point zero six times plus point one zero, minus point zero zero six zero.
- Month two to three: plus point zero one times plus point one five, plus point zero zero one five.
- Month three to four: plus point one three times minus point zero five, minus point zero zero six five.
- Month four to five: plus point zero five times minus point one five, minus point zero zero seven five.
- Month five to six: minus point one three times minus point zero five, plus point zero zero six five.
- The products add to minus point zero one two zero.
- The negative sum already tells us that high levels tend to be followed by falls.

17 Example 3: the slope and the intercept

Slope = sum of products ÷ sum of squared level gaps

Sum of products	from the last slide	**−0.0120**
Sum of squared level gaps	(36+1+169+25+169) ÷ 10,000	**0.0400**
Slope β	−0.0120 ÷ 0.0400	**−0.30**
Intercept α	−0.03 − (−0.30 × 0.40)	**+0.09**
Fitted line	change = 0.09 − 0.30 × C(t−1)	**the fit**

IN THE LESSON

- The slope is the sum of products divided by the sum of squared level gaps.
- The sum of products, from the last slide, is minus point zero one two zero.
- For the denominator, square the five level gaps.
- In ten thousandths that is thirty six, one, one hundred and sixty nine, twenty five and one hundred and sixty nine, which add to four hundred, so point zero four zero zero.
- The slope beta is minus point zero one two zero divided by point zero four zero zero, which is minus point three zero.
- The intercept alpha is the mean change minus beta times the mean level: minus point zero three minus minus point three zero times point four zero.
- That is minus point zero three plus point one two, or plus point zero nine.
- So the fitted line is: change equals point zero nine minus point three zero times the lagged level.

18 Example 3: speed, mean and autocorrelation

Fit: change = 0.09 − 0.30 × C(t−1). Monthly data, dt = 1 month

Speed a	−β = −(−0.30)	**0.30 a month**
Long-run mean μ	α ÷ a = 0.09 ÷ 0.30	**0.30**
Autocorrelation	1 − a × dt = 1 − 0.30	**0.70**
Gap shrinks each month	by the factor 1 − a	**× 0.70**

IN THE LESSON

- Now read the model off the fit.
- The data are monthly and we measure time in months, so d t is one.
- The speed of reversion is minus beta, minus minus point three zero, so point three zero a month: thirty per cent of the gap closes each month.
- The long run mean is alpha divided by the speed, point zero nine divided by point three zero, which is point three zero.
- That matches the normal state level we saw earlier.
- The matching autocorrelation, the correlation between one month's value and the next, is one minus ay times d t, so one minus point three zero, or point seven zero.
- Equivalently, the expected gap from the mean shrinks by a factor of point seven zero each month.

19 Autocorrelation: the level on its own lag

C(t) = α + (1 + β) × C(t−1) + noise
Slope = 1 + β = 1 − a = autocorrelation

Intercept, α	Same intercept as before
Slope, β	Slope of the change regression
Autocorrelation	Kestrel: 1 − 0.30 = 0.70

IN THE LESSON

- There is a second way to see the autocorrelation.
- Add C of t minus one to both sides of the change regression.
- The level today equals alpha, plus one plus beta times the level last period, plus noise.
- So if you regress the level on its own lag, the intercept is the same alpha, and the slope is one plus beta, which is one minus ay.
- For Kestrel that is one minus point three zero, or point seven zero, the autocorrelation we found.
- The two regressions carry the same information.
- Just check which one you have been given: a slope near minus point three on changes, or a slope near point seven on levels.

