CHARTA LEARNING · LESSON NOTES

Portfolio Credit Risk

FRM Part II

1 A story: two customers, one bad year

- Marlow Bank lends to Cedar Retail and Dune Shipping
- Worry: do they fail in the same year?
- Default indicator: 1 if the firm defaults, 0 if not
- Default correlation = correlation of the two indicators

IN THE LESSON

- Picture Marlow Bank, with two customers we will call Cedar Retail and Dune Shipping.
- Cedar fails in about one and a half years out of a hundred, and Dune in about four.
- The bank cares about more than each chance alone.
- It cares whether the two fail in the same year, because a recession that hits both brings two losses at once.
- So we need a number for how much two defaults go together.
- Over a chosen horizon, usually one year, record each borrower as a one if it defaults and a zero if not.
- The default correlation is the ordinary correlation between those two ones and zeros.

2 Default correlation

Corr = (JDP − p1×p2) ÷ √(p1(1−p1) × p2(1−p2))
Independent: JDP = p1 × p2, so Corr = 0
JDP = p1×p2 + ρ × √(p1(1−p1) × p2(1−p2))

Default probabilities p1, p2	Each firm's chance of default
Joint default probability	Chance both default together
Default correlation	Zero if defaults are independent
Correlation, rho	Another name for Corr here

IN THE LESSON

- Now the formula.
- Let p one and p two be the default probabilities and let J D P be the joint default probability, the chance both default in the same period.
- The default correlation, Corr, is J D P minus p one times p two, divided by the square root of p one times one minus p one, times p two times one minus p two.
- The top line is the surplus over independence.
- If defaults are independent, J D P is just p one times p two, the top line is zero, and so is the correlation.
- Rearranged, with rho the default correlation, J D P equals p one times p two plus rho times that same square root.

3 Example 1: Cedar and Dune

Cedar p1 = 1.5%, Dune p2 = 4%, JDP = 0.12%

Independent JDP	1.5% × 4%	**0.06%**
Numerator	0.12% − 0.06%	**0.06%**
Denominator	√(0.015×0.985×0.04×0.96)	**0.02382**
Default correlation	0.0006 ÷ 0.02382	**2.5%**
Reverse: Corr 5% gives JDP	0.0006 + 5% × 0.02382	**0.18%**

IN THE LESSON

- Now numbers.
- Cedar has a default probability of one point five per cent and Dune four per cent.
- Marlow's model says both fail in the same year with probability point one two per cent.
- Independence would give point zero six per cent, so the numerator is point zero zero zero six.
- The denominator is the square root of point zero one five times point nine eight five times point zero four times point nine six, about point zero two three eight two.
- The default correlation is two point five per cent.
- In reverse, a correlation of five per cent gives a joint probability of point zero zero zero six plus five per cent times point zero two three eight, about point one eight per cent.

4 Facts and traps about default correlation

Cedar and Dune: right answer and the usual errors

Point	Cedar and Dune
Correct default correlation	2.5%
Highest possible, PDs differ	60.5%, not 100%
Error: drop the p1 × p2 term	5.0%
Error: drop the square root	105.8%, impossible
Typical size in practice	A few per cent

IN THE LESSON

- Four facts for the exam.
- First, default correlations are small, often a few per cent, and well below the asset return correlation that drives them, because a yes or no event discards information.
- Second, they depend on the default probabilities and the horizon.
- Third, they cannot reach one when the probabilities differ.
- Here the joint probability can be at most one point five per cent, the smaller probability, which caps the correlation at about sixty point five per cent.
- Fourth, the usual errors.
- Dropping the p one times p two term gives five per cent.
- Dropping the square root gives about one hundred and six per cent, which is impossible and should warn you.

5 Four weaknesses of a correlation framework

- Nonlinearity: one linear number hides tail dependence
- Sparse data: few joint defaults, large sampling error
- Instability: shifts with horizon, PD level, downturns
- Dimensionality: N(N − 1) ÷ 2 pairs to calibrate

IN THE LESSON

- A portfolio framework built on correlation has weaknesses, and the exam asks you to name them.
- First, nonlinearity: correlation is a single linear summary, so it cannot describe tail dependence, and two portfolios with the same correlation can differ widely in the chance that many names default together.
- Second, sparse data: defaults are rare, so default correlations come from very few joint events and carry large sampling error.
- Third, instability: correlations depend on the horizon and on the default probability level, and they rise in downturns, so calm period estimates understate stress losses.
- Fourth, dimensionality, which we count next.

