CHARTA LEARNING · LESSON NOTES

Credit Risk in Derivatives and Portfolios

FRM Part II

1 A story: Marlow Bank and Corvane Metals

- A loan: the amount at risk is known
- A derivative: the amount at risk moves with the market
- Counterparty owes you: you lose if it defaults
- You owe the counterparty: you still pay in full

IN THE LESSON

- Marlow Bank has agreed an interest rate swap with Corvane Metals, a mining company.
- Compare this with an ordinary loan.
- If Marlow lends Corvane a million dollars, the amount at risk is known.
- A derivative is stranger.
- Suppose the swap is worth plus three million dollars to Marlow today, so Corvane owes Marlow.
- If Corvane defaults now, Marlow loses that amount.
- Now suppose the swap is worth minus three million to Marlow.
- Marlow owes Corvane, and if Corvane defaults Marlow must still pay in full.
- So Marlow loses nothing from the default itself.
- Credit risk exists only when the contract is an asset to you, and because its value moves with markets, the amount at risk is uncertain.

2 Credit exposure on a derivative

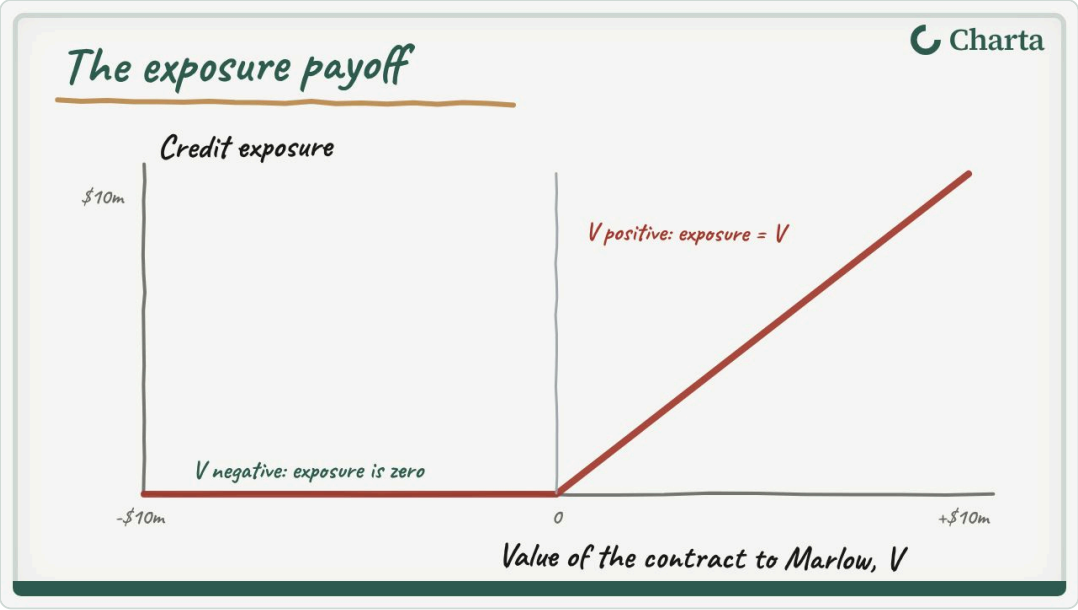
Exposure = max(V , 0)

Current value, V	Positive: owed to you
The floor of zero	You never gain from the default
Credit exposure	What is lost if it defaults now

IN THE LESSON

- Now the formal idea.
- Call the contract's current value to you V.
- Only a positive value is at risk, so the credit exposure to the counterparty at any time is the larger of V and zero.
- When V is positive, the counterparty owes you, and a default costs you V.
- When V is negative, you owe the counterparty, and the exposure is zero.
- Notice two things.
- First, this is a floor: exposure cannot be negative, because a default never gives you a gain on a contract you owe on.
- Second, the exposure is a property of one moment.
- Tomorrow V will be different, so the exposure changes with every move in the market variables behind the contract.

3 The exposure payoff



IN THE LESSON

- This chart draws the rule as a picture.
- Along the bottom is the value of the contract to Marlow, from minus ten million dollars on the left to plus ten million on the right.
- Up the side is the credit exposure.
- On the left of the middle line the contract is a liability to Marlow, and the exposure is flat at zero.
- Default then costs nothing.
- To the right of the middle the exposure rises one for one with the value.
- The shape is the payoff of a bought call option: your exposure is like an option you own, worth more the more the contract can move in your favour.