20 Example 4: forecasting Kestrel's correlation

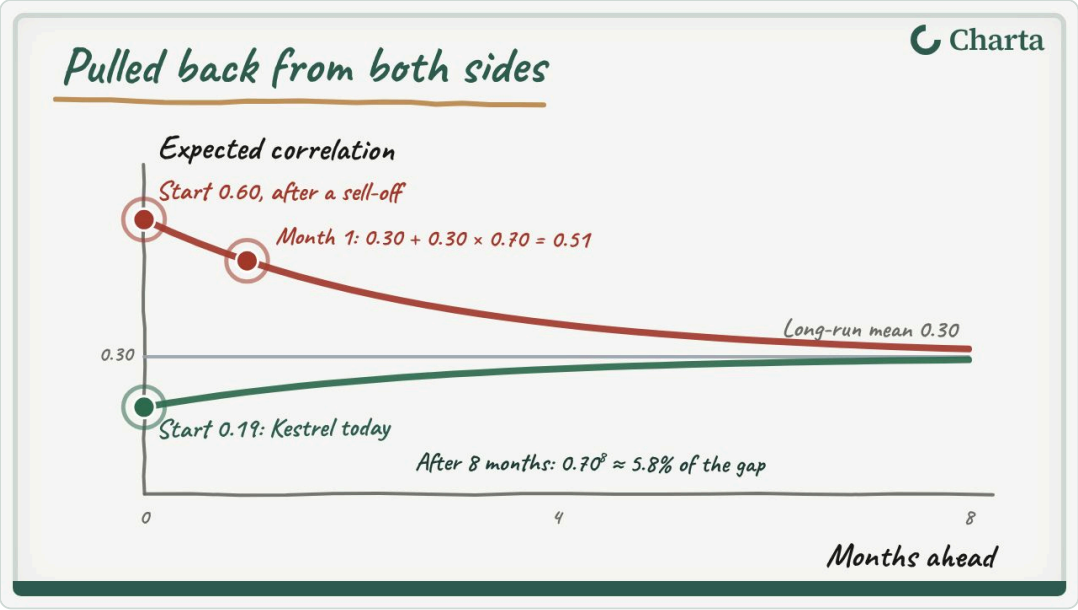
Latest reading (month 6): 0.19. μ = 0.30, a = 0.30 a month

Gap today	0.19 − 0.30	**−0.11**
Expected change, month 7	0.09 − 0.30 × 0.19	**+0.033**
Expected level, month 7	0.30 − 0.11 × 0.70	**0.223**
Expected level, month 8	0.30 − 0.077 × 0.70	**0.2461**
Gap after 2 months	−0.11 × 0.70^2	**−0.0539**

IN THE LESSON

- Now use the model to forecast.
- Kestrel's latest reading, at month six, is point one nine, below the long run mean of point three zero.
- The gap today is point one nine minus point three zero, which is minus point one one.
- The expected change for month seven is point zero nine minus point three zero times point one nine, which is point zero nine minus point zero five seven, or plus point zero three three.
- So the expected level for month seven is point two two three.
- With the gap, point three zero minus point one one times point seven zero gives the same point two two three.
- For month eight the gap is minus point zero seven seven times point seven zero, minus point zero five three nine, so the expected level is point three zero minus point zero five three nine, which is point two four six one.
- After two months the gap is minus point one one times point seven zero squared, which is minus point zero five three nine.

21 Pulled back from both sides



IN THE LESSON

- This chart shows the expected path of correlation over eight months, with a speed of point three zero a month and a long run mean of point three zero, the grey line.
- The upper curve starts at point six zero, as just after a sell-off spike.
- One month later its gap of point three zero has shrunk to point two one, so the expected level is point five one.
- The lower curve starts from Kestrel's point one nine and climbs towards the mean.
- Both curves close thirty per cent of their gap each month, so they move fast at first and then slowly.
- After eight months, point seven zero to the power eight, about five point eight per cent of each starting gap, is left.
- Mean reversion pulls from both sides.

22 When the time step is not one

Same monthly fit, but the speed is quoted per year: dt = 1/12

Slope	β = −a × dt	**−0.30**
Speed per year	a = −β ÷ dt = 0.30 × 12	**3.6**
Long-run mean	α ÷ (a × dt) = 0.09 ÷ 0.30	**0.30**
Autocorrelation	1 − a × dt = 1 − 3.6 ÷ 12	**0.70**

The speed depends on the time unit; the autocorrelation does not

IN THE LESSON

- Sometimes the speed is quoted per year while the data are monthly.
- Then the time step d t is one twelfth of a year.
- The slope is still minus ay times d t, and it is still minus point three zero.
- So the speed per year is minus beta divided by d t, point three zero times twelve, which is three point six.
- The long run mean does not change: alpha divided by the product of ay and d t, point zero nine divided by point three zero, still point three zero.
- And the autocorrelation is one minus ay times d t, one minus three point six divided by twelve, still point seven zero.
- The speed number depends on the unit of time, but the autocorrelation between successive observations does not.
- Always check the time step before quoting a speed.