6 Example 2: counting pairs and joint events

N obligors need N(N − 1) ÷ 2 correlations

10 borrowers	10 × 9 ÷ 2	**45 pairs**
250 borrowers (Marlow)	250 × 249 ÷ 2	**31,125 pairs**
1,000 borrowers	1000 × 999 ÷ 2	**499,500 pairs**
Joint defaults expected, 20 years	20 × 1% × 1%	**0.002**

IN THE LESSON

- With N obligors there are N times N minus one, over two, distinct pairs, each needing a correlation.
- Ten borrowers give forty five pairs.
- Marlow's two hundred and fifty corporate borrowers give thirty one thousand, one hundred and twenty five.
- A thousand borrowers need four hundred and ninety nine thousand, five hundred, and all must fit into one consistent matrix.
- Now sparse data.
- Two firms that each default one year in a hundred should, if independent, default together one year in ten thousand, so over twenty years we expect point zero zero two joint defaults, almost surely none.
- In an exam scenario, name the weakness: instability, sparse data, nonlinearity or dimensionality.

7 Correlation changes the shape, not the mean

- EL = sum of PD × LGD × exposure: no correlation in it
- Unexpected loss does depend on correlation
- Higher correlation: more joint defaults, fatter right tail
- Smaller diversification benefit, so more capital

IN THE LESSON

- Suppose Marlow has measured the correlations.
- What do they do?
- First, what they do not change.
- Expected loss is the sum, over the loans, of default probability times loss given default times exposure.
- No correlation appears, so it is the same whether defaults move together or not.
- Correlation changes the shape of the loss distribution.
- Higher default correlation makes joint defaults more likely, fattens the right tail where many loans fail together, and shrinks the diversification benefit.
- So unexpected loss, which equity must cover, grows with correlation while the mean stays put.

8 Unexpected loss of two loans

Sigma = $\sqrt{(\sigma_1^2 + \sigma_2^2 + 2 \times \rho \times \sigma_1 \times \sigma_2)}$
$\rho = 0$: $\sqrt{(\sigma_1^2 + \sigma_2^2)}$ $\rho = 1$: $\sigma_1 + \sigma_2$

Stand-alone loss SD	Each loan's own spread of loss
Correlation, rho	Their default correlation
Unexpected loss of the pair	Grows as rho grows

IN THE LESSON

- Take two loans with stand alone loss standard deviations sigma one and sigma two, and default correlation rho.
- The portfolio's loss standard deviation, the unexpected loss, is the square root of sigma one squared plus sigma two squared plus two times rho times sigma one times sigma two.
- Read it from the cross term.
- With rho equal to zero it vanishes and the risks combine like the sides of a right angle triangle.
- With rho equal to one the root becomes sigma one plus sigma two, the plain sum, with no diversification.
- Anything between sits below the sum, and the gap is the diversification benefit.

9 Example 3: the pair of loans

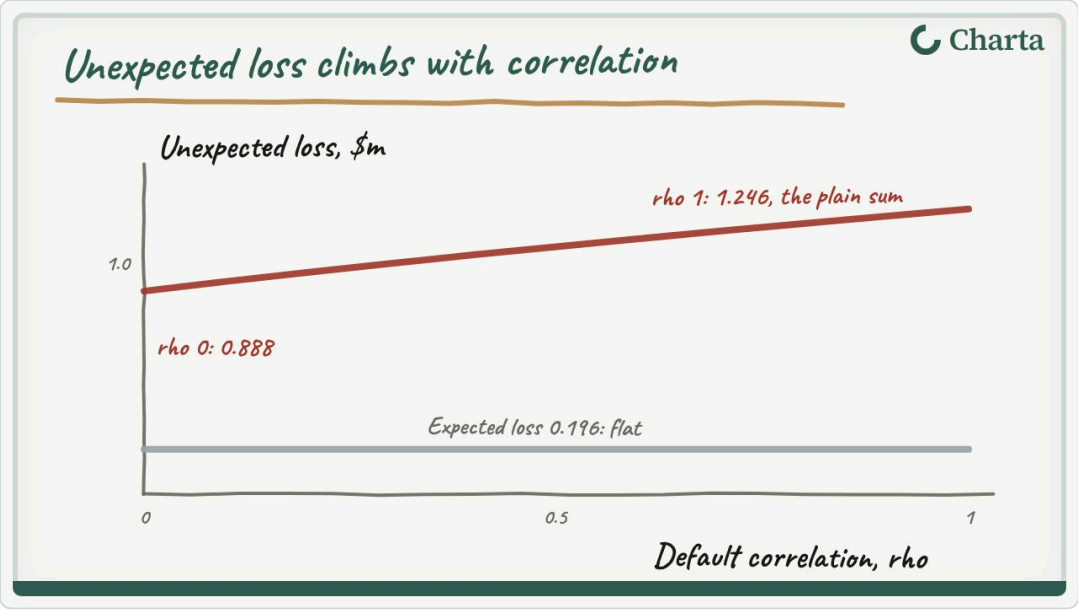
Alder $10m, LGD 50%, PD 2%. Birch $8m, LGD 40%, PD 3%. Amounts in $ millions