4 Which contracts carry how much risk

Credit risk depends on whether the contract can be an asset

Contract	Credit risk	Reason
Bought option	One-sided	Always an asset
Written option, paid	None	Only a liability
Forward or swap	Two-sided	Value changes sign
Exchange-traded	Little	Clearing house margins
Currency vs rate swap	Currency higher	Principal, FX moves

IN THE LESSON

- Now judge the contracts.
- A bought option is always an asset to its buyer, so its risk is one sided, and real.
- An option you have written and been paid for is only ever a liability to you.
- You hold no claim on the counterparty, so there is no credit risk.
- Forwards and swaps can change sign as markets move, so the risk is two sided: sometimes you are owed, sometimes you owe.
- Exchange traded contracts carry little credit risk, as the clearing house margins both sides.
- Finally, a currency swap is riskier than an interest rate swap of the same size.
- Principal is exchanged at the end, and exchange rates are volatile, so the value can swing far more.

5 Example 1: exposure is value, not notional

Swap, notional $200m. Value to Marlow +$3.2m, later −$1.8m

Exposure at +$3.2m	max(3.2, 0)	**$3.2m**
Share of notional	3.2 ÷ 200	**1.6%**
Exposure at −$1.8m	max(−1.8, 0)	**$0**
Wrong: use the notional	200 ÷ 200	**100%**

IN THE LESSON

- Marlow has a swap with Corvane with a notional of two hundred million dollars.
- Marlow's valuation team says the swap is worth plus three point two million dollars to Marlow today.
- The credit exposure is the larger of three point two million and zero, which is three point two million.
- That is one and six tenths per cent of the notional.
- Now suppose markets move and the swap is worth minus one point eight million.
- The exposure falls to zero, even though the notional has not changed.
- The exam point: measure exposure by the positive replacement value, never by the notional.

6 Credit and debt valuation adjustments

- No default value V ignores that anyone can fail
- CVA: PV of expected loss if the counterparty defaults
- DVA: PV of the loss they bear if you default
- CVA uses their default chance; DVA uses your own

IN THE LESSON

- A derivative has a price that ignores default, called the no default value.
- But Corvane could fail before the swap ends, and the market knows it.
- So the true value of the contract to Marlow is lower than the no default value, by the credit valuation adjustment, or C V A.
- It is the present value of the expected loss from the counterparty defaulting before the contract ends.
- There is a mirror image.
- Marlow can fail too, and then Corvane would bear a loss on what Marlow owes it.
- The present value of that expected loss, seen from Corvane's side, is the debt valuation adjustment, or D V A.
- It uses Marlow's own default probabilities.
- So each adjustment belongs to a different party's chance of default.

7 The CVA formula

CVA = Σ (1 − R) × q_i × v_i

Recovery rate, R	Share recovered after default
Default chance in interval i	Default chance in interval i
Expected exposure, vi	PV of expected positive exposure
CVA	Sum over all intervals

IN THE LESSON

- To compute C V A, split the life of the contract into intervals, for example years.
- For each interval i, three things matter.
- First, the probability that the counterparty defaults in that interval, called q i.
- Second, the present value of the expected positive exposure at that time, called v i.
- It is expected positive exposure because only a positive value can be lost, as we saw with the floor of zero.
- Third, the recovery rate R, the share of the claim that is recovered after a default.
- So the loss given default is one minus R.
- Multiply the three together for each interval and add up over the whole life of the contract.
- The sum is the C V A.

8 Example 2: Marlow's CVA on Corvane

Recovery 40%. Corvane q: 2%, 2.5%, 3%. PV of exposure: $4m, $5m, $3m

Year 1	0.6 × 2% × 4m	**$48,000**
Year 2	0.6 × 2.5% × 5m	**$75,000**
Year 3	0.6 × 3% × 3m	**$54,000**
CVA	48,000 + 75,000 + 54,000	**$177,000**

IN THE LESSON

- Marlow's swap with Corvane has three years left, and we use yearly intervals.
- Recovery is forty per cent, so one minus R is sixty per cent.
- Corvane's chance of default in each year is two per cent, two and a half per cent and three per cent.
- The present values of expected positive exposure are four, five and three million dollars.
- Year one gives point six times point zero two times four, which is forty eight thousand dollars.
- Year two gives point six times point zero two five times five, seventy five thousand.
- Year three gives point six times point zero three times three, fifty four thousand.
- The C V A is the sum, one hundred and seventy seven thousand dollars.
- That is how much value the swap loses because Corvane might fail.

9 Value after both adjustments

Adjusted value = V − CVA + DVA

Current value, V	No default value of the contract
CVA	Counterparty's default: subtract
DVA	Your own default: add
Adjusted value	What the contract is really worth

IN THE LESSON

- Writing V for the no default value of the contract, the value after both adjustments is V minus C V A plus D V A.
- The signs deserve care.
- C V A is a loss to you, because the counterparty may fail, so it is subtracted.
- D V A is a gain to you, so it is added.
- Why a gain?
- If you default, part of what you owe is never paid, so your liability is worth less.
- This is why D V A rises when your own credit spread widens: the market thinks you are more likely to fail, and your liabilities are worth less.
- It feels odd that bad news about yourself is a gain, but that is the arithmetic.

