CHARTA LEARNING · LESSON NOTES

Measuring Credit Risk

FRM Part I

1 A story: the bank and the bad year

- Normal years: a few loans fail, and pricing covers it
- Rare bad year: many fail together, losses jump
- Capital is the cushion of the bank's own money
- Credit risk measurement sizes that cushion

IN THE LESSON

- Imagine a small bank with a thousand loans.
- In most years a few borrowers fail, the bank expects that, and it prices the loans so that the interest it earns covers those failures.
- That is the normal cost of doing business.
- But once in a long while a recession hits, many borrowers fail together, and the losses are far bigger than usual.
- If the bank has no cushion, it collapses, and its depositors and lenders are hurt.
- So the bank holds a cushion of its own money, called capital.
- The big question is simple.
- How large should that cushion be?
- The rest of this lesson is the answer.

2 Economic capital

Economic capital = Loss at confidence level − Expected loss

Loss at confidence level	Loss exceeded in 0.03% of years
Expected loss	The average loss, paid by pricing
Economic capital	Equity held against the rest

IN THE LESSON

- Economic capital is the bank's own answer to the cushion question.
- The bank picks a horizon, normally one year, and a confidence level.
- It models the distribution of its possible losses, across credit, market and operational risk, allowing for diversification between them.
- Then it reads off the loss at the chosen confidence level and subtracts the loss it already expects, because expected loss is paid for by pricing and provisions.
- What remains is the unexpected part, which equity must absorb.
- The confidence level is tied to the credit rating the bank wants.
- A bank targeting a double A rating, whose chance of default in a year is about three hundredths of one per cent, would use roughly the ninety nine point nine seven per cent level.

3 Example 1: sizing economic capital

Illustration: one-year loss, 99.97% level. Amounts in $ millions

Loss at 99.97% level	from the loss model	**640**
Expected loss	covered by pricing	**90**
Economic capital	640 − 90	**550**

IN THE LESSON

- Here is a simple case.
- Suppose the bank's model of one-year losses says that in all but three out of ten thousand years, the loss is below six hundred and forty million.
- That is the ninety nine point nine seven per cent level, matching the double A target.
- The average, or expected, loss is ninety million.
- Economic capital is the difference.
- Six hundred and forty minus ninety is five hundred and fifty million.
- Notice that the bank does not hold capital against the ninety million, because it expects it and charges customers for it.
- The five hundred and fifty million is the cushion for the unexpected.
- Everything later in this lesson is about measuring these two ingredients properly.

4 Economic versus regulatory capital

Two different numbers for two different purposes

	Economic | Regulatory
Set by	The bank | Supervisor
Method	Own models | Prescribed formulas
Applies to	One bank | All banks alike
Used for	Pricing, RAROC, allocation | Minimum floor

IN THE LESSON

- Regulatory capital is a different thing, and the exam trap is to treat the two as the same number.
- Regulatory capital is the minimum that supervisors require.
- They set it with prescribed formulas and parameters, and the same rules apply to every bank.
- Economic capital belongs to one bank.
- It uses that bank's own models, its own mix of loans, and its own chosen rating target.
- Because it is specific to the bank, it is the number used for pricing a loan, for judging how well a business unit performs, for example with a measure called R A R O C, which compares profit with capital used, and for deciding where capital should be allocated.
- The two numbers can differ in either direction.

5 Why defaults move together

- Borrowers share one economy: recession hits many at once
- Pairwise default correlation is small: a few per cent or less
- Independent defaults would give a steady yearly rate
- Real rates swing from far below to several times average

IN THE LESSON

- Now we ask how strongly loan defaults depend on one another.
- Borrowers share the same economy.
- In a recession many firms weaken at once, and in a boom very few fail.
- So defaults are positively correlated.
- The surprise is how small the number is.
- The correlation between the default events of any two borrowers is only a few per cent, or even less.
- It sounds negligible, but across thousands of loans it matters enormously.
- If defaults were truly independent, a large portfolio would show nearly the same default rate every year, because the luck would average out.
- In reality, observed default rates swing from well below the average to several times above it.

6 Default rate variance with correlation

Var(default rate) = p(1 − p) × [1/n + (n − 1)/n × ρ]
Large n: Var → ρ × p(1 − p)

Default probability, p	Chance one loan defaults
Number of loans, n	How many loans in the pool
Correlation, rho	Pairwise default correlation
Variance of default rate	The spread of yearly default rate

IN THE LESSON

- Here is how to measure that swing.
- With independent loans, the standard deviation of the default rate is the square root of the fraction p times one minus p, over n.
- It shrinks as n grows.
- Now add a pairwise correlation rho.
- The variance of the default rate becomes p times one minus p, times the sum of two parts: one over n, plus n minus one over n times rho.
- The first part vanishes as the portfolio grows.
- The second part does not.
- As n becomes very large, the variance approaches rho times p times one minus p, which is not zero.
- That floor is the risk that no amount of diversification across loans can remove.