Expected loss, any ρ	0.100 + 0.096	**0.196**
Unexpected loss, ρ = 0	$\sqrt{(0.700^2 + 0.546^2)}$	**0.888**
Unexpected loss, ρ = 0.3	adds $2 \times 0.3 \times 0.700 \times 0.546$	**1.009**
Unexpected loss, ρ = 1	0.700 + 0.546	**1.246**
Diversification gain at 0.3	1.246 − 1.009	**0.237**

IN THE LESSON

- Take two Marlow loans.
- Alder Foods: ten million dollars, loss given default fifty per cent, default probability two per cent.
- Birch Steel: eight million, forty per cent, three per cent.
- Expected losses are one hundred thousand dollars and ninety six thousand, so one hundred and ninety six thousand in all, whatever the correlation.
- Stand alone standard deviations, exposure times loss given default times the square root of p times one minus p, are seven hundred thousand and about five hundred and forty six thousand.
- At correlation zero the pair's unexpected loss is about eight hundred and eighty eight thousand.
- At point three it is about one million and nine thousand.
- At one it is the plain sum, about one million, two hundred and forty six thousand.

10 Unexpected loss climbs with correlation



IN THE LESSON

- This chart plots the pair's unexpected loss as default correlation rises from zero to one.
- Expected loss stays flat at one hundred and ninety six thousand dollars.
- The unexpected loss starts at about eight hundred and eighty eight thousand with independence, rises at every point and bends over gently, and reaches one million, two hundred and forty six thousand at correlation one, where diversification has gone.
- A practical warning: the correlation we can estimate is usually the calm period figure, and calm period correlations understate stress, so the true correlation may sit further right on this curve.

11 Credit VaR and correlation

- Credit VaR = loss quantile at the confidence level − EL
- Normally rises with correlation, up to the plain sum
- Small, lumpy portfolios: step-function quantile, can fall
- Always quote the confidence level and the horizon

IN THE LESSON

- Now the figure risk managers quote, Credit V a R: the loss quantile at a chosen confidence level, less the expected loss.
- As the right tail fattens, the quantile rises, so Credit V a R normally rises with correlation, and at perfect correlation reaches the undiversified sum.
- But beware small, lumpy portfolios.
- With few loans the loss takes only a few values, so the quantile is a step function, and it can even fall when correlation rises.
- And a Credit V a R figure means nothing without its confidence level and horizon: ninety nine per cent and ninety nine point nine per cent over one year are very different numbers.

12 Example 4: when Credit VaR falls

Two loans, loss $5m each if default, PD 2%, 97.5% level. $ millions

Expected loss	2 × 2% × 5	**0.2**
Independent: chance of no loss	0.98 × 0.98	**96.04%**
Perfect correlation: no loss	both or neither	**98%**
97.5% quantile: ind., perfect	next step up; zero	**5, then 0**
Credit VaR: ind., perfect	5 − 0.2; 0 − 0.2	**4.8, then −0.2**

IN THE LESSON

- Here is the lumpy case.
- Marlow holds two loans, each losing five million dollars on default, each with a two per cent default probability.
- Expected loss is point two million.
- If defaults are independent, the chance of no loss is point nine eight squared, ninety six point zero four per cent.
- That is short of a ninety seven point five per cent level, so the quantile is the next step, five million, and Credit V a R is four point eight million.
- Now make defaults perfectly correlated.
- The chance of no loss is ninety eight per cent, which is enough, so the quantile falls to zero and Credit V a R is minus point two million.
- More correlation, less Credit V a R: do not trust quantiles from tiny portfolios.

13 The single factor model: the idea

- Each firm: shared economy plus its own luck
- Common factor M = state of the economy
- Default when the asset return falls below a threshold
- A bad M raises every firm's default chance at once

IN THE LESSON

- Thousands of pairwise correlations are hopeless to estimate, so practitioners build correlation from one shared ingredient, the economy.
- Each Marlow borrower's fortunes have two parts: the state of the economy, shared by everybody, and the firm's own luck.
- That is the single factor model.
- Each firm has a standardised asset return that rises when its assets do well, and it defaults when that return falls below a threshold.
- When the economy turns bad, every firm's return is pushed down together, so every firm is more likely to default at once.
- That is how correlation arises, without specifying a single pair.