10 Example 3: adjusting the swap's value

V = $2,000,000. Marlow q: 1%, 1.5%, 2.5%. Exposure to Corvane PV: $3m, $2m, $1m

CVA (from Example 2)	Corvane's default	**$177,000**
DVA	0.6 × (1%×3 + 1.5%×2 + 2.5%×1)m	**$51,000**
Adjusted value	2,000,000 − 177,000 + 51,000	**$1,874,000**
Wrong: subtract DVA	2,000,000 − 177,000 − 51,000	**$1,772,000**

IN THE LESSON

- Return to the swap.
- Its no default value to Marlow is two million dollars, and the C V A is one hundred and seventy seven thousand.
- Now D V A, from Marlow's side.
- Marlow's own default chances are one per cent, one and a half per cent and two and a half per cent, and the present values of the exposure it creates for Corvane are three, two and one million dollars.
- With recovery of forty per cent, D V A is point six times the sum of point zero one times three, point zero one five times two and point zero two five times one, which is fifty one thousand dollars.
- The adjusted value is two million minus one hundred and seventy seven thousand plus fifty one thousand, one million eight hundred and seventy four thousand.
- Subtracting D V A by mistake gives one million seven hundred and seventy two thousand.

11 Spreads and default probability: the idea

- Risky bond yield minus risk-free yield = credit spread
- The spread pays for the expected loss from default
- Higher default chance or lower recovery: wider spread
- Work backwards from the spread to the hazard rate

IN THE LESSON

- Where do the default probabilities in a C V A come from?
- One source is the market for bonds.
- Suppose Corvane has a bond that pays a higher yield than a risk free bond of the same maturity.
- The extra yield is the credit spread, and it compensates the investor for the loss expected from default.
- The higher the chance of default, and the lower the recovery, the bigger the spread has to be.
- So a spread carries information about default probability, and we can work backwards from it.
- To do that, we need a measure of how fast default arrives, and that is the hazard rate, the annual default intensity.

12 Spread, hazard rate and default probability

Hazard rate = Spread ÷ (1 − R)
PD by T = 1 − e^(−λT)

Credit spread	Yield over the risk-free rate
Recovery rate, R	Share recovered after default
Hazard rate, lambda	Annual default intensity
Default probability by T	Survival is e^(−λT)

IN THE LESSON

- The standard approximation says the spread equals the average hazard rate, lambda, times one minus the recovery rate R.
- Rearranged, the average hazard rate implied by an observed spread is the spread divided by one minus R.
- Over a horizon of T years, the cumulative probability of default is one minus the exponential of minus lambda times T.
- The exponential of minus lambda T is the survival probability over the same period, so default and survival add to one.
- If lambda is small and T short, this is close to lambda times T, but with a long horizon the exponential form stops the probability from ever passing one hundred per cent.

13 Example 4: Corvane's five year probability

Spread 3%, recovery 40%, horizon 5 years

Hazard rate	3% ÷ (1 − 40%)	**5% a year**
λ × T	5% × 5	**0.25**
Survival probability	e^(−0.25)	**77.88%**
Default probability	1 − 77.88%	**22.12%**

IN THE LESSON

- Corvane's five year bond yields three per cent more than the risk free rate, so the spread is three per cent.
- The average hazard rate is three per cent divided by sixty per cent, which is five per cent a year.
- Over five years, lambda times T is five per cent times five, which is point two five.
- The survival probability is the exponential of minus point two five, about seventy seven point eight eight per cent.
- So the cumulative probability of default is one minus that, about twenty two point one two per cent.
- The mistake to avoid is reading the spread, three per cent, as the default probability.
- It is a rate: divide by the loss given default, then compound over the horizon.

14 Example 5: a term structure of spreads

Spreads: 2 years 2.0%, 4 years 3.0%. Recovery 40%

λT at 2 years	(2% ÷ 0.6) × 2	**0.0667**
λT at 4 years	(3% ÷ 0.6) × 4	**0.2000**
Forward hazard, years 3-4	(0.2 − 0.0667) ÷ 2	**6.67%**
Conditional default chance	1 − e^(−0.1333)	**12.48%**

IN THE LESSON

- Now suppose we see spreads for two maturities.
- The two year spread is two per cent and the four year spread is three per cent, with recovery still forty per cent.
- The two year average hazard rate is two divided by point six, three point three three per cent, so lambda T is point zero six six seven.
- The four year average hazard rate is five per cent, so lambda T is point two.
- The difference, point one three three three, is the hazard accumulated over two years, so the forward hazard rate is six point six seven per cent.
- The conditional chance of defaulting in years three and four, given survival to year two, is one minus the exponential of minus point one three three three, twelve point four eight per cent.