7 Example 2: independence versus correlation

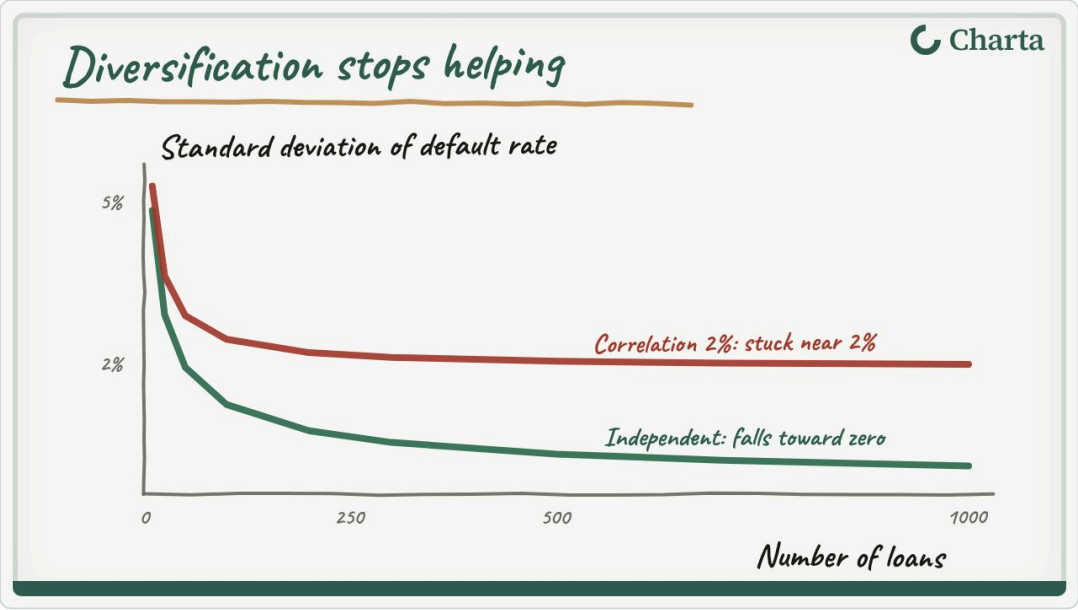
1,000 loans, p = 2%. Compare ρ = 0 with ρ = 2%

Independent: SD of rate	√(0.02×0.98 ÷ 1000)	**0.44%**
Correlated: SD of rate	√(0.0196 × 0.02098)	**2.03%**
Floor for huge n	√(0.02 × 0.0196)	**1.98%**
Ratio	2.03 ÷ 0.44	**4.6 times**

IN THE LESSON

- Take a thousand loans, each with a two per cent chance of default.
- If they were independent, the variance of the default rate is point zero two times point nine eight, divided by a thousand.
- The standard deviation is about point four four per cent.
- Now let the correlation be two per cent.
- The bracket becomes one over a thousand plus nine hundred and ninety nine over a thousand times point zero two, which is about point zero two one.
- Multiply by point zero one nine six and take the root: the standard deviation is about two per cent, nearly five times larger.
- Correlation did not change the average default rate, but it fattened the tail.

8 Diversification stops helping



IN THE LESSON

- This chart shows the same point as a picture.
- Along the bottom is the number of loans, from a few up to a thousand.
- Up the side is the standard deviation of the yearly default rate.
- The green line is the independent case.
- It falls steadily as loans are added, because bad luck on one loan is cancelled by good luck on another.
- The red line has a pairwise correlation of two per cent.
- It starts at nearly the same place, but it flattens out near two per cent and refuses to fall further.
- After a few hundred loans, adding more barely helps.
- The lesson for capital is clear: it cannot be set on the assumption that defaults are independent.

9 Expected loss

EL = PD × LGD × EAD
LGD = 1 − recovery rate

Probability of default	Chance the borrower defaults
Loss given default	Share not recovered
Exposure at default	Amount owed at default

IN THE LESSON

- Now to measuring the loss a bank should expect.
- Expected loss is the average credit loss over a horizon, usually one year.
- For a single exposure it is the product of three things.
- The probability of default, which is how likely the borrower fails.
- The loss given default, which is the share of the exposure that is not recovered, so it equals one minus the recovery rate.
- And the exposure at default, the amount owed when failure happens.
- For a facility that is only partly drawn, the exposure at default is the drawn amount plus a conversion factor times the undrawn commitment, because borrowers often draw more as they get into trouble.