14 The model: asset return and threshold

$R_i = \beta_i \times M + \sqrt{(1 - \beta_i^2)} \times e_i$
Default when $R_i < k_i$, $k_i = N^{-1}(PD_i)$

Beta, factor loading	Sensitivity to the economy
Common factor, M	Shared by every firm, std normal
Own shock, e	The firm's own luck, std normal
Default threshold, k	Return below this means default

IN THE LESSON

- Here is the model.
- Firm i's standardised asset return, R i, equals beta i times M, plus the square root of one minus beta i squared, times e i.
- M is the common factor, such as the state of the economy.
- The term e i is the firm's idiosyncratic shock.
- Both are standard normal, and the shocks are independent of M and of each other.
- Beta i is the firm's sensitivity to the economy, and the square root weight keeps the return standard normal.
- The firm defaults when R i falls below its threshold k i, the inverse standard normal of its default probability.
- For two per cent that is about minus two point zero five four.

15 Conditional default probability

CondPD = N((kᵢ − βᵢ × M) ÷ √(1 − βᵢ²))
Average over M gives back the plain PD

Default threshold, k	N⁻¹(PD) of the firm
Beta, factor loading	Loading on the economy
Common factor, M	Low M = bad economy
Conditional default probability	Chance of default, given M

IN THE LESSON

- Fix the economy at a value of M.
- Then the firms' own shocks are independent, so defaults are independent too, and each firm has a conditional default probability: N of the quantity, k i minus beta i times M, all divided by the square root of one minus beta i squared.
- Averaging over every value of M recovers the ordinary default probability.
- A bad outcome, a low M, lifts every firm's probability at once.
- And the model needs one beta per firm, instead of N times N minus one, over two, pairwise correlations.
- Common errors: flipping the sign of beta i times M, and omitting the square root in the denominator.

16 Example 5: default chance in a bad economy

Alder: β 0.4, PD 2%. Birch: β 0.7, PD 3%

Alder threshold	N⁻¹(2%)	**−2.054**
Alder, bad economy M = −2	(−2.054 + 0.8) ÷ 0.9165	**8.6%**
Birch, bad economy M = −2	(−1.881 + 1.4) ÷ 0.7141	**25.0%**
Alder, average economy M = 0	−2.054 ÷ 0.9165	**1.25%**
Wrong sign, Alder M = −2	(−2.054 − 0.8) ÷ 0.9165	**0.09%**

IN THE LESSON

- Run the model on two Marlow borrowers.
- Alder Foods has beta point four and default probability two per cent.
- Birch Steel has beta point seven and three per cent.
- Alder's threshold is minus two point zero five four.
- In a bad economy, M equal to minus two, the numerator is minus two point zero five four plus point eight, divided by the root of point eight four, which is point nine one six five.
- That gives minus one point three six eight, so Alder's conditional default probability is about eight point six per cent.
- Birch, with the larger beta, reaches twenty five per cent.
- In an average economy Alder falls to one point two five per cent.
- Flip the sign by mistake and Alder shows point zero nine per cent in the bad economy, which is absurd.

17 What beta means

- Beta = loading on the factor M = correlation with M
- Beta² explained by M; 1 − beta² is the firm's own share
- Negative beta: a counter-cyclical firm
- Standardised beta = Cov ÷ (SD of return × SD of factor)
- Cov 0.018, SDs 20% and 15%: beta 0.6, raw slope 0.8

IN THE LESSON

- Beta is the loading of the firm's standardised asset return on the common factor, and it equals the correlation between that return and the factor.
- Beta squared is the share of the return's variance explained by the factor; one minus beta squared is the firm's own idiosyncratic share.
- A beta can be negative: a counter cyclical firm, which does better when the economy weakens, has a negative correlation with the factor.
- Warning: a regression slope in raw units is not the standardised beta.
- Divide the covariance with the factor by the product of both standard deviations.
- A covariance of point zero one eight, with return and factor standard deviations of twenty and fifteen per cent, gives beta point six, while the raw slope is point eight.

18 Asset return correlation from two betas

$\rho(i, j) = \beta_i \times \beta_j$
Equal betas: $\rho = \beta^2$

Asset return correlation	Correlation of the asset returns
Beta of firm i	Loading of firm i on M
Beta of firm j	Loading of firm j on M

IN THE LESSON

- Why does this give a correlation between any two firms?
- Because the idiosyncratic shocks are independent, the only link between firm i and firm j is the common factor.
- So the asset return correlation is rho equals beta i times beta j.
- With equal betas, rho is beta squared.
- From just N betas we get every one of the N times N minus one, over two, pairwise correlations, and they fit together automatically.
- The typical errors are averaging the betas and multiplying their squares.
- The correlation is the product of the betas themselves, with their signs.

