CHARTA LEARNING · LESSON NOTES

Correlation Basics

FRM Part II

1 A story: when co-movement changes

- Larkspur Bank: a trading desk in calm years
- Hedges, mortgage slices, quantos, VaR capital
- Crisis: hedge leaks, safe slices lose, capital up
- Cause: a change in how prices move together
- Loss from adverse changes in correlation

IN THE LESSON

- Meet Larkspur Bank and its trading desk.
- In calm years the desk held a portfolio of bank shares and hedged it with an index future, because the two moved closely together.
- It also held slices of mortgage pools, sold quanto options to clients, and held capital against its V a R.
- Then a crisis came.
- The hedge leaked, the slices thought to be safe lost money, and the capital charge jumped.
- None of these losses came from one price on its own.
- They came from a change in how prices moved together.
- That is financial correlation risk: the risk of loss from an adverse change in the correlation between two or more financial variables, such as asset returns, exchange rates or the defaults of different borrowers.

2 What correlation measures

ρ = Cov(X, Y) ÷ (σ_X × σ_Y)
Lies between −1 and +1; linear dependence only

Covariance of X and Y	Do X and Y move together?
Volatility of X, σX	Standard deviation of X
Volatility of Y, σY	Standard deviation of Y
Correlation, ρ	The measure; risk is its change

IN THE LESSON

- First, the measure itself.
- Covariance tells us whether two variables, X and Y, tend to move in the same direction, but its size depends on their units.
- Dividing by the product of the two standard deviations, sigma X and sigma Y, removes the units.
- The result is the correlation, written rho.
- So rho equals the covariance of X and Y, divided by sigma X times sigma Y.
- It always lies between minus one and plus one.
- This is the Pearson correlation, and it captures only linear dependence, meaning how well a straight line describes the relationship.
- Keep two words apart.
- Correlation is the measure.
- Correlation risk is the exposure to a change in that measure.

3 Example 1: a hedge that weakens

Bank shares σ = 20%; index future σ = 15%; covariance 0.027

Correlation, calm	0.027 ÷ (0.20 × 0.15)	**0.90**
Variance removed	0.90^2	**81%**
Volatility left	20% × √(1 − 0.81)	**8.72%**
Stress: ρ 0.6, hedge reset	0.6^2 removed	**36%**
Volatility left, stress	20% × √(1 − 0.36)	**16.0%**

Calm-market estimates of ρ are unreliable in stress

IN THE LESSON

- Larkspur's bank shares have a volatility of twenty per cent, the index future fifteen per cent, and their covariance is point zero two seven.
- The correlation is point zero two seven divided by point two times point one five, which is point nine.
- With the best hedge ratio, a hedge removes a fraction rho squared of the variance.
- Point nine squared is eighty one per cent.
- The volatility left is twenty per cent times the square root of one minus point eight one, about eight point seven two per cent.
- Now a crisis comes.
- Most correlations rise in stress, but the link between this hedge and the shares weakens, and the correlation falls to point six.
- Even if Larkspur resets the hedge to the best ratio, it removes only point six squared, thirty six per cent.
- The volatility left is twenty per cent times the square root of point six four, sixteen per cent.
- The hedge weakens as correlation falls, and a correlation estimated in calm markets proved unreliable in stress.

4 Example 2: zero correlation, full dependence

X is −1, 0 or +1, each with probability 1/3; $Y = X^2$

Mean of X	(−1 + 0 + 1) ÷ 3	**0**
Mean of XY = mean of X^3	(−1 + 0 + 1) ÷ 3	**0**
Covariance	E[XY] − E[X] × E[Y]	**0**
Correlation	0 ÷ (σX × σY)	**0**
Yet Y is fixed by X	know X, know Y exactly	**dependent**

IN THE LESSON

- Let X take the values minus one, zero and plus one, each with probability one third, and let Y equal X squared.
- This classic case shows that, because Pearson correlation sees only straight lines, a correlation of zero does not mean independence.
- The mean of X is minus one plus zero plus one, over three, which is zero.
- The mean of X times Y is the mean of X cubed, again minus one plus zero plus one over three, so zero.
- The covariance is the mean of X Y minus the mean of X times the mean of Y, which is zero minus zero, so zero.
- The correlation is therefore zero.
- Yet Y is completely fixed by X: if you know X, you know Y exactly.
- The relationship is a curve, not a line.