10 Example 3: expected loss of a credit line

Limit $10m: $6m drawn, $4m undrawn. CCF 50%, PD 2%, recovery 40%

Exposure at default	6m + 50% × 4m	**$8m**
Loss given default	1 − 40%	**60%**
Expected loss	2% × 60% × 8m	**$96,000**
Wrong: recovery for LGD	2% × 40% × 8m	**$64,000**

IN THE LESSON

- A company has a ten million dollar credit line.
- It has drawn six million, and four million is undrawn.
- The bank estimates that half of the undrawn part would be used by the time of default, so the conversion factor is fifty per cent.
- Exposure at default is six plus half of four, which is eight million.
- If lenders typically recover forty per cent, loss given default is sixty per cent.
- With a two per cent default probability, expected loss is point zero two times point six times eight million, which is ninety six thousand dollars.
- The common mistake is to use the recovery rate in place of loss given default.
- Here that would give sixty four thousand dollars, a third too low.

11 Expected loss: adding up and the spread

Second loan: $5m, PD 1%, LGD 45%. Spread 1.2%, recovery 40%

Second loan EL	1% × 45% × 5m	**$22,500**
Both loans	96,000 + 22,500	**$118,500**
PD from spread	1.2% ÷ (1 − 40%)	**2%**
Three-year cumulative PD	$1 - 0.98^3$	**5.88%**

IN THE LESSON

- Three more facts about expected loss.
- First, expected losses add across loans, whatever the correlation between them, because averages are linear.
- A second loan with five million exposure, a one per cent default probability and loss given default of forty five per cent has expected loss of twenty two thousand five hundred dollars.
- The pair has expected loss of one hundred and eighteen thousand five hundred.
- Second, a credit spread of one point two per cent with forty per cent recovery implies a default probability of about two per cent, since the spread is roughly probability times loss given default.
- Third, over several years use the cumulative default probability, about five point nine per cent over three years, not a simple sum.

12 Unexpected loss

- UL = loss at a high percentile, minus expected loss
- EL is covered by pricing and provisions
- UL is covered by equity capital
- Some texts quote UL as a standard deviation: check which

IN THE LESSON

- Expected loss is only the average.
- Real losses bounce around it, and the part that matters for survival is the bounce.
- Unexpected loss is the potential loss beyond what is expected.
- It is the amount by which the loss at a high confidence level, such as the ninety nine point ninth percentile over a year, exceeds the expected loss.
- This is exactly the quantity we called economic capital a few minutes ago.
- Expected loss is covered by pricing and provisions.
- Unexpected loss is covered by equity.
- One warning: some textbooks quote unexpected loss as the standard deviation of loss, which is a different measure from the high-percentile figure.
- Always check which meaning is in use.

13 Standard deviation of one loan's loss

UL (SD) = EAD × LGD × √(PD × (1 − PD))

Exposure at default	Amount owed at default
Loss given default	Share not recovered
Probability of default	Chance the borrower defaults

IN THE LESSON

- For a single loan with fixed loss given default and fixed exposure, there are only two outcomes in a year: default or no default.
- Here we use unexpected loss in its standard deviation sense, and it has a neat form.
- It is the exposure times the loss given default, times the square root of p times one minus p, where p is the default probability.
- The square root term is the standard deviation of a yes or no event.
- Be careful with this formula.
- It is the standard deviation for one exposure, treated as default or no default over one period.
- It is not a general portfolio formula, and it is not the regulatory capital formula.

14 Combining two standard deviations

UL (SD) of pair = √($\sigma_1^2 + \sigma_2^2 + 2\rho\sigma_1\sigma_2$)
Less than $\sigma_1 + \sigma_2$ unless $\rho = 1$

Two loans, sigma 1 and 2	Each loan's risk
Correlation, rho	Their default correlation
Combined risk	Risk of the pair

IN THE LESSON

- Expected losses add up across loans, but unexpected losses do not.
- Suppose two loans have standard deviations of loss sigma one and sigma two, and their correlation is rho.
- The combined standard deviation is the square root of sigma one squared, plus sigma two squared, plus two times rho times sigma one times sigma two.
- This is less than simply adding the two, unless the correlation is exactly one.
- That gap is the benefit of diversification.
- It reduces unexpected loss, and so it reduces the capital needed.
- It never reduces expected loss, because the expected loss of the pair is always the sum of the two expected losses.

15 Example 4: unexpected loss of two loans

Loan 1: $8m, LGD 60%, PD 2%. Loan 2: $5m, LGD 45%, PD 1%. ρ = 0.2

UL (SD), loan 1	4.8m × √0.0196	**$672,000**
UL (SD), loan 2	2.25m × √0.0099	**$223,872**
Combined SD, ρ = 0.2	$\sqrt{(\sigma_1^2 + \sigma_2^2 + 0.4\sigma_1\sigma_2)}$	**$749,586**
Plain sum	672,000 + 223,872	**$895,872**

IN THE LESSON

- Use our two loans.
- The first has exposure eight million, loss given default sixty per cent and default probability two per cent.
- Its unexpected loss is four point eight million times the square root of point zero one nine six, which is point one four, giving six hundred and seventy two thousand dollars.
- The second has exposure five million, loss given default forty five per cent and default probability one per cent.
- Two point two five million times the square root of point zero zero nine nine gives about two hundred and twenty four thousand.
- With a correlation of point two, the combined figure is about seven hundred and fifty thousand, below the plain sum of about eight hundred and ninety six thousand.