23 Reading the slope

With dt = 1: speed = −slope, autocorrelation = 1 − speed

Slope β	Speed a	Autocorrelation	Reversion
0	0	1	None
−0.05	0.05	0.95	Very slow
−0.30	0.30	0.70	Kestrel
−0.80	0.80	0.20	Fast

IN THE LESSON

- This table links the slope, the speed and the autocorrelation, with a time step of one.
- Read each row in turn.
- A slope of zero means a speed of zero and an autocorrelation of one: no pull at all, and correlation wanders freely.
- A slope of minus point zero five means a speed of point zero five and an autocorrelation of point nine five: very slow reversion, where this month's correlation is almost the same as last month's.
- Kestrel's slope of minus point three zero gives a speed of point three zero and an autocorrelation of point seven zero.
- A slope of minus point eight zero gives a speed of point eight zero and an autocorrelation of point two zero: fast reversion, where most of any gap closes within one period.
- The exam point: a slope near zero means slow reversion and an autocorrelation near one.

24 Mistakes to avoid with the regression

Kestrel's fit: intercept 0.09, slope −0.30, time step 1 month

Wrong: speed = slope	drops the minus sign	**−0.30**
Right: speed = −slope	−(−0.30)	**0.30**
Wrong: mean = intercept	reports α as the mean	**0.09**
Right: mean = α ÷ a	0.09 ÷ 0.30	**0.30**
Autocorrelation	1 − a, not a itself	**0.70**

IN THE LESSON

- Two errors cost marks again and again.
- The first is dropping the minus sign.
- Reporting the slope itself as the speed gives minus point three zero, a negative speed, which would mean correlation is pushed away from its mean rather than pulled back.
- The right answer is minus the slope, which is point three zero.
- The second error is reporting the intercept as the mean.
- The intercept, point zero nine, is ay times mu, not mu itself.
- The right mean is point zero nine divided by point three zero, which is point three zero.
- Finally, keep the speed and the autocorrelation apart.
- The autocorrelation is one minus ay, point seven zero, not ay itself.