5 Where correlation risk shows up

Wherever a position or a hedge relies on co-movement

Area	Where correlation enters
Investments	Diversification, asset allocation
Trading	Multi-asset and credit derivatives
Risk management	Portfolio VaR, hedge ratios
Regulation	Capital rules that combine risks
Whole system	Stability: firms failing together

IN THE LESSON

- Correlation risk appears wherever a position or a hedge relies on co-movement.
- In investments, diversification and asset allocation depend on assets not all falling together.
- In trading, the prices of multi-asset options and of credit derivatives, such as collateralised debt obligations, depend on correlation inputs.
- In risk management, portfolio V a R and hedge ratios are built from correlations.
- In regulation, capital rules combine risks using correlations.
- And for the whole financial system, stability depends on banks and markets not failing together.
- In each of these areas, the trouble is the same: correlations estimated in quiet times change when they matter most.

6 Correlation and the 2007-2009 crisis

- Pools cut into equity, mezzanine, senior
- Gaussian copula pricing: one correlation input
- Input low, constant, from calm regional history
- House prices fell US-wide: defaults correlated
- Losses reached tranches rated AAA

IN THE LESSON

- Now the crisis.
- Banks pooled mortgages into mortgage backed securities and collateralised debt obligations, C D Os, and cut each pool into tranches.
- The equity tranche takes the first losses, the mezzanine tranche the next, and the senior tranche only the last.
- These were priced mainly with the Gaussian copula, a model that ties together the default times of the loans through a single correlation input.
- That input was low, held constant, and calibrated on a calm history of regionally diverse mortgages.
- Then U S house prices fell across the whole country at once.
- Defaults became strongly correlated, and losses reached tranches that had been thought safe, including tranches rated triple A.

7 Example 3: the same pool, two correlations

Pool $100m, EL 5%. Equity 0-10, mezzanine 10-20, senior 20-100 ($m)

Pool expected loss, any ρ	5% × 100	**$5.0m**
ρ ≈ 0: pool loses 5 a year	all of it hits equity	**equity $5.0m**
ρ ≈ 0: senior expected loss	losses never reach 20	**$0**
ρ = 1: all or nothing, p = 5%	equity: 5% × 10	**$0.5m**
ρ = 1: senior expected loss	5% × 80	**$4.0m**

ρ = 1: mezzanine 5% × 10 = $0.5m; total 0.5 + 0.5 + 4.0 = $5.0m

IN THE LESSON

- Larkspur holds tranches of a one hundred million dollar pool with an expected loss of five per cent.
- Equity covers losses from zero to ten million, mezzanine ten to twenty, and senior twenty to one hundred.
- The pool's expected loss is five million whatever the correlation.
- With correlation near zero and many loans, the pool loses about five million every year.
- All of it hits the equity tranche, which loses five million a year.
- Losses never reach twenty, so the senior tranche expects nothing.
- Now take correlation of one: all loans default together with probability five per cent, or none do.
- The equity tranche's expected loss is five per cent times ten, half a million.
- The senior tranche's is five per cent times eighty, four million.
- The mezzanine also expects half a million, so the total is still five million.

8 What the crisis teaches

- Higher ρ: losses cluster in bad states
- Equity tranche gains value: long correlation
- Senior tranche loses value: short correlation
- Low ρ hid senior risk; tranches rated AAA
- Exam: pool EL ignores ρ; tranche losses depend on it

IN THE LESSON

- What does the example teach?
- When correlation rises, losses cluster in bad states instead of being spread thinly across all states.
- That is good news for the equity tranche, whose expected loss fell from five million to half a million, so its value rises with correlation: the equity tranche is long correlation.
- It is bad news for the senior tranche, whose expected loss rose from nothing to four million: the senior tranche is short correlation.
- So a model that understated correlation understated the risk of senior tranches, and those tranches carried triple A ratings.
- Remember the exam point.
- Expected pool loss does not depend on correlation, but the losses of each tranche do.