16 The binomial model of defaults

Defaults: mean = np, SD = √(np(1 − p))
Loss: mean = L·np, SD = L·√(np(1 − p))
At least one default: $1 - (1 - p)^n$

Loss per default, L	Exposure times LGD
Mean loss	Average loss
Standard deviation of loss	Spread of the loss
Normal multiplier, z	2.326 at the 99% level

IN THE LESSON

- The binomial model is the simplest model of a portfolio.
- Treat the portfolio as n independent trials, and give every loan the same default probability p.
- The number of defaults has a mean of n times p, and a variance of n times p times one minus p.
- The standard deviation is the square root of that variance.
- The chance of at least one default is one minus the quantity one minus p, to the power n.
- If every default costs the same amount L, which is exposure times loss given default, the loss has a mean of L times n times p, and a standard deviation of L times the square root of n p one minus p.
- For a big portfolio the shape is close to normal, so the loss at a confidence level is the mean plus z times the standard deviation, with z equal to two point three two six at ninety nine per cent.

17 Example 5: a binomial loss estimate

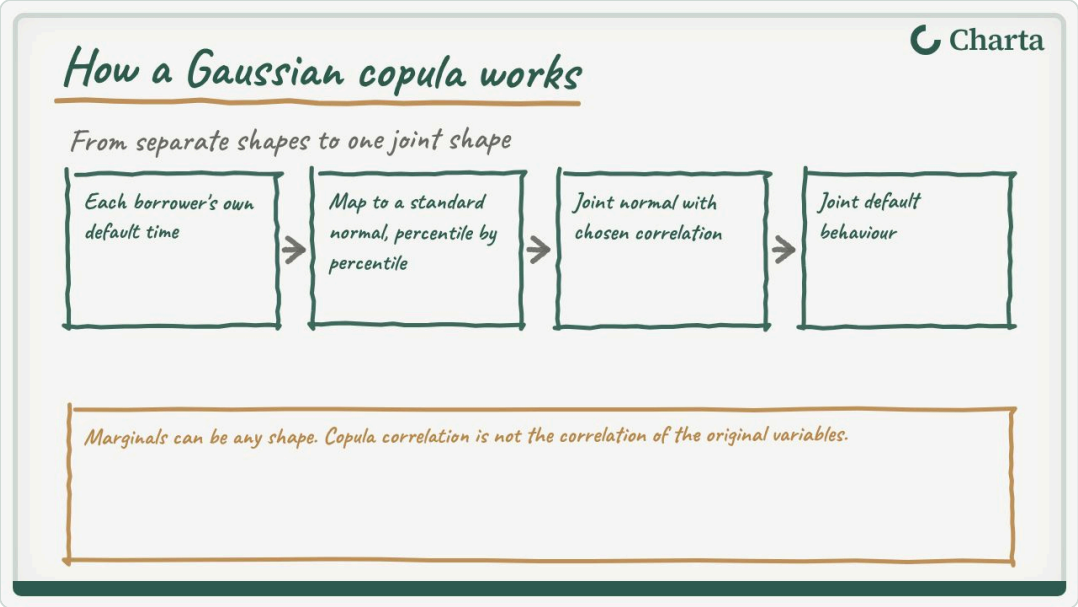
1,000 loans, p = 2%, loss $0.5m per default. Amounts in $ millions

Expected defaults	1000 × 2%	**20**
SD of defaults	√(1000×0.02×0.98)	**4.427**
Mean loss	0.5 × 20	**10.00**
SD of loss	0.5 × 4.427	**2.214**
99% loss	10 + 2.326 × 2.214	**15.15**

IN THE LESSON

- Take a thousand loans, each with a two per cent default probability and a loss of half a million dollars per default.
- The expected number of defaults is twenty.
- The variance is a thousand times point zero two times point nine eight, which is nineteen point six, so the standard deviation is about four point four three defaults.
- Multiply by the half million loss per default: the mean loss is ten million and the standard deviation is about two point two one million.
- At the ninety nine per cent level, add two point three two six standard deviations to the mean, giving about fifteen point one five million.
- The usual errors are forgetting the loss per default and reporting variance instead of standard deviation.