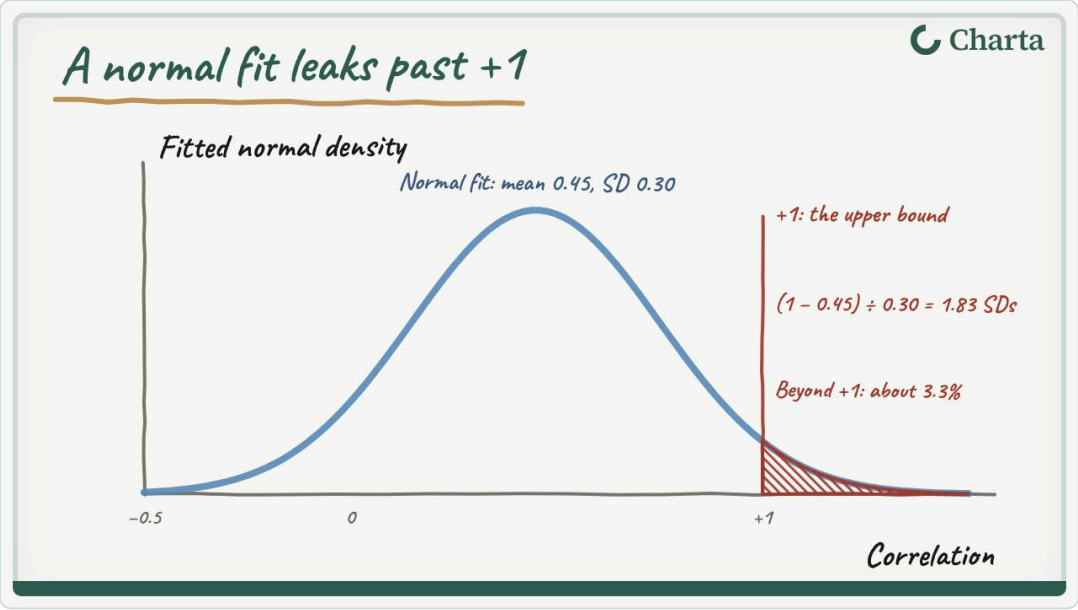
25 Correlations live between −1 and +1

- Correlations always lie between −1 and +1
- Normal: unbounded, puts probability on impossible values
- Normal: symmetric, but correlation data are often skewed
- So the normal fits correlation data poorly
- Better: bounded or more flexible shapes

IN THE LESSON

- Now the last outcome: which statistical distribution best describes correlation data.
- Suppose Kestrel collects many monthly correlation estimates and asks what shape their distribution has.
- The natural first guess is the normal distribution, the bell curve.
- But correlations are bounded: they always lie between minus one and plus one.
- A normal distribution has no bounds, so it puts some probability on impossible values beyond plus one or below minus one.
- It is also symmetric, while correlation data are often skewed, bunched on one side with a longer tail on the other.
- So the normal fits correlation data poorly, and researchers look for bounded or more flexible shapes.

26 A normal fit leaks past +1



IN THE LESSON

- Here is the problem in a picture.
- Kestrel fits a normal distribution to its equity correlations, with a mean of point four five and a standard deviation of point three zero.
- The curve is the fitted bell shape.
- The vertical line marks plus one, the highest value a correlation can take.
- The hatched area beyond it is the probability the normal gives to impossible values.
- Plus one is one minus point four five, divided by point three zero, which is one point eight three standard deviations above the mean.
- The normal puts about three point three per cent of its probability beyond that point, so about one estimate in thirty would be an impossible correlation above one.
- A bounded distribution puts zero probability there.

27 Goodness-of-fit tests

- Fit each candidate, then run a goodness-of-fit test
- Kolmogorov-Smirnov, Anderson-Darling, chi-squared
- Statistic = distance between sample and fitted curve
- KS: largest gap between the cumulative distributions
- Best fit: smallest statistic, or largest p-value

IN THE LESSON

- How do researchers choose between candidate distributions?
- They fit each one to the data and run a goodness of fit test.
- Three common ones are the Kolmogorov Smirnov test, the Anderson Darling test and the chi squared test.
- Each produces a test statistic that measures the distance between the sample and the fitted distribution.
- The Kolmogorov Smirnov statistic, for example, is the largest gap between the sample's cumulative distribution and the fitted one.
- A smaller distance means a closer fit.
- So the smallest statistic, or equivalently the largest p value, marks the best fit.
- The p value is the chance of seeing a distance at least this large if the fitted distribution were true.

28 Example 5: picking the best fit

Kestrel's equity correlations: Kolmogorov-Smirnov statistic for three fits

Normal	unbounded, symmetric	**0.110**
Beta	bounded, two shape parameters	**0.052**
Johnson SB	bounded, four parameters	**0.031**
Best fit	smallest statistic	**Johnson SB**
Wrong reading	largest statistic	**Normal**

IN THE LESSON

- Kestrel's team fits three candidates to its equity correlations and computes the Kolmogorov Smirnov statistic for each, using figures made up for our example.
- The normal, which is unbounded and symmetric, has a statistic of point one one zero.
- The Beta distribution, which is bounded and has two shape parameters, has point zero five two.
- The Johnson S B distribution, which is bounded and has four parameters, has point zero three one.
- The best fit has the smallest statistic, so Johnson S B wins.
- The frequent slip is to read the result the other way round and pick the largest statistic.
- Here that would choose the normal, the worst fit of the three.