9 Quanto options

- Quanto: paid in another currency, fixed rate
- Larkspur: call on a euro index, paid in dollars
- Price depends on ρ of index and FX rate
- ρ enters through the risk-neutral drift
- FX quoted as domestic per foreign: $ per €

IN THE LESSON

- Correlation enters any option price whose payoff depends on more than one random variable.
- Start with a quanto option, which pays in a currency other than that of its underlying asset, at a fixed exchange rate.
- Larkspur sells a client a call on a euro stock index, paid in dollars at a fixed rate, so the client takes no currency risk.
- But the bank does, and the price now depends on the correlation between the index and the exchange rate.
- That correlation enters through an adjustment to the risk-neutral drift of the index.
- The sign of the adjustment depends on how the exchange rate is quoted.
- We quote it as domestic currency per foreign currency, here dollars per euro, and rho is the correlation between the index and that rate.

10 The quanto forward

F = S × e^{(rF − ρ × σS × σFX) × T}
Higher ρ lowers the quanto forward and the quanto call

Asset price, S	In the asset's own currency
Foreign risk-free rate, rF	Foreign risk-free rate
Volatilities σS and σFX	Asset and exchange rate
Asset-FX correlation, ρ	Asset and $-per-€ rate

IN THE LESSON

- Here is the quanto forward price.
- Write S for the asset price, in its own currency, r F for the foreign risk-free rate, sigma S and sigma F X for the volatilities of the asset and the exchange rate, rho for their correlation, and T for the maturity in years.
- The quanto forward is S times e to the power of r F minus rho times sigma S times sigma F X, all times T.
- The term rho times sigma S times sigma F X is the quanto adjustment to the drift.
- The adjustment is subtracted, so when rho is positive it lowers the drift, and a higher correlation lowers the quanto forward, and with it the value of a quanto call.
- When rho is zero the adjustment vanishes.

11 Example 4: Larkspur's quanto forward

Index 4,000 (in euros). rF 3%, σS 20%, σFX 10% ($ per €), T = 1 year

Drift with ρ = 0	3% − 0	**3.0%**
Quanto forward, ρ = 0	4,000 × e^{0.03}	**4,121.82**
Drift with ρ = +0.5	3% − 0.5 × 20% × 10%	**2.0%**
Quanto forward, ρ = +0.5	4,000 × e^{0.02}	**4,080.81**
Quanto forward, ρ = −0.5	drift 3% + 1%: 4,000 × e^{0.04}	**4,163.24**

IN THE LESSON

- The index stands at four thousand points.
- The euro risk-free rate is three per cent, the index volatility twenty per cent, the volatility of the dollar per euro rate ten per cent, and the maturity one year.
- With a correlation of zero, the drift is just three per cent, and the quanto forward is four thousand times e to the power point zero three, about four thousand one hundred and twenty one point eight two.
- With a correlation of plus point five, the drift is three per cent minus point five times twenty per cent times ten per cent, which is two per cent.
- The forward falls to four thousand times e to the power point zero two, about four thousand and eighty point eight one.
- With minus point five, the drift is three per cent plus one per cent, four per cent, and the forward rises to four thousand times e to the power point zero four, about four thousand one hundred and sixty three point two four.

12 Options on several assets

Read the payoff before deciding the sign

Payoff	If ρ rises
Basket call	Gains
Call on the minimum (worst-of)	Gains
Joint-probability digital	Gains
Best-of call	Loses
Spread or exchange option	Loses

IN THE LESSON

- Read the payoff before deciding the sign.
- For exotic options on several assets, the direction depends on it.
- A basket call pays on the average of several assets.
- Higher correlation makes the basket more volatile, because the assets no longer offset each other, so the call gains.
- A call on the minimum, or worst of, pays only if every asset does well, which is likelier when they move together, so it gains.
- A joint probability digital pays if two events both happen, which is also likelier with high correlation, so it gains.
- A best of call pays on the strongest asset.
- It loves dispersion, so higher correlation hurts it.
- Spread and exchange options pay on the difference between two assets, and the difference becomes less volatile as correlation rises, so they lose.