18 How a Gaussian copula works



IN THE LESSON

- The binomial model assumes independence, so it understates the spread and the tail.
- With a thousand loans and a correlation of just two per cent, the default rate varied by about two per cent, which is twenty defaults, not four.
- We need a way to bring correlation in, and that is a copula, a tool for joining separate distributions into one joint distribution while controlling the dependence.
- Look at the boxes from left to right.
- Start with each borrower's own distribution, for example the probability that the firm has defaulted by time t.
- These marginal distributions can have any shape.
- Map each one, percentile by percentile, onto a standard normal variable.
- Then assume those normal variables are jointly normal, with a chosen correlation.
- Finally, read the answer back: the dependence structure lives in the normal world, and the original shapes are kept.
- Note that the copula correlation is not the correlation of the original variables.

19 The one-factor model and its uses

$$X = \sqrt{\rho} \times M + \sqrt{(1 - \rho)} \times Z$$

Borrower's normal variable	Compared with a default threshold
Common factor, M	Shared by every borrower
Own shock, Z	Different for each borrower
Correlation, rho	Root rho: weight on common factor

IN THE LESSON

- The Gaussian copula is used to model default correlation among borrowers, to value credit products such as collateralised debt obligations and nth-to-default swaps, and inside portfolio models such as Vasicek and CreditMetrics.
- In the one-factor version each borrower's normal variable is built from two independent standard normal pieces.
- M is a common factor, the state of the economy.
- Z is the borrower's own shock.
- A weight of root rho goes on the common factor and root one minus rho on the own shock.
- Borrowers with a bad common factor tend to fail together.
- The main criticism is that the Gaussian shape gives weak dependence in the extreme tails, and a fixed correlation that proved too low in stress.

20 Example 6: default chance in a bad economy

PD 2%, ρ = 0.25, bad economy M = −2

Default threshold	$N^{-1}(2\%)$	**−2.054**
Common part	$\sqrt{0.25} \times (-2)$	**−1.000**
Own shock must be below	$(-2.054 + 1) \div \sqrt{0.75}$	**−1.217**
Chance of default	$N(-1.217)$	**11.2%**

IN THE LESSON

- Here is the copula at work.
- A borrower has a two per cent default probability.
- In the normal world, it defaults if its variable X falls below the point where the normal cumulative curve equals two per cent, which is minus two point zero five four.
- Let rho be point two five and suppose the economy is in a bad state, with M equal to minus two.
- The common part contributes root point two five, which is a half, times minus two, giving minus one.
- For X to fall below minus two point zero five four, the own shock must be sufficiently negative.
- The result is that the chance of default, given this bad economy, is about eleven per cent, more than five times the two per cent average.

21 The Vasicek worst case default rate

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

Worst case default rate	Default rate in a bad year
Confidence level, X	Use 99.9% in Basel
Correlation, rho	Zero gives WCDR equal to PD
Credit capital	Loss at the level, minus EL

IN THE LESSON

- The Vasicek model uses that one-factor setup for a large portfolio of similar loans.
- Every borrower has the same default probability and the same copula correlation rho with one common factor.
- Because the portfolio is large, each borrower's own shock averages out, so once the economy is fixed the default rate is fixed too.
- A bad economy is a low value of the factor, and its worst value at confidence X is minus inverse N of X, which becomes the plus sign below.
- At confidence X, the worst case default rate in a year is N of the quantity: inverse N of PD, plus root rho times inverse N of X, all divided by root one minus rho.
- Credit capital is then exposure times loss given default times the worst case rate minus PD.
- The Basel formula uses a confidence level of ninety nine point nine per cent.

22 Example 7: Vasicek capital

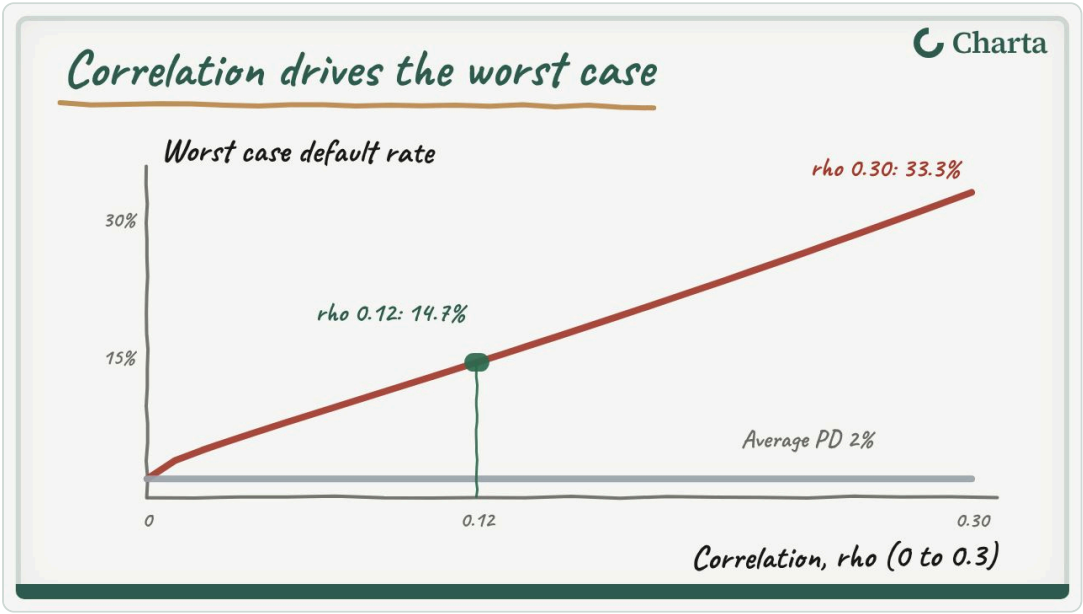
Portfolio $1,000m, PD 2%, ρ = 0.12, LGD 50%, confidence 99.9%