19 Example 6: Alder and Birch

Alder $\beta = 0.4$, Birch $\beta = 0.7$

Variance explained by M	0.4^2 and 0.7^2	**16%, 49%**
Asset return correlation	0.4×0.7	**0.28**
Alder with a β = −0.3 firm	$0.4 \times (-0.3)$	**−0.12**
Wrong: average, then square	$(0.4 + 0.7) \div 2$; then 0.55^2	**0.55; 0.3025**
Wrong: multiply the squares	0.16×0.49	**0.0784**

IN THE LESSON

- Back to Alder and Birch.
- Alder's factor explains sixteen per cent of the variance of its return, Birch's forty nine per cent.
- The asset return correlation is point four times point seven, point two eight.
- A counter cyclical firm with beta minus point three would have correlation with Alder of minus point one two.
- Now the wrong answers.
- Averaging the betas gives point five five, and squaring gives point three zero two five.
- Multiplying the squares gives point one six times point four nine, point zero seven eight four.
- Both differ from the right answer, point two eight.
- Use the product of the two betas.

20 Example 7: joint default probability

Alder β 0.4, PD 2%. Birch β 0.7, PD 3%. Asset correlation 0.28

Thresholds	$N^{-1}(2\%)$, $N^{-1}(3\%)$	**−2.054, −1.881**
Both default if M = −2	8.6% × 25.0%	**2.1%**
Both default if M = 0	1.25% × 0.42%	**0.005%**
Average over all M	bivariate normal routine	**0.212%**
If independent	2% × 3%	**0.060%**

IN THE LESSON

- Now the joint default probability of Alder and Birch.
- The asset return correlation is point two eight.
- The thresholds are minus two point zero five four and minus one point eight eight one.
- The joint default probability is the bivariate normal chance that both returns fall below their thresholds, or equivalently the average over M of the product of the two conditional default probabilities.
- In a bad economy, M equal to minus two, the product is eight point six per cent times twenty five per cent, about two point one per cent.
- In an average economy it is tiny.
- Averaged over every economy, weighted by likelihood and found with a bivariate normal routine, it is point two one two per cent, against point zero six per cent if independent: about three and a half times as likely.

21 Example 7 continued: default correlation

JDP = 0.2124%, p1 = 2%, p2 = 3%

Numerator	0.2124% − 0.06%	**0.1524%**
Denominator	√(0.0196 × 0.0291)	**0.02388**
Default correlation	0.001524 ÷ 0.02388	**6.4%**
Birch defaults, given Alder did	0.2124% ÷ 2%	**10.6%**
Alder defaults, given Birch did	0.2124% ÷ 3%	**7.1%**

IN THE LESSON

- Now the default correlation.
- The numerator is point two one two four per cent minus point zero six per cent, point one five two four per cent.
- The denominator is the square root of point zero one nine six times point zero two nine one, point zero two three eight eight.
- The default correlation is six point four per cent, far below the asset return correlation of twenty eight per cent.
- The joint probability also answers conditional questions.
- Given that Alder defaulted, Birch defaults with probability the joint probability divided by Alder's own, ten point six per cent, against three per cent unconditionally.
- Given that Birch defaulted, Alder's chance is seven point one per cent.

22 Asset correlation is not default correlation

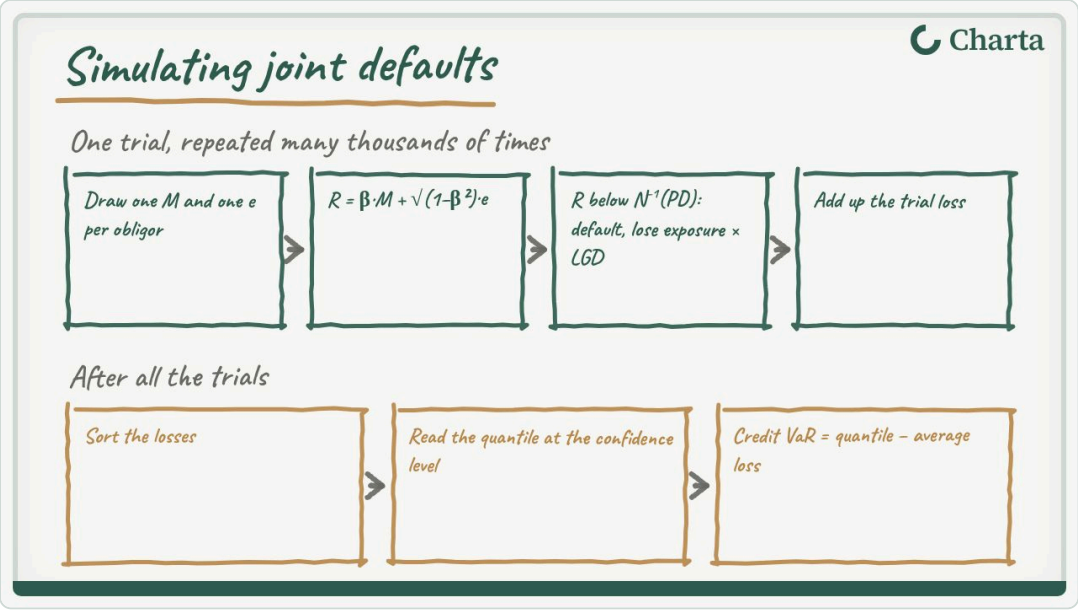
Three numbers for Alder and Birch, in the order we compute them