13 Example 5: spread versus basket

Two assets: σ1 = 30%, σ2 = 25%. Basket 50/50. Compare ρ = 0.2 and ρ = 0.8

Spread vol, ρ = 0.2	√(0.09 + 0.0625 − 0.03)	**35.0%**
Spread vol, ρ = 0.8	√(0.1525 − 0.12)	**18.0%**
Basket vol, ρ = 0.2	√(0.0225 + 0.015625 + 0.0075)	**21.4%**
Basket vol, ρ = 0.8	√(0.038125 + 0.03)	**26.1%**

Spread variance has − 2ρσ1σ2: the spread option loses as ρ rises

IN THE LESSON

- Take two assets with volatilities of thirty and twenty five per cent, and an equal weighted basket of the two, and compare correlations of point two and point eight.
- For a spread, the relevant variance is sigma one squared plus sigma two squared, minus two times rho times sigma one times sigma two.
- At a correlation of point two, that is point zero nine plus point zero six two five minus point zero three, and the square root gives thirty five per cent.
- At point eight the cross term is point one two, and the spread volatility falls to about eighteen per cent.
- For the basket, the cross term is added instead.
- At point two, the variance is point zero two two five plus point zero one five six two five plus point zero zero seven five, a volatility of about twenty one point four per cent.
- At point eight, it rises to about twenty six point one per cent.
- Same assets, same rise in correlation, opposite effects.

14 The correlation swap

- Fixed correlation for the realised average
- Swap rate set so the value is zero at inception
- No principal changes hands
- Buyer receives notional × (realised − fixed); negative: pays
- So the buyer is long correlation

IN THE LESSON

- How can Larkspur trade correlation itself?
- With a correlation swap.
- It exchanges a fixed correlation, called the swap rate, for the realised average pairwise correlation of a set of assets over the life of the contract.
- The swap rate is set so that the swap has no value at inception, and no principal changes hands.
- At maturity the buyer receives the notional times realised correlation minus fixed correlation, and the seller pays it.
- If realised correlation ends below the fixed rate, the payment runs the other way.
- So the buyer gains when correlation rises: the buyer is long correlation.

15 The realised average correlation

Average ρ = 2 ÷ (n² − n) × Σ_{i<j} ρ_{ij}
Buyer receives Notional × (realised − fixed)

Number of pairs	n assets give n(n − 1) ÷ 2 pairs
Realised average correlation	Average over all the pairs
Fixed correlation, the swap rate	Set so value is zero at start
Notional	No principal changes hands

IN THE LESSON

- With n assets there are n times n minus one, divided by two, distinct pairs.
- The realised average is the average of the correlations of all those pairs.
- In symbols, the average equals two divided by n squared minus n, times the sum of the pairwise correlations rho i j, taking each pair once.
- The two and the n squared minus n together are just one over the number of pairs.
- The payoff to the buyer is the notional times the realised average minus the fixed correlation.
- The fixed correlation, the swap rate, is set so the swap is worth zero at the start, and the notional only scales the payment: no principal changes hands.

16 Example 6: settling a correlation swap

4 stocks, fixed rate 0.30, notional $5m. Larkspur is the buyer

Number of pairs	4 × 3 ÷ 2	**6**
Sum of pair correlations	0.2 + 0.3 + ... + 0.7	**2.7**
Realised average	2 ÷ (16 − 4) × 2.7	**0.45**
Buyer receives	$5m × (0.45 − 0.30)	**$750,000**
Wrong: 2 ÷ 16 × 2.7 = 0.3375	$5m × (0.3375 − 0.30)	**$187,500**

IN THE LESSON

- Larkspur buys a one year correlation swap on four stocks, with a fixed rate of point three zero and a notional of five million dollars.
- Four stocks give four times three over two, six pairs.
- Over the year the six realised pairwise correlations are point two, point three, point four, point five, point six and point seven, which sum to two point seven.
- The realised average is two divided by sixteen minus four, times two point seven, which is point four five.
- Larkspur, the buyer, receives five million times point four five minus point three zero, seven hundred and fifty thousand dollars.
- The common slip is to divide by n squared, sixteen.
- That gives an average of point three three seven five, and a payment of five million times point three three seven five minus point three zero, only one hundred and eighty seven thousand five hundred dollars.

17 Who buys a correlation swap, and why

- Used to hedge or to speculate on correlation
- Dealer short basket calls loses if ρ rises: buys
- Dispersion trader loses if ρ rises: buys
- Expect ρ to rise: buy. Expect it to fall: sell
- The buyer is long correlation

IN THE LESSON

- Correlation swaps are used to hedge correlation risk or to speculate on it.
- Take two natural buyers.
- First, a dealer who has sold basket call options.
- We saw that basket calls gain value when correlation rises, so the dealer who is short them loses, and buying the swap pays off in exactly that case.
- Second, a dispersion trader, who sells options on an index and buys options on its members.
- That trade profits when the members move apart, so it loses when correlation rises, and the swap hedges it.
- A speculator who expects correlation to rise buys too, and one who expects it to fall sells.
- In every case, remember the direction: the buyer is long correlation.