Numerator	−2.0537 + 0.3464 × 3.0902	**−0.9833**
Divide by √0.88	−0.9833 ÷ 0.9381	**−1.048**
Worst case default rate	N(−1.048)	**14.7%**
Capital	1000 × 0.5 × (14.7% − 2%)	**$63.6m**

IN THE LESSON

- Let us use it.
- A bank has a portfolio of one billion dollars with a default probability of two per cent, correlation point one two, and loss given default of fifty per cent.
- First, inverse N of point zero two is minus two point zero five four.
- Inverse N of point nine nine nine is positive three point zero nine, and root point one two is point three four six, so that term adds one point zero seven.
- The numerator is minus point nine eight.
- Divide by root point eight eight, which is point nine three eight, to get minus one point zero four eight.
- The normal cumulative value is fourteen point seven per cent.
- So credit capital is one billion times point five times the gap between fourteen point seven per cent and two per cent, about sixty three point six million dollars.

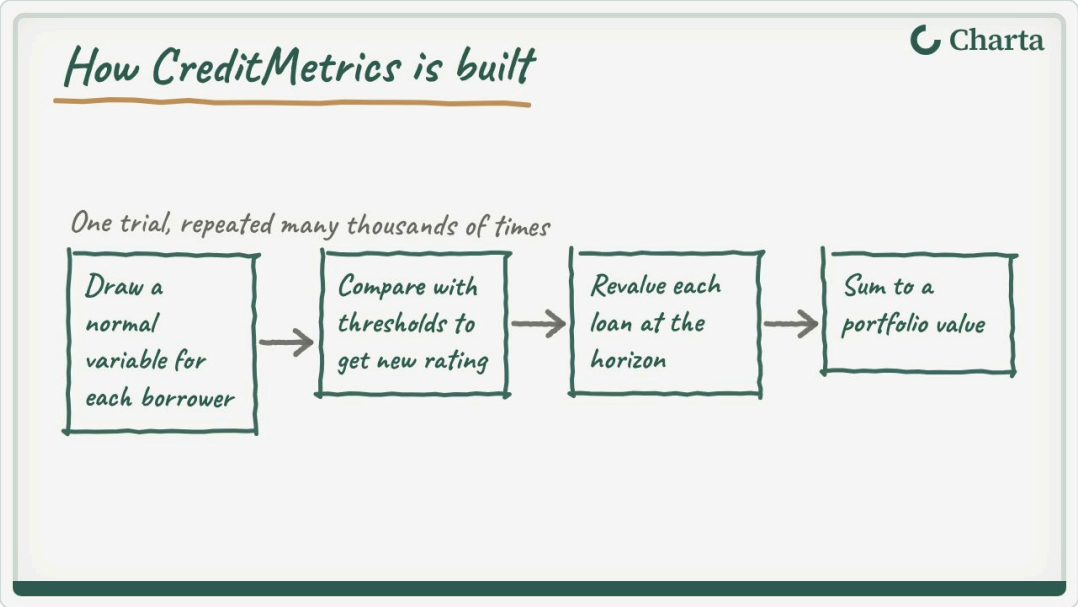
23 Correlation drives the worst case



IN THE LESSON

- This chart shows how sensitive the Vasicek worst case default rate is to the correlation.
- Along the bottom is rho, from zero to thirty per cent.
- Up the side is the worst case default rate at the ninety nine point nine per cent level, for a borrower with a two per cent default probability.
- With zero correlation, the line starts at two per cent: the worst case is just the average, because there is no common factor.
- Then it climbs steeply.
- At point one two it is almost fifteen per cent, and at point three it is over thirty three per cent.
- Correlation is an input nobody can observe directly, yet capital depends heavily on it.

24 How CreditMetrics is built



IN THE LESSON

- The Vasicek model counts only defaults.
- CreditMetrics goes further, because a loan can lose value when the borrower is merely downgraded.
- Follow the boxes.
- Each borrower has a credit rating, with one-year probabilities of moving to every other rating, including default.
- Correlation between borrowers comes from a factor model for their asset or equity returns, applied through the Gaussian copula.
- In each Monte Carlo trial, every loan gets a simulated normal variable, compared with thresholds to decide its new rating.
- Every loan is revalued at the horizon, and the portfolio is summed.
- After many trials we have a distribution of portfolio value, and economic capital is read from it.