Step	Alder and Birch
1. Asset correlation, ρ = β × β	28%
2. Joint default probability	0.212%
3. Default correlation, Corr	6.4%
Mistake	Reporting 28% as Corr

IN THE LESSON

- Put the two correlations side by side, because confusing them is the commonest mistake here.
- The asset return correlation, rho, comes from the betas: twenty eight per cent.
- The joint default probability comes from the bivariate normal, or the average over the factor: point two one two per cent.
- The default correlation, Corr, comes from that joint probability by the formula we met at the start: only six point four per cent.
- Rho describes how smooth asset returns move together; Corr describes how yes or no default events move together.
- Reporting the asset correlation as the default correlation overstates dependence badly.
- Check which one the question asks for.

23 Simulating joint defaults



IN THE LESSON

- When no closed form exists, we simulate.
- In each trial, draw one common factor M, shared by every obligor, and an independent own shock e for each obligor.
- Form each obligor's asset return, beta times M plus the square root of one minus beta squared, times e.
- If it is below the threshold, the inverse standard normal of the default probability, the obligor defaults and loses its exposure times its loss given default.
- Add the losses for the trial's portfolio loss, and repeat many thousands of times.
- Then sort the losses, read off the quantile at the chosen confidence level, and subtract the average loss to get Credit V a R.

24 Example 8: one simulated trial

M = −2; shocks −0.5, −0.9, −1.0. Thresholds −2.054, −1.881, −2.170

Alder: 0.4M + 0.9165e	−0.800 − 0.458	**−1.26: safe**
Birch: 0.7M + 0.7141e	−1.400 − 0.643	**−2.04: default**
Cedar: 0.5M + 0.866e	−1.000 − 0.866	**−1.87: safe**
Trial loss	Birch: 8m × 40%	**$3.2m**
Same shocks, no factor	Birch R = −0.643	**$0 lost**

Cedar loan: PD 1.5%, beta 0.5, $12m, LGD 45%

IN THE LESSON

- One trial for a three loan Marlow portfolio: Alder, Birch and a new loan, Cedar, with default probability one point five per cent, beta point five, exposure twelve million dollars and loss given default forty five per cent.
- The draw is a bad economy, M equal to minus two, with own shocks of minus point five, minus point nine and minus one.
- Alder's return is minus one point two six, above its threshold, so it survives.
- Birch's is minus one point four minus point six four, minus two point oh four, below its threshold of minus one point eight eight one, so Birch defaults and loses eight million times forty per cent, three point two million.
- Cedar's is minus one point eight seven, above minus two point one seven, so it survives.
- Drop the common factor and the same shocks give no default: correlation vanishes and Credit V a R is understated.

25 From trials to Credit VaR

Illustrative simulation output. Amounts in $ millions

Trials run		**10,000**
Trials beyond the 99.9% quantile	10,000 × 0.1%	**10**
99.9% quantile	10th largest loss	**41.5**
Average loss over all trials		**5.2**
Credit VaR	41.5 − 5.2	**36.3**

With 1,000 trials the tail is one trial: very noisy

IN THE LESSON

- Reading the results.
- Marlow runs ten thousand trials and sorts the losses.
- At the ninety nine point nine per cent level, point one per cent of trials lie beyond the quantile, which is ten of ten thousand, so the quantile is about the tenth largest loss.
- Say it is forty one point five million dollars.
- The average loss over all trials is five point two million, so Credit V a R is thirty six point three million.
- The evidence is thin: ten trials decide the answer.
- With only a thousand trials the quantile would rest on the single largest loss and swing wildly from run to run.
- Tail estimates need plenty of trials.