18 Portfolio VaR and correlation

$VaR_p = \sqrt{(VaR_1^2 + VaR_2^2 + 2\rho \times VaR_1 \times VaR_2)}$
Benefit = VaR_1 + VaR_2 − VaR_p

Standalone VaRs	Same horizon and confidence
Correlation of the two, ρ	Elliptical returns assumed
Diversification benefit	Zero when ρ = +1

IN THE LESSON

- Now to capital.
- Portfolio V a R combines the standalone V a Rs with the correlation between them.
- Assume returns are normal, or more generally elliptical, and that both V a Rs use the same horizon and confidence level.
- For two long positions with standalone V a Rs, V a R one and V a R two, and correlation rho, portfolio V a R is the square root of V a R one squared plus V a R two squared plus two times rho times V a R one times V a R two.
- It rises with correlation.
- The diversification benefit is the sum of the standalone V a Rs less portfolio V a R.
- At a correlation of one, portfolio V a R is the simple sum and the benefit is zero.

19 Example 7: Larkspur's desk VaR

Two long books. 1-day 99% VaRs: equities $1.897m, bonds $2.530m

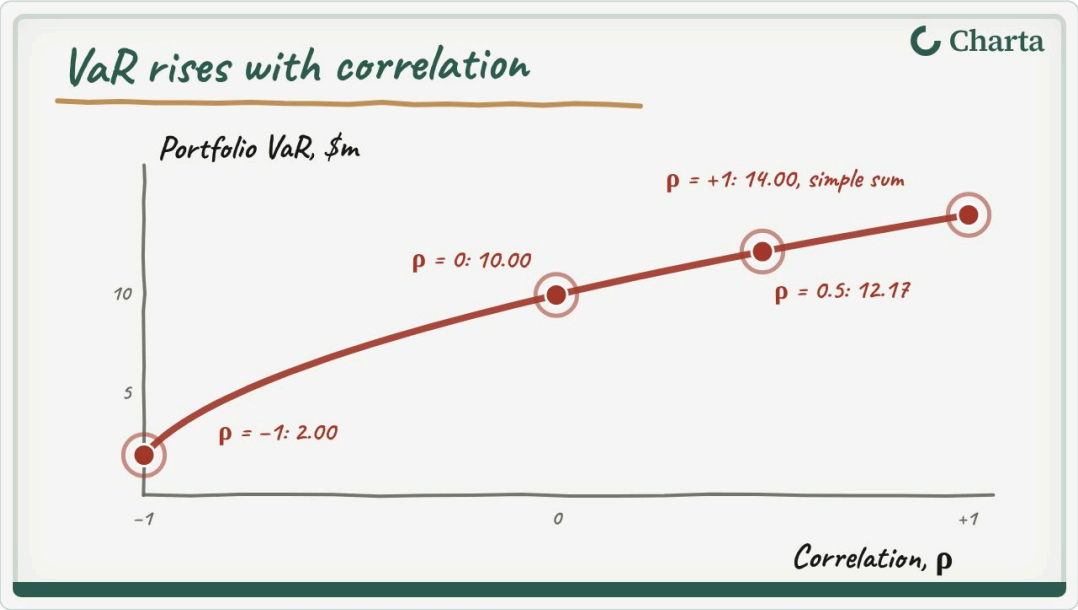
Equities, 10-day 99% VaR	1.897 × √10	**$6.00m**
Bonds, 10-day 99% VaR	2.530 × √10	**$8.00m**
Portfolio VaR, ρ = 0	√(36 + 64)	**$10.00m**
Portfolio VaR, ρ = 0.5	√(36 + 64 + 2 × 0.5 × 48)	**$12.17m**
Portfolio VaR, ρ = 1	6 + 8	**$14.00m**

Diversification benefit at ρ = 0.5: 14.00 − 12.17 = $1.83m

IN THE LESSON

- Larkspur's desk holds two long books, equities and bonds.
- Their one day ninety nine per cent V a Rs are one point eight nine seven million and two point five three zero million dollars.
- The capital rules use ten day V a R, obtained from one day V a R by the square root of time, so we multiply each by the square root of ten.
- That gives six million for equities and eight million for bonds.
- With a correlation of zero, portfolio V a R is the square root of thirty six plus sixty four, ten million.
- With a correlation of point five, the cross term adds forty eight, and portfolio V a R is the square root of one hundred and forty eight, about twelve point one seven million.
- With a correlation of one, it is the simple sum, fourteen million.
- At point five the diversification benefit is fourteen minus twelve point one seven, one point eight three million.