15 Risk-neutral probabilities and the exam trap

- Spreads give risk-neutral probabilities
- They exceed historical, real-world default rates
- The spread also pays risk premium, liquidity and tax
- Strip out non-credit parts before converting

IN THE LESSON

- A last warning.
- Probabilities backed out of spreads are risk neutral.
- They are the probabilities that, together with market prices, make the numbers fit.
- They are usually larger than the real world probabilities seen in historical default records, because a spread contains more than expected default loss.
- It also pays for risk premium, because investors dislike bearing default risk, for liquidity, because the bond may be hard to sell, and for tax effects.
- So never read the spread as the default probability, and strip out the non credit components before converting.
- For pricing a C V A, risk neutral probabilities are the consistent choice.
- For forecasting actual defaults, real world ones are the better guide.

16 Credit risk mitigants

- They cut the loss, not the chance of default
- Close-out netting: contracts combine into one net claim
- Collateral: cash or securities, after a haircut
- Downgrade trigger: demand collateral or end on downgrade

IN THE LESSON

- Marlow cannot easily change how likely Corvane is to default.
- What it can change is how much it loses if Corvane does.
- That is what credit risk mitigants do.
- They reduce the loss from a given counterparty without changing its probability of default.
- There are three main ones.
- Close out netting lets all the contracts with one counterparty be terminated and combined into one net claim, so positive and negative values offset.
- Collateral, usually cash or securities under a margin agreement, covers part of the claim, after a haircut that reduces the collateral's value.
- Downgrade triggers let you demand collateral, or end the contracts, if the counterparty's rating falls below a set level.

17 Mitigants side by side

Mitigants reduce loss, and each leaves a residual risk

	Netting	Collateral
Acts on	Exposure	Unsecured exposure
Effect	Offsets values	Covers a claim
Needs	Legal agreement	Margin agreement
Residual risk	Legal enforceability	Volatility, delay
Order of use	First	Second, after haircut

IN THE LESSON

- Set the first two mitigants side by side.
- Netting acts on the exposure: it replaces the sum of positive values with the single net value, which is never larger.
- Collateral acts on what remains: it lowers the unsecured exposure, but only by the haircut value of what is held.
- Downgrade triggers work by timing, since they let you act early.
- In credit analysis all three act on exposure at default and loss given default, never on the default probability.
- Each also has residual risks.
- Netting depends on legal enforceability in the counterparty's country.
- Collateral can fall in value, or arrive late.
- And there is wrong way risk: if the collateral is shares of Corvane itself, it loses value just when Corvane fails.

18 Loss after mitigation

Loss = E × (1 − R)
E = net value − collateral after haircut, floored at zero

Unsecured exposure, E	After netting and collateral
Recovery rate, R	Applied to what remains
Loss	Credit loss if it defaults

IN THE LESSON

- Here is how the mitigants enter the loss.
- Let E be the unsecured exposure after netting and collateral, and R the recovery rate.
- The loss is E times one minus R.
- The order matters, and it is the exam point.
- First apply netting to get the net claim.
- Then apply the collateral, after its haircut, to get E.
- Only then apply recovery, and only to what is left.
- Do not apply recovery to the whole claim first.
- Recovery applies only to the part that is still unsecured.
- Mitigants shrink E and sometimes the loss given default.
- They leave the default probability alone.

19 Example 6: netting and collateral at work

Trades +$14m, −$6m, +$5m. Collateral $10m, haircut 20%. Recovery 40%

Net claim	14 − 6 + 5	**$13m**
Collateral after haircut	10 × (1 − 20%)	**$8m**
Unsecured exposure E	13 − 8	**$5m**
Loss	5 × (1 − 40%)	**$3.0m**
With no mitigants	19 × (1 − 40%)	**$11.4m**

IN THE LESSON

- Marlow has three trades with Corvane under an enforceable netting agreement.
- The values to Marlow are plus fourteen, minus six and plus five million dollars.
- Without netting, Marlow would count only the positive trades, nineteen million.
- With netting the claim is fourteen minus six plus five, thirteen million.
- Corvane has posted collateral with a market value of ten million, and the haircut is twenty per cent, so the collateral counts as eight million.
- The unsecured exposure is thirteen minus eight, five million.
- Recovery is forty per cent, so the loss on default is five million times sixty per cent, three million dollars.
- With no netting or collateral, the loss would be nineteen million times sixty per cent, eleven point four million.

20 Why default correlation matters

- Expected loss depends only on individual default chances
- Correlation decides how likely defaults are to cluster
- The tail, economic capital and senior tranches rise with it
- Correlation changes the tail, not the mean

IN THE LESSON

- Until now we looked at one counterparty.
- A bank has thousands, and what matters is whether they fail together.
- Default correlation measures how likely defaults are to cluster.
- It does not affect the portfolio's expected loss, which depends only on individual default probabilities.
- It matters for the tail.
- If defaults are independent, the number of failures stays near its average.
- If they cluster, there are years with very few defaults and years with very many, and the bad years are far worse.
- So the tail of the loss distribution, and with it the economic capital the bank must hold, rises sharply with correlation.
- So does the risk of senior tranches, which are hurt only when many names fail at once.