26 Granularity

- Granular: many small loans. Concentrated: a few large
- Expected loss does not depend on granularity
- Credit VaR does: granular books diversify own luck away
- Concentrated books keep extra idiosyncratic risk
- Systematic risk stays however many names

IN THE LESSON

- Granularity is how evenly exposure is spread across names: a granular portfolio has many small exposures, a concentrated one a few large ones.
- Take two Marlow books of the same total size, one with a hundred small loans, the other with ten big ones.
- For given default probabilities, loss given default and correlation, expected loss is identical, because it is just a sum.
- Credit V a R is not.
- In the granular book each firm's own luck is a small piece, and the pieces average out, leaving only the systematic factor risk.
- In the concentrated book one idiosyncratic default can take a big bite, which raises Credit V a R.
- But the factor risk never disappears.

27 Measuring granularity: the Herfindahl index

HHI = Σ (share of each name)²
Effective number of names = 1 ÷ HHI

Share of exposure	Name's share of total exposure
Herfindahl index, HHI	Low = granular, high = lumpy
Effective number of names	Equal-sized loans it behaves like

IN THE LESSON

- We measure granularity with the Herfindahl index, H H I: square each name's share of total exposure and add up the squares.
- For N equal exposures each share is one over N, so the index is one over N.
- One divided by the index is therefore the effective number of equal sized exposures.
- A low index means a granular portfolio, a high index a concentrated one, and the effective number translates an uneven book into the number of equal loans it behaves like.
- A single name book has an index of one and an effective number of one.

28 Example 9: three loan books

Each loan book totals $1,000m

100 loans of $10m	100×0.01^2	**HHI 0.01**
10 loans of $100m	10×0.1^2	**HHI 0.10**
$400m plus six of $100m	$0.4^2 + 6 \times 0.1^2$	**HHI 0.22**
Effective numbers	1÷0.01, 1÷0.10, 1÷0.22	**100, 10, 4.5**

IN THE LESSON

- Each Marlow loan book totals one thousand million dollars.
- Book X has a hundred loans of ten million.
- Each share is one hundredth, squared point zero zero zero one, and a hundred of them add to point zero one: effective number one hundred.
- Book Y has ten loans of a hundred million.
- Each share squared is point zero one, ten of them make point one, effective number ten.
- Book Z is lopsided: one loan of four hundred million and six of a hundred million.
- The squared shares give point one six plus six times point zero one, point two two, and one over point two two is four point five.
- Though it holds seven loans, Z behaves like only four and a half equal ones.

29 The floor under Credit VaR

Loss rate = $N\left(\left(N^{-1}(PD) + \sqrt{\rho} \times N^{-1}(q)\right) \div \sqrt{(1-\rho)}\right)$

Credit VaR = Loss rate × LGD × exposure − EL

Default probability	Each name's default chance
Correlation, rho	Asset return correlation
Confidence level, q	For example 99%
Limiting loss rate	Share of exposure lost at level q

IN THE LESSON

- As a portfolio becomes very granular, idiosyncratic risk diversifies away, and the loss rate per unit of exposure, before loss given default, at confidence q, is N of the quantity inverse N of the default probability plus root rho times inverse N of q, all divided by root one minus rho.
- This is the Vasicek worst case default rate.
- Multiply by loss given default and exposure for the loss at the quantile, and subtract expected loss for Credit V a R.
- The key point: this limit is not zero.
- Adding names never removes systematic risk, so Credit V a R has a floor above zero, and a concentrated book sits above it.

30 Example 10: the granular floor

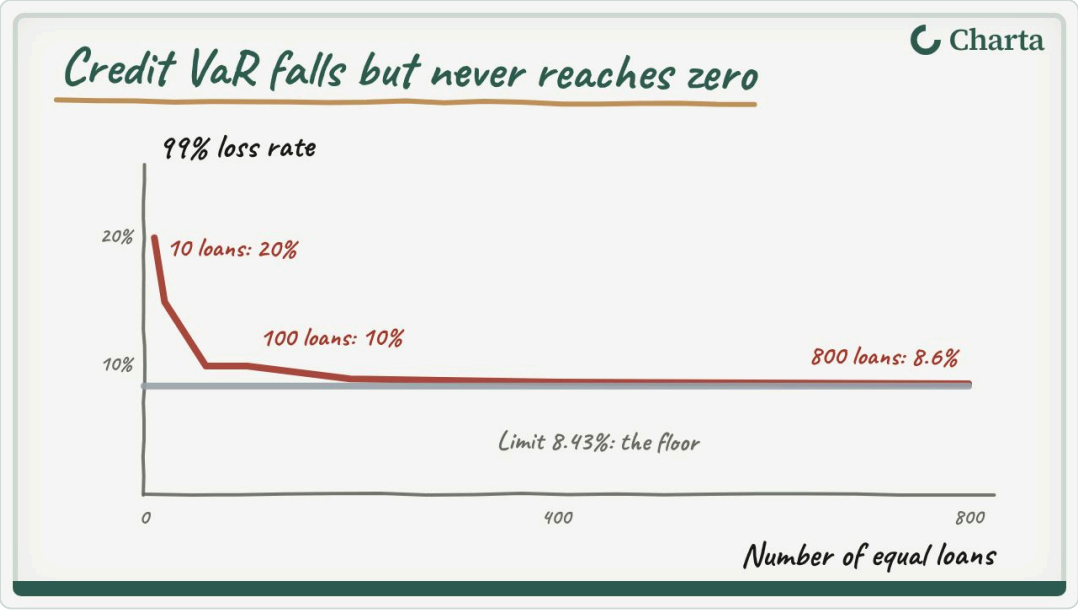
Exposure $1,000m, PD 1.5%, ρ = 0.15, LGD 45%, q = 99%