25 Example 8: thresholds and one trial

BBB, illustrative: default 2%, BB 5%, BBB 88%, A 5%. $\rho = 0.2$

Thresholds	N^{-1} of 2%, 7%, 95%	**−2.05,−1.48,1.645**
Mild trial: M=−1.5, Z=0.5	0.447×−1.5 + 0.894×0.5	**−0.22**
Result of mild trial	between −1.48 and 1.645	**stays BBB**
Severe trial: M=−3, Z=−1	0.447×−3 + 0.894×−1	**−2.24 : default**

IN THE LESSON

- Here is a borrower rated triple B.
- Suppose its one-year probabilities are: default two per cent, a downgrade to double B five per cent, staying triple B eighty eight per cent, and an upgrade five per cent.
- These numbers are for illustration.
- The thresholds are inverse N of the cumulative probabilities from the bottom: minus two point zero five four for default, minus one point four seven six for double B or worse, and one point six four five for triple B or worse.
- With correlation point two, a mild trial gives X of minus point two two, so the borrower stays triple B.
- A severe trial gives minus two point two four, below the default threshold.
- A defaulted loan is worth only its recovery value.

26 CreditMetrics: from trials to capital

- Defaulted loan: worth its recovery value
- Other loans: discount at the new rating's spread
- Capital = mean value − value at the 99.9th percentile
- Unlike default-only models, it captures downgrade losses
- Example: 980 − 880 = capital of 100 ($ millions)

IN THE LESSON

- After the trials, the bank has thousands of possible portfolio values one year ahead.
- A defaulted loan is valued at its recovery value.
- A loan that migrated to a new rating is valued by discounting its remaining cash flows at the spread for that rating, so a downgraded loan is worth less even though it has not defaulted.
- Economic capital is then the mean portfolio value minus the value at the chosen percentile, for example the ninety nine point ninth.
- If the mean value is nine hundred and eighty million and the bottom one in a thousand outcome is eight hundred and eighty million, capital is one hundred million.
- A default-only model would miss the losses from downgrades.

27 Euler allocation of risk

Risk = Σ $x_i \times \partial R/\partial x_i$
Contribution of loan i = $x_i \times \mathrm{Cov}(L_i, L_P) \div \sigma_P$

Exposure to loan i	Amount lent to loan i
Rate of change of risk	Risk added per extra dollar
Contribution of loan i	Adds up exactly to total risk

IN THE LESSON

- Once we know the risk of the whole portfolio, how much does each loan contribute?
- Euler's theorem answers this.
- It applies to any risk measure that doubles when every position is doubled, which is called homogeneous of degree one.
- Standard deviation, value at risk and expected shortfall all qualify.
- Then total risk equals the sum, over all loans, of the exposure to the loan times the rate at which risk changes as that exposure changes.
- Each term is that loan's contribution, and the contributions add up exactly to the total.
- For standard deviation, the contribution of loan i is its exposure times the covariance of its loss with the portfolio loss, divided by the portfolio standard deviation.

28 Example 9: contributions to portfolio risk

$\sigma_1 = 3m$, $\sigma_2 = 4m$, $\rho = 0.5$. Amounts in $ millions

Portfolio SD	√(9 + 16 + 2×0.5×12)	**6.083**
Contributions	15 ÷ 6.083 and 22 ÷ 6.083	**2.466, 3.617**
Check: sum of contributions	2.466 + 3.617	**6.083**
Incremental risk, loan 2	6.083 − 3	**3.083**

IN THE LESSON

- Take two loans whose loss standard deviations are three and four million, with correlation one half.
- The portfolio variance is nine plus sixteen plus two times a half times three times four, which is thirty seven, so the portfolio standard deviation is about six point oh eight million.
- The covariance of loan one with the portfolio is nine plus six, fifteen, and for loan two it is sixteen plus six, twenty two.
- Dividing by six point oh eight gives contributions of two point four seven and three point six two, which sum exactly to six point oh eight.
- Compare incremental risk, the portfolio's risk with loan two minus without it, which is three point oh eight.
- It does not sum to the total.

29 Exposure on a derivative

Exposure = max(net value − collateral , 0)

Net value of the contracts	All trades in the netting set
Collateral held	Value posted by counterparty
Credit exposure	What can be lost if it defaults

IN THE LESSON

- Now consider derivatives.
- On a loan, the exposure is essentially known: it is the amount lent, and it falls as the borrower repays.
- A derivative is different.
- Its value depends on future market variables, so the exposure at default is uncertain and has to be simulated, and it evolves over the life of the contract.
- It can also be positive or negative.
- The bank only loses if the counterparty defaults while the contract is worth a positive amount to the bank, because a negative value must still be paid.
- Netting agreements, collateral and close-out terms change the exposure, and are modelled for each netting set, as the larger of the net value less collateral, and zero.