21 Joint default probability

Joint = PDA × PDB + ρ × √(PDA(1 − PDA) × PDB(1 − PDB))

Chance A defaults, PDA	Default chance of name A
Chance B defaults, PDB	Default chance of name B
Correlation, rho	Default correlation of the pair
Joint default probability	Chance both default

IN THE LESSON

- For two names, A and B, the chance that both default over the same period depends on their default correlation.
- Start from the independent answer, the default probability of A times that of B.
- Then add a term for correlation: rho times the square root of P D A times one minus P D A, times P D B times one minus P D B.
- With zero correlation the second term disappears and we get the independent product.
- With positive correlation the joint probability is higher, so defaults cluster.
- The root term is the product of the two standard deviations of the yes or no default events, which is why it appears.
- This is the starting point for any pair of names in a portfolio.

22 Example 7: Corvane and a second miner

PD of Corvane 4%, Delmar 3%, correlation 0.2

Independent joint	4% × 3%	**0.12%**
Root term	√(0.04×0.96×0.03×0.97)	**0.03343**
Correlation term	0.2 × 0.03343	**0.67%**
Joint probability	0.12% + 0.67%	**0.79%**
Versus independent	0.79% ÷ 0.12%	**6.6 times**

IN THE LESSON

- Take Corvane, with a four per cent default probability, and a second miner, Delmar Resources, with three per cent.
- If they were independent, the chance that both default is point zero four times point zero three, which is point one two per cent.
- Now suppose their default correlation is point two.
- The square root term is the root of point zero four times point nine six times point zero three times point nine seven.
- That is about point zero three three four.
- Multiply by point two and the correlation term is about point six seven per cent.
- The joint default probability is point one two per cent plus point six seven per cent, about point seven nine per cent, more than six times the independent figure.
- Neither single default probability changed.

23 Reduced form and structural models

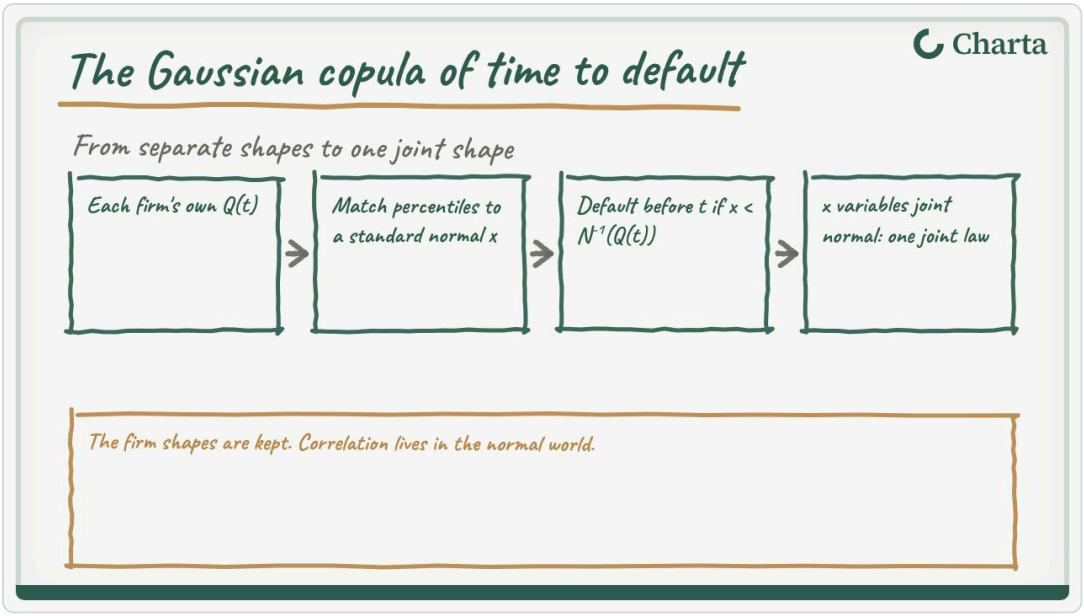
Two ways to model default and its correlation

	Reduced form	Structural
Default is	Random, at a hazard rate	Assets fall below debt
Correlation from	Linked intensities	Correlated assets
Strength	Fits market spreads	Economic intuition
Weakness	Low unless intensities swing	Needs balance sheets

IN THE LESSON

- There are two families of models for default correlation.
- Reduced form models treat default as a random event that arrives at a rate, the hazard rate we met earlier.
- They create correlation by linking the intensities or the default times of different names.
- They calibrate to market spreads, which is a strength, but they tend to give low default correlation unless the intensities are very volatile.
- Structural models are of the Merton type.
- They say a firm defaults when the value of its assets falls below its debt, and the correlation comes from correlated asset values.
- They have economic intuition, but they need balance sheet data.
- Remember the exam point: correlation changes the tail, not the mean.