Numerator	$-2.1701 + 0.3873 \times 2.3263$	**−1.2691**
Divide by √0.85	−1.2691 ÷ 0.9220	**−1.3765**
Loss rate N(−1.3765)		**8.43%**
Loss at the quantile	8.43% × 45% × 1000	**$37.9m**
Credit VaR floor	37.9 − 6.75 (EL)	**$31.2m**

IN THE LESSON

- Compute the floor for Marlow's book of one thousand million dollars: default probability one point five per cent, asset correlation fifteen per cent, loss given default forty five per cent, confidence ninety nine per cent.
- Inverse N of point zero one five is minus two point one seven zero one.
- Inverse N of point nine nine is two point three two six three, and root point one five is point three eight seven three, so that term adds point nine zero one.
- The numerator is minus one point two six nine one.
- Dividing by root point eight five, point nine two two, gives minus one point three seven six five, and the normal function gives eight point four three per cent.
- The loss at the quantile is eight point four three per cent times forty five per cent times a thousand million, thirty seven point nine million.
- Less expected loss of six point seven five million, the floor is thirty one point two million.

31 Credit VaR falls but never reaches zero



IN THE LESSON

- This chart shows the ninety nine per cent loss rate, for the same default probability and correlation, as the number of equal loans grows.
- With ten loans it is twenty per cent.
- With a hundred it is ten per cent.
- With eight hundred it is about eight point six per cent.
- The grey line is the limit we just computed, eight point four three per cent.
- The curve falls, flattens as it nears that floor, and stays above it.
- The steps at small numbers are the lumpy quantiles we met earlier.
- Spreading exposure more evenly lowers Credit V a R, but cannot get below the systematic floor.

32 Recap: the main ideas

- Default correlation: small, from JDP
- Weak points: nonlinear, sparse, unstable, many pairs
- Correlation: EL unchanged, UL and Credit VaR up
- Beta: rho = beta i × beta j; JDP from the factor model
- Simulate for Credit VaR; granular books still have a floor

IN THE LESSON

- Let us pull it together.
- Default correlation is the correlation of default indicators, computed from the joint default probability, and it is small.
- A framework built on correlation suffers from nonlinearity, sparse data, instability and dimensionality.
- Correlation leaves expected loss unchanged and raises unexpected loss and normally Credit V a R, though lumpy portfolios can behave oddly.
- In the single factor model each firm has a beta, the asset correlation is the product of two betas, and the joint default probability and default correlation follow.
- Simulation reads Credit V a R from a quantile.
- Concentration raises Credit V a R, and systematic risk is a floor that no number of names removes.

33 Try three quick questions

1. p = 5% and 2%, JDP = 0.3%. Default correlation?
   Answer: (0.3% − 0.1%) ÷ 3.05% ≈ 6.6%.
2. Betas 0.6 and 0.5. Asset correlation? Default correlation larger?
   Answer: 0.6 × 0.5 = 0.30. Default correlation is smaller.
3. SDs $6m and $8m, ρ = 0.5. Portfolio SD?
   Answer: √(36 + 64 + 48) ≈ $12.2m, below $14m.

IN THE LESSON

- Now try three quick questions.
- Pause the video, think, then listen to the answers.
- Question one.
- Two firms have default probabilities of five per cent and two per cent, and a joint default probability of point three per cent.
- What is the default correlation?
- Independence would give point one per cent, so the numerator is point two per cent.
- The denominator is the square root of point zero four seven five times point zero one nine six, about three point oh five per cent.
- The answer is about six point six per cent.
- Question two.
- In a single factor model two firms have betas of point six and point five.
- What is the asset return correlation, and is the default correlation larger or smaller?
- Point three, and smaller.
- Question three.
- Two loans have loss standard deviations of six million and eight million dollars, with correlation point five.
- What is the portfolio standard deviation?
- The root of one hundred and forty eight, about twelve point two million, below the plain sum of fourteen million.