20 VaR rises with correlation



IN THE LESSON

- This chart shows Larkspur's portfolio V a R as the correlation between the two books moves from minus one to plus one, with the standalone V a Rs fixed at six and eight million.
- At plus one, on the right, it reaches fourteen million, the simple sum, with no diversification at all.
- At point five it is twelve point one seven million.
- At zero it is ten million.
- At minus one it falls to the difference between the two, just two million, because the books offset each other.
- For long positions, a higher correlation always raises portfolio V a R, and the diversification benefit shrinks to nothing at plus one.

21 From VaR to the capital charge

Capital = max(VaR_{yesterday} , k × average VaR_{60 days})
k ≥ 3; 10-day 99% VaR = 1-day VaR × √10

Yesterday's VaR	Yesterday's 10-day 99% VaR
Multiplier, k	At least 3
60-day average VaR	Last 60 business days

IN THE LESSON

- How does this reach capital?
- Under the original V a R based internal models approach, the general market risk charge is the higher of two numbers.
- The first is yesterday's V a R.
- The second is a multiplier, k, times the average V a R over the last sixty business days.
- The multiplier is at least three.
- The V a R used is ten day, ninety nine per cent V a R, usually obtained from one day V a R by the square root of time rule, multiplying by the square root of ten.
- Because portfolio V a R depends on correlation, a change in correlation flows straight into the capital charge, magnified by the multiplier.
- Basel two point five later added a stressed V a R term of the same form.

22 Example 8: correlation moves the capital charge

VaRs $6m and $8m. Calm ρ = 0. Stress: ρ = 0.5 for 60 days. k = 3

Calm: portfolio VaR	√(36 + 64)	**$10.00m**
Calm charge	max(10.00, 3 × 10.00)	**$30.00m**
Stress: portfolio VaR	√(36 + 64 + 48)	**$12.17m**
Stress charge	max(12.17, 3 × 12.17)	**$36.50m**
Extra capital from ρ alone	36.50 − 30.00	**$6.50m**

IN THE LESSON

- Larkspur's books still have V a Rs of six and eight million.
- Suppose they are steady, so yesterday's V a R equals the sixty day average, and the multiplier is three.
- In calm markets the correlation is zero and portfolio V a R is ten million.
- The charge is the higher of ten million and three times ten million, so thirty million dollars.
- Now stress arrives and the correlation rises to point five and stays there for sixty days.
- Portfolio V a R becomes the square root of thirty six plus sixty four plus forty eight, about twelve point one seven million.
- The charge is the higher of twelve point one seven and three times twelve point one seven, about thirty six point five zero million.
- Nothing about the positions changed, yet the bank needs six point five zero million more capital, from correlation alone.

23 Two cautions about correlation and VaR

- With a short: the cross term turns negative
- Long 6, short 8, ρ = 0.5: √(36 + 64 − 48) = 7.21
- Same pair, ρ = 0.9: √(100 − 86.4) = 3.69
- Higher ρ now lowers VaR: better hedge
- ρ jumps in crises; capital jumps too

IN THE LESSON

- Two cautions.
- First, the rule that higher correlation means higher V a R holds for two long positions.
- When one position is short, the cross term changes sign.
- Suppose Larkspur is long equities with a V a R of six million and short a related book with a V a R of eight million.
- At a correlation of point five, portfolio V a R is the square root of thirty six plus sixty four minus forty eight, about seven point two one million.
- At point nine, it is the square root of one hundred minus eighty six point four, about three point six nine million.
- Higher correlation now lowers V a R, because the short position hedges better.
- Second, correlations tend to jump in crises, so capital can jump with them, just when it is hardest to raise.