24 The Gaussian copula of time to default



IN THE LESSON

- Now a standard tool for joining separate default distributions into one.
- Each firm has its own cumulative default probability, Q i of t, the chance that it defaults before time t.
- These are estimated one firm at a time and can have any shape.
- The Gaussian copula maps each one onto a standard normal variable by matching percentiles.
- Firm i defaults before t when its normal variable x i is below the inverse normal of Q i of t.
- So a firm with a five per cent chance of defaulting is mapped to the point below which five per cent of the normal curve lies.
- Then the dependence between the x variables is set as multivariate normal.
- The copula joins the separate shapes, and the normal world carries the correlation.

25 The one-factor Gaussian copula

$x = \sqrt{\rho} \times M + \sqrt{(1 - \rho)} \times Z$
PD given $F = N[(N^{-1}(Q(T)) - \sqrt{\rho} \times F) \div \sqrt{(1 - \rho)}]$

Firm's normal variable, x	Default if below $N^{-1}(Q(T))$
Common factor, M	Shared by every firm
Own shock, Z	Different for each firm
Correlation, rho	Root rho multiplies F
Cumulative default chance, Q	The firm's own estimate

IN THE LESSON

- To make this practical, the one factor version builds each firm's variable from two pieces.
- M is a common factor, the state of the economy shared by every firm.
- Z is an independent shock, specific to that firm.
- Both are standard normal.
- The variable x is root rho times M plus root one minus rho times Z, and every pair of firms then has correlation rho.
- Now fix the common factor at a value F.
- The probability that a firm defaults by T is N of the quantity: inverse N of Q of T, minus root rho times F, all divided by root one minus rho.
- A negative F is a bad economy and raises the probability.
- Take care: root rho multiplies F, not rho itself.

26 Example 8: Corvane in good and bad times

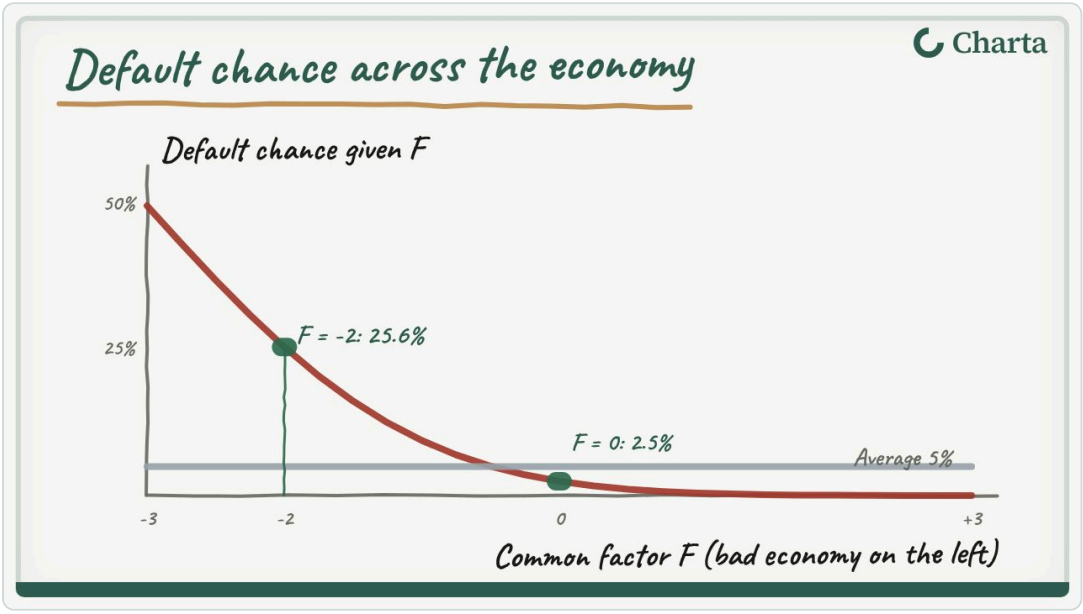
Q = 5%, ρ = 0.3. Compare F = −2 with F = +2

Threshold	$N^{-1}(5\%)$	**−1.645**
Bad, F = −2	(−1.645 + 0.548×2) ÷ 0.837	**−0.657**
Default chance, bad	N(−0.657)	**25.57%**
Default chance, good (F = +2)	N((−1.645 − 1.095) ÷ 0.837)	**0.05%**

IN THE LESSON

- Corvane has a cumulative five per cent chance of defaulting by the horizon, and the copula correlation is point three.
- The default threshold is the inverse normal of point zero five, minus one point six four five.
- Root point three is point five four eight, and root point seven is point eight three seven.
- In a bad economy, F is minus two.
- The numerator is minus point five five, divided by point eight three seven gives minus point six five seven, and the normal cumulative value is twenty five point five seven per cent.
- In a good economy, F is plus two.
- The numerator is minus one point six four five minus one point zero nine five, minus two point seven four, divided by point eight three seven, minus three point two seven five, so about point zero five per cent.
- The conditional chance swings enormously around the five per cent average.