29 Example 6: reading p-values

Kestrel's credit team: Anderson-Darling p-values for default correlations

Normal	very poor fit	**p = 0.01**
Generalised extreme value	flexible, often skewed	**p = 0.12**
Johnson SB	bounded, four parameters	**p = 0.40**
Best fit	largest p-value	**Johnson SB**

IN THE LESSON

- Kestrel's credit team does the same for default correlations, meaning the correlations between the default events of different borrowers.
- This time the software reports Anderson Darling p values, again made up for our example.
- The normal has a p value of point zero one, a very poor fit.
- The generalised extreme value distribution, a flexible and often skewed shape, has a p value of point one two.
- The Johnson S B distribution, which is bounded and has four parameters, has a p value of point four zero.
- With p values the rule turns round: a larger p value means the data sit closer to the fitted curve, so the best fit is the largest p value, and Johnson S B wins.
- The study the reading reports also found Johnson S B the best fit for default correlations.

30 The three pairings to remember

Bounded or flexible shapes beat the normal

Correlations of	Best fit	Shape
Equities	Johnson SB	Flexible, 4 parameters, bounded
Bonds	Generalised extreme value	Flexible, can be skewed
Defaults	Johnson SB	Flexible, 4 parameters, bounded

IN THE LESSON

- Here are the empirical findings to remember, from the study the reading reports.
- For equity correlations, it found the best fit was the Johnson S B distribution, a flexible four parameter distribution that is bounded at both ends.
- For bond correlations, it found the generalised extreme value distribution, usually shortened to G E V, a flexible family that can take a skewed shape.
- For default correlations, it again found Johnson S B.
- The Beta distribution is another bounded shape that is sometimes used for correlations, but it was not the winner here.
- Notice what the winners have in common.
- Each is either bounded, like the true range of a correlation, or flexible enough to bend to the data.
- None of them is the normal.
- Learn the three pairings exactly, because exam questions test them directly.

31 Putting it together for Kestrel

- Run VaR with recession-level correlations too
- Stress is jumpier, not just higher: test several values
- Forecast with mean reversion: a = 0.30, μ = 0.30
- Simulate from a bounded fit, e.g. Johnson SB for equity
- So no simulated correlation passes ±1

IN THE LESSON

- Let us put the three ideas together for Kestrel's risk team.
- First, it runs its V a R with recession level correlations as well as calm ones, since the calm matrix understates downturn risk.
- Second, it expects correlations in stress to be jumpier, not just higher, so it tests more than one stressed value.
- Third, it forecasts correlation with its mean reversion model: a speed of point three zero a month towards a long run level of point three zero.
- Fourth, when it simulates correlations, it draws from a bounded fit, such as Johnson S B for equity correlations, rather than from a normal.
- That way a simulated correlation never passes plus or minus one.

32 Try three quick questions

1. Intercept 0.10, slope −0.25, dt = 1. Speed, mean, autocorrelation?
   Answer: a = 0.25; μ = 0.10 ÷ 0.25 = 0.40; autocorr 0.75.
2. KS statistics 0.08, 0.05, 0.11. Best fit? Usual winner for equity?
   Answer: The smallest, 0.05. Equity: Johnson SB.
3. Recession versus expansion: equity correlation and its volatility?
   Answer: Both rise: calm inputs understate downturn VaR.

IN THE LESSON

- Now try three quick questions.
- Pause the video, think, then listen to the answers.
- Question one.
- A regression of the monthly change in correlation on its lagged level gives an intercept of point one zero and a slope of minus point two five, with a time step of one.
- What are the speed of reversion, the long run mean and the autocorrelation?
- The speed is point two five, the mean is point one zero divided by point two five, which is point four zero, and the autocorrelation is point seven five.
- Question two.
- Three fits to equity correlations give Kolmogorov Smirnov statistics of point zero eight, point zero five and point one one.
- Which fits best, and which distribution usually wins for equity correlations?
- The fit with point zero five, the smallest, and Johnson S B.
- Question three.
- Compared with an expansion, what happens in a recession to equity correlations and to their volatility?
- Both rise, so calm period correlations understate downturn V a R.