24 Correlation in market risk and credit risk

- Market risk: co-movement of risk factors
- Drives portfolio VaR, diversification, hedges
- Drives the value of multi-asset derivatives
- They change, usually upwards in stress
- Credit: default correlation; more years of big losses

IN THE LESSON

- Correlation risk affects both of the main risk types.
- In market risk, it is the co-movement of risk factors, such as share prices, interest rates and exchange rates.
- Portfolio V a R, diversification and hedge effectiveness all depend on it, as we have seen, and so does the value of every multi-asset derivative.
- The trouble is that these correlations change, and usually upwards in stress.
- In credit risk, correlation appears as default correlation, the tendency of obligors, meaning borrowers and counterparties, to default together.
- A loan book where defaults are correlated has more years of very large losses, even if each loan is just as risky as before.

25 Joint default probability

P(both) = p_A × p_B + ρ × √(p_A(1 − p_A) × p_B(1 − p_B))
Expected loss unchanged; joint defaults more likely

Default probability of A	Obligor A, on its own
Default probability of B	Obligor B, on its own
Default correlation, ρ	Tendency to default together

IN THE LESSON

- Here is how default correlation enters.
- Take two obligors, A and B, with default probabilities p A and p B, and default correlation rho.
- The probability that both default is p A times p B, plus rho times the square root of p A times one minus p A, times p B times one minus p B.
- The first term is the joint probability if the defaults were independent.
- The second term adds to it in proportion to the correlation.
- Notice what does not change: each obligor's own default probability, and therefore the expected number of defaults and the expected loss.
- Correlation moves probability towards the outcomes where both fail together.

26 Example 9: two borrowers that fail together

Borrower A: pA = 2%. Borrower B: pB = 3%. Compare ρ = 0 and ρ = 0.3

Joint default, ρ = 0	2% × 3%	**0.06%**
Square root term	√(0.0196 × 0.0291)	**0.02388**
Joint default, ρ = 0.3	0.0006 + 0.3 × 0.02388	**0.78%**
At least one fails, ρ = 0.3	5% − 0.78% (4.94% at ρ = 0)	**4.22%**
Expected defaults, any ρ	0.02 + 0.03	**0.05**

Joint default ≈ 13 × likelier; ρ fattens the loss tail; EL unchanged

IN THE LESSON

- Larkspur lends to two firms, A with a default probability of two per cent, and B with three per cent.
- If defaults are independent, the chance that both fail is two per cent times three per cent, point zero six per cent.
- Now set the default correlation at point three.
- The square root term is the square root of point zero one nine six times point zero two nine one, about point zero two three eight eight.
- The joint default probability becomes point zero zero zero six plus point three times point zero two three eight eight, about point seven eight per cent.
- The chance that at least one fails is five per cent minus point seven eight, about four point two two per cent, down from four point nine four at zero correlation.
- The expected number of defaults is point zero five in both cases.
- So the joint default is roughly thirteen times likelier, correlation fattens the tail of losses, and expected loss is unchanged.

27 Where default correlation bites

- CDO tranches: equity long ρ, senior short ρ
- First-to-default: higher ρ, fewer payouts
- Wrong-way: exposure up as credit falls (linked-bank CDS)
- Correlation fattens the loss tail
- Exam: shape changes, expected loss does not

IN THE LESSON

- Default correlation drives several credit products and risks.
- In C D O tranches, as we saw, the equity tranche is long correlation and the senior tranche is short it.
- In a first to default basket, which pays when the first of several names defaults, a higher correlation makes at least one default less likely, just as in our example, so the basket is less likely to pay.
- Then there is wrong way risk, where the exposure to a counterparty grows as its credit quality falls.
- An example is buying credit protection on a bank from another bank closely linked to it.
- Across all of these, correlation fattens the tail of the loss distribution.
- The exam point: correlation changes the shape of the loss distribution, not the expected loss.

28 Correlation and systemic risk

- Systemic risk: one firm's distress spreads
- Correlation is the channel
- Crisis: asset, default and firm correlations rise
- Diversification fails; losses arrive together
- Diversified firms can still move together

IN THE LESSON

- Now the wider system.
- Systemic risk is the risk that distress at one institution or market spreads and brings down much of the financial system.
- Correlation is the channel.
- In a crisis, the correlations among asset prices, among defaults and among institutions all rise.
- Diversification that works in normal times fails, and losses arrive together.
- Each firm can be well diversified on its own, yet the firms can be highly correlated with each other, because they hold the same assets or lend to one another.
- If Larkspur and its peers all own the same mortgage tranches, one shock hits them all at once.