27 Default chance across the economy



IN THE LESSON

- This chart shows the conditional default probability for Corvane as the common factor moves.
- Along the bottom is F, the state of the economy, from minus three on the left, a severe recession, to plus three on the right, a boom.
- Up the side is the chance that Corvane defaults, given that value of F.
- The grey line is the unconditional average of five per cent.
- The red curve falls steadily from the left.
- At minus two it is above twenty five per cent.
- At the middle, F equal to zero, it is below the average, about two and a half per cent.
- In a boom it is nearly zero.
- Default risk is concentrated in the bad states, and that is exactly the clustering that correlation creates.

28 Credit VaR from the Gaussian copula

$WCDR = N\left[\left(N^{-1}(PD) + \sqrt{\rho} \times N^{-1}(X) \right) \div \sqrt{(1 - \rho)} \right]$
$UL = EAD \times LGD \times (WCDR - PD)$

Worst case default rate	Default rate in a bad year
Confidence level, X	The percentile of the worst case
Correlation, rho	Copula correlation, any pair
Unexpected loss	Loss at the level, minus expected

IN THE LESSON

- Both methods turn default correlation into a loss distribution and read credit value at risk as a percentile.
- The first is the Gaussian copula, the Vasicek approach, which suits a large portfolio of similar loans.
- Let rho be the copula correlation and X the confidence level.
- The worst case default rate at that level is N of the quantity: inverse N of P D, plus root rho times inverse N of X, all divided by root one minus rho.
- The loss at that level is exposure at default times loss given default times that rate.
- This is a default only approach, and credit value at risk subtracts the expected loss, so the unexpected loss is E A D times L G D times the gap between the worst case rate and P D.

29 Example 9: credit VaR of Corvane's loan book

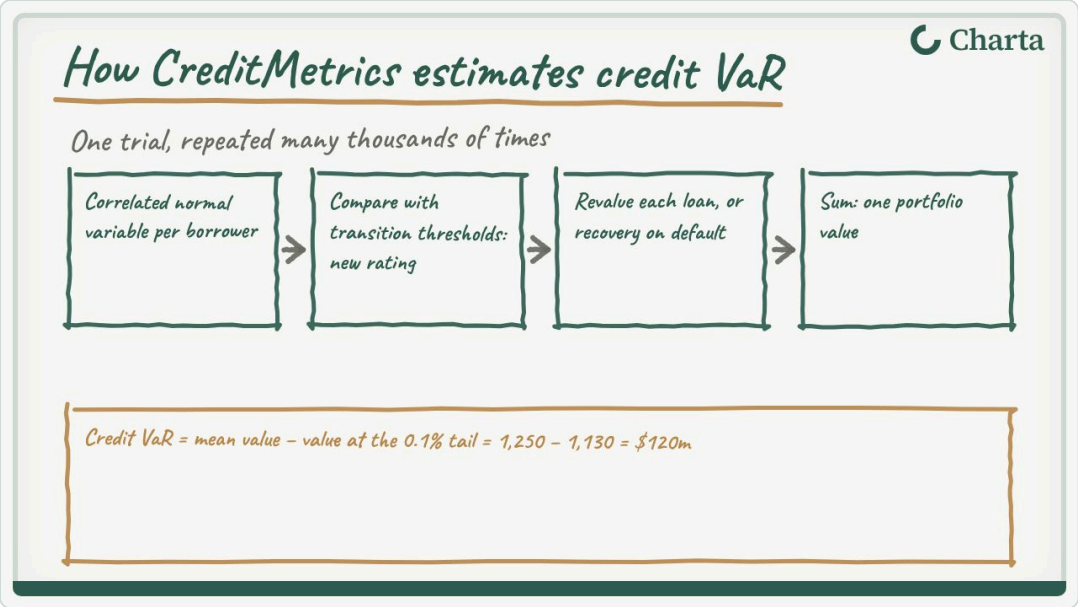
Exposure $800m, PD 3%, LGD 55%, ρ = 0.15, level 99.9%

Numerator	$-1.881 + 0.387 \times 3.090$	**−0.684**
Divide by √0.85	$-0.684 \div 0.922$	**−0.742**
Worst case default rate	$N(-0.742)$	**22.91%**
Unexpected loss	$800 \times 0.55 \times (22.91\% - 3\%)$	**$87.6m**