30 Example 10: netting and collateral

Three trades with one counterparty. Amounts in $ millions

Without netting	12 + 3	**15**
Net value	12 − 5 + 3	**10**
Collateral held	posted by counterparty	**4**
Exposure	max(10 − 4, 0)	**6**

Wrong-way risk: exposure peaks when default is likely

IN THE LESSON

- Suppose the bank has three trades with one counterparty, covered by a legally enforceable netting agreement.
- The trades are worth plus twelve, minus five and plus three million to the bank.
- Without netting, the bank would count only the positive trades, twelve plus three, which is fifteen million.
- With netting, the values offset: twelve minus five plus three is ten million.
- If the counterparty has posted four million of collateral, the exposure is ten minus four, which is six million.
- Notice there is also wrong-way risk: exposure may be highest just when the counterparty is most likely to default.
- So capital needs a joint model of market and credit risk, not the simple loan formulas.

31 Why credit risk is hard to measure

- Rare defaults mean few observations and noisy estimates
- A one-in-a-thousand-year loss cannot be observed
- Unobservable correlations that drive the answer
- Recoveries fall when defaults rise: stress LGD higher

IN THE LESSON

- Finally, why is all of this hard to get right?
- Defaults are rare, especially for high-quality borrowers, so estimates of default probabilities, correlations and recovery rates rest on few observations and are noisy.
- Capital models aim at a one-year, ninety nine point nine per cent loss, a one-in-a-thousand-year event, which we have never observed.
- Loans are not traded, so there are no market prices, asset correlations cannot be seen, and results are very sensitive to them.
- Recovery rates fall just when default rates rise, so average loss given default understates stress losses.
- The honest conclusion is to treat every output as an estimate with wide error bands.

32 Practical difficulties and what to do

Each weakness points to a response

Difficulty	Response
Cannot backtest a 1-in-1000 event	Stress tests and judgement
Procyclical PD-based capital	Through-the-cycle views
Big single names	Model concentration separately
Unobservable correlation	Sensitivity analysis

IN THE LESSON

- Here are more difficulties, each with a sensible response.
- Backtesting is nearly impossible: with twenty five years of data, a model that is badly wrong will usually show no exceedance of the ninety nine point nine per cent loss, just as a correct one will.
- Point-in-time default probabilities rise in recessions, so capital is procyclical, meaning it rises exactly when banks can least afford it.
- Large single-name concentrations break the granularity assumption of models such as Vasicek.
- Because model outputs are uncertain, test how sensitive the result is to each input, and do not rely on a single number with false precision.

33 Recap: the main ideas

- Economic capital = loss at level − EL; not regulatory
- EL = PD × LGD × EAD adds; UL needs correlation
- Binomial: independent. Copula: correlation. Vasicek: WCDR
- CreditMetrics adds downgrades; Euler shares out risk
- Derivative exposure is simulated; all outputs are estimates

IN THE LESSON

- Let us pull it together.
- Economic capital is the bank's own estimate of the cushion needed for unexpected losses at a confidence level, and it differs from regulatory capital.
- Defaults are mildly but positively correlated, which keeps the default rate volatile however many loans there are.
- Expected loss is probability times loss given default times exposure and is priced.
- Unexpected loss is the rest and is held in equity, and it does not simply add.
- The binomial model assumes independence, the Gaussian copula adds correlation, Vasicek gives a worst case default rate, and CreditMetrics adds downgrades.
- Euler shares risk across loans, derivatives need exposure simulation, and every estimate is uncertain.

34 Try three quick questions

1. Exposure $10m, recovery 30%, PD 3%. Expected loss?
   Answer: LGD 70%. 3% × 70% × 10m = $210,000.
2. 100 loans, p = 10%, independent. SD of defaults?
   Answer: √(100 × 0.1 × 0.9) = √9 = 3 defaults.
3. Why is the combined SD of loss below the sum of the two SDs?
   Answer: Correlation below one, so diversification helps.

IN THE LESSON

- Now try three quick questions.
- Pause the video, think, then listen to the answers.
- Question one.
- A loan has an exposure of ten million dollars, a recovery rate of thirty per cent and a default probability of three per cent.
- What is the expected loss?
- Loss given default is seventy per cent, so three per cent times seventy per cent times ten million is two hundred and ten thousand dollars.
- Question two.
- One hundred independent loans each have a default probability of ten per cent.
- What is the standard deviation of the number of defaults?
- The square root of one hundred times point one times point nine is the square root of nine, which is three defaults.
- Question three.
- Why is the combined standard deviation of loss below the sum of the two standard deviations?
- Because correlation is below one, so diversification helps.