29 Correlation and concentration risk

- Concentration: one name, sector or region
- Those exposures are more highly correlated
- Adding similar positions gives little benefit
- n equal positions: $\sigma \times \sqrt{(1/n + (1 - 1/n) \times \rho)}$
- However large n, never below $\sigma \times \sqrt{\rho}$

IN THE LESSON

- Concentration risk is the risk of loss from exposure that is heavily weighted to one name, one sector or one region.
- Exposures of that kind are more highly correlated, since the same shock hits them all.
- So adding more positions of the same kind gives little benefit.
- A simple formula shows why.
- For n equal positions, each with volatility sigma and average pairwise correlation rho, portfolio volatility is sigma times the square root of one over n, plus one minus one over n, times rho.
- As n grows, the first part vanishes, but the second does not.
- With positive correlation, however many positions you add, volatility never falls below sigma times the square root of rho.

30 Example 10: the floor under diversification

Equal positions, each σ = 25%. One region: ρ = 0.16. Spread out: ρ = 0.04

10 names, one region	25% × √(0.1 + 0.9 × 0.16)	**12.35%**
100 names, one region	25% × √(0.01 + 0.99 × 0.16)	**10.26%**
Floor, one region	25% × √0.16	**10.0%**
10 names, spread out	25% × √(0.1 + 0.9 × 0.04)	**9.22%**
Floor, spread out	25% × √0.04	**5.0%**

Correlation sets the limit on diversification

IN THE LESSON

- Larkspur compares two loan books of equal positions, each with a volatility of twenty five per cent.
- One is concentrated in a single region, with an average correlation of point one six.
- With ten names, volatility is twenty five per cent times the square root of point one plus point nine times point one six, about twelve point three five per cent.
- With one hundred names, it is twenty five per cent times the square root of point zero one plus point nine nine times point one six, about ten point two six per cent.
- However many names, the floor is twenty five per cent times the square root of point one six, ten per cent.
- The other book is spread across regions, with correlation point zero four.
- Ten names give about nine point two two per cent, and its floor is just five per cent.

31 Mistakes to avoid

- Zero correlation is not independence
- Pool EL ignores ρ; tranche losses depend on it
- Quanto: higher ρ lowers forward ($ per €)
- Swap: divide by $(n^2 - n) \div 2$; buyer long ρ
- Higher ρ: VaR up for longs, down with a short

IN THE LESSON

- Before the questions, here are the mistakes examiners look for.
- First, a correlation of zero does not mean independence, because Pearson correlation sees only linear links.
- Second, the expected loss of a pool does not depend on correlation, but tranche losses do.
- Third, in a quanto with the exchange rate quoted as domestic per foreign currency, a higher correlation lowers the forward and the call.
- Fourth, the average in a correlation swap divides by n squared minus n, over two, and the buyer is long correlation.
- Fifth, higher correlation raises V a R for two long positions but lowers it when one position is short.

32 Try three quick questions

1. VaRs $4m and $3m, both long, ρ = 0.5. Portfolio VaR?
   Answer: √(16 + 9 + 12) = √37 ≈ $6.08m.
2. Swap: 3 assets, fixed 0.35, notional $2m; pairs 0.3, 0.4, 0.5. Buyer gets?
   Answer: Average 0.4. $2m × 0.05 = $100,000.
3. pA = pB = 10%, default correlation 0.2. Joint default probability?
   Answer: 0.01 + 0.2 × 0.09 = 2.8%.

IN THE LESSON

- Now try three quick questions.
- Pause the video, think, then listen to the answers.
- Question one.
- Two long positions have standalone V a Rs of four million and three million dollars, and a correlation of point five.
- What is portfolio V a R?
- It is the square root of sixteen plus nine plus twelve, the square root of thirty seven, about six point zero eight million dollars.
- Question two.
- A correlation swap on three assets has a fixed rate of point three five and a notional of two million dollars.
- The realised pair correlations are point three, point four and point five.
- What does the buyer receive?
- The average is point four, so the buyer receives two million times point zero five, one hundred thousand dollars.
- Question three.
- Two obligors each have a default probability of ten per cent and a default correlation of point two.
- What is the joint default probability?
- It is point zero one plus point two times point zero nine, two point eight per cent.