IN THE LESSON

- Marlow has eight hundred similar loans, total exposure eight hundred million dollars.
- Each has a default probability of three per cent, loss given default of fifty five per cent, and the copula correlation is point one five.
- We want the ninety nine point nine per cent level.
- The inverse normal of point zero three is minus one point eight eight one.
- Root point one five is point three eight seven, times three point zero nine for the inverse normal of point nine nine nine, adds one point one nine seven.
- The numerator is minus point six eight four.
- Divided by root point eight five, point nine two two, that is minus point seven four two.
- The worst case default rate is the normal cumulative value, twenty two point nine one per cent.
- Unexpected loss is eight hundred million times point five five times the gap between twenty two point nine one and three per cent, about eighty seven point six million dollars.

30 How CreditMetrics estimates credit VaR



IN THE LESSON

- The Gaussian copula formula counts only defaults.
- CreditMetrics goes further, because a loan can lose value when the borrower is merely downgraded.
- Follow the boxes.
- A transition matrix gives the probabilities of each rating moving to every other rating, or to default, over the horizon, usually one year.
- CreditMetrics simulates correlated normal variables for the borrowers and compares each one with thresholds taken from the transition matrix, to give a new rating.
- Each position is then revalued at that rating, or at recovery value on default.
- Many trials give a distribution of portfolio values.
- Credit value at risk is the percentile of value below the mean.
- If the mean is one thousand two hundred and fifty million dollars and the bottom one in a thousand outcome is one thousand one hundred and thirty million, the figure is one hundred and twenty million.

31 Gaussian copula versus CreditMetrics

Two ways to turn correlation into credit VaR

	Copula (Vasicek)	CreditMetrics
Events	Defaults only	Defaults and migration
Portfolio	Large, similar loans	Any mix of ratings
Method	One-factor formula	Monte Carlo trials
Result	WCDR, then UL	Percentile below mean
Assumes	Granular, one factor	Transition matrix

IN THE LESSON

- Put the two side by side.
- Both start from default correlation and both read a percentile of the loss distribution.
- The Gaussian copula, in its Vasicek form, is a closed formula for a large portfolio of similar loans.
- It assumes granularity, meaning no single loan is large, and it assumes one common factor.
- It counts defaults only, and it measures the unexpected loss as the worst case loss minus the expected loss.
- CreditMetrics is a simulation.
- It can handle a mixed portfolio of different ratings and sizes, and it captures migration losses, the fall in value when a loan is downgraded.
- The exam point: CreditMetrics captures migration, and the formula does not.

32 Recap: the main ideas

- Exposure = max(V, 0): replacement value, never notional
- Adjusted = V − CVA + DVA; spread = hazard × (1 − R)
- Netting, collateral, triggers cut E and LGD, not PD
- Correlation changes the tail; joint default rises with ρ
- Copula/Vasicek: defaults only; CreditMetrics adds migration

IN THE LESSON

- In summary.
- A derivative creates credit risk only when its value to you is positive, so exposure is the larger of V and zero, measured by replacement value and never by notional.
- C V A is the present value of expected loss from the counterparty failing, D V A is the mirror image from your own default, and the adjusted value is V minus C V A plus D V A.
- A spread gives the hazard rate by dividing by one minus recovery.
- Netting, collateral and downgrade triggers cut exposure and loss, not default probability.
- Correlation changes the tail, not the mean.
- Reduced form models use hazard rates, structural ones use asset values.
- The one factor copula gives conditional default chances, and Vasicek and CreditMetrics turn them into credit value at risk.

33 Try three quick questions

1. Spread 2.4%, recovery 40%. Hazard rate? One-year PD?
   Answer: 2.4% ÷ 0.6 = 4%. PD = 1 − e^(−0.04) = 3.92%.
2. Net claim $30m, collateral after haircut $12m, recovery 25%. Loss?
   Answer: E = 18m. Loss = 18 × 0.75 = $13.5m.
3. V = $6m, CVA = $0.5m, DVA = $0.2m. Adjusted value?
   Answer: 6 − 0.5 + 0.2 = $5.7m.

IN THE LESSON

- Now try three quick questions.
- Question one.
- A bond has a spread of two point four per cent and the recovery rate is forty per cent.
- Hazard rate, and one year default probability?
- The hazard rate is two point four divided by point six, which is four per cent.
- The one year probability is one minus the exponential of minus point zero four, about three point nine two per cent.
- Question two.
- A counterparty's trades net to thirty million dollars, collateral after the haircut is twelve million, and recovery is twenty five per cent.
- What is the loss?
- Exposure is eighteen million, and the loss is eighteen times point seven five, thirteen point five million.
- Question three.
- A contract's no default value is six million, C V A is half a million and D V A is two hundred thousand.
- Adjusted value?
- Six minus point five plus point two is five point seven million dollars.