CHARTA LEARNING · LESSON NOTES

Bottom-Up Correlation Models

FRM Part II

1 A story: two names, one question

- Kestrel holds Ashford Motors and Brook Shipping bonds
- One-year default chance: Ashford 5%, Brook 3%
- Committee: what is the chance both default?
- Unrelated: 5% × 3% = 0.15%. Shared driver: up to 3%
- The answer depends on how the names move together

IN THE LESSON

- Meet Kestrel Capital, a credit fund.
- It holds bonds issued by Ashford Motors, a car maker, and Brook Shipping, a freight firm.
- Its analysts have studied each company on its own, and they estimate a one year default chance of five per cent for Ashford and three per cent for Brook.
- The risk committee asks a simple question: what is the chance that both default within the year?
- If the two firms had nothing to do with each other, the answer would be five per cent times three per cent, just point one five per cent.
- But both depend on world trade.
- If they share a driver, the chance could be as high as three per cent, twenty times larger.
- The answer depends on how the names move together, and that is what this lesson models.

2 Marginal and joint distributions

- Marginal: one variable's own distribution
- Cumulative probability: chance of being at or below a value
- Joint: the chance of events for both together
- Dependence: how one outcome shifts the other's chances
- The marginals alone do not fix the joint

IN THE LESSON

- First, some words.
- A marginal distribution is the distribution of one variable on its own, such as Ashford's default time, ignoring everything else.
- A cumulative probability is the chance that the variable is at or below a given value.
- For example, the chance that Ashford has defaulted by the end of one year is its cumulative default probability, five per cent.
- A joint distribution gives the chance of events for both variables together, such as both firms defaulting by the year end.
- Dependence is how the outcome of one variable shifts the chances for the other.
- The key point is that the two marginals alone do not fix the joint.
- Many different joint distributions share exactly the same marginals, and they differ only in their dependence.

3 Why correlation alone is not enough

- Correlation fixes the joint if jointly normal, not in general
- X normal, Y = X^2: correlation 0, yet X fixes Y
- Default times and fat-tailed returns are not normal
- We need a tool that carries all of the dependence

IN THE LESSON

- You might think correlation settles the matter.
- It does not.
- In general, correlation does not fix the joint distribution.
- Jointly normal variables are the standard case where it does.
- Here is a simple example.
- Let X be standard normal, and let Y be X squared.
- The correlation between them is zero, because large positive and large negative values of X give the same Y, and the effects cancel.
- Yet Y is completely determined by X.
- Zero correlation, total dependence.
- In finance, the variables we care about are rarely normal.
- Default times are not normal, and returns often have fat tails, meaning extreme outcomes are more common than a normal curve allows.
- So we need a tool that carries all of the dependence, not just one number.
- That tool is a copula.

4 Sklar's theorem

Joint = C(u, v)
u = P(X ≤ x), v = P(Y ≤ y)

Marginal probability u	Marginal cumulative prob. of X
Marginal probability v	Marginal cumulative prob. of Y
Copula, C	Carries all the dependence
Joint probability	Both at or below together

IN THE LESSON

- Here is the formal idea.
- A copula is a function that joins separate marginal distributions into one joint distribution, while carrying all of the dependence between the variables.
- Take two variables, X and Y.
- Call the marginal cumulative probability of X the letter u: the chance that X is at or below some value, little x.
- Call the one for Y the letter v.
- Each lies between zero and one.
- Let C be the copula.
- Sklar's theorem says that the joint cumulative probability, the chance that X is at or below little x and Y is at or below little y at the same time, can always be written as C of u and v.
- The marginals enter only through u and v.
- Everything about dependence lives in C.

5 What a copula lets a risk manager do

- Step 1: calibrate each marginal on its own
- For example a default time, or a fat-tailed return
- Step 2: choose the dependence, the copula, separately
- Bottom-up: single default probabilities become joint ones
- These drive CDO tranches and portfolio loss

IN THE LESSON

- This split is what makes copulas so useful.
- A risk manager works in two steps.
- Step one: calibrate each asset's own distribution, using whatever fits best.
- For a bond it might be a default time, and for a share it might be a fat tailed return.
- Step two: choose the dependence structure separately, by picking a copula and its parameter.
- Neither step disturbs the other.
- In credit, this is the heart of the bottom up approach, which builds portfolio risk upwards from the individual names.
- Copulas turn individual default probabilities into joint default probabilities, and those drive the tranches of a C D O and the portfolio loss.
- Recall that a C D O slices a pool of credit risk into layers called tranches.

6 Three anchor cases

Every copula lies between the two bounds

Dependence	Joint probability C(u, v)
Independence	u × v
Perfect positive: upper bound	min(u, v)
Most negative: lower bound	max(u + v − 1, 0)
Any copula	Between the two bounds

IN THE LESSON

- Every copula lies between two bounds, and three anchor cases show why.
- With independence, the joint probability is simply u times v.
- With perfect positive dependence, the two variables move in lockstep, so both are at or below their values whenever the less likely event happens, and the joint probability is the smaller of u and v.
- That is the upper bound.
- The lower bound, reached with the most negative dependence possible, is the larger of two numbers: u plus v minus one, and zero.
- Any copula at all gives a joint probability between the lower and the upper bound.
- A joint probability outside these bounds is a sign of an error in the calculation.

7 Example 1: bounds for two bond returns

Monthly bond returns at or below a chosen level: u = 0.8, v = 0.7

Independence	0.8 × 0.7	**0.56**
Upper bound	min(0.8, 0.7)	**0.70**
Lower bound	max(0.8 + 0.7 − 1, 0)	**0.50**
Range for any copula	lower to upper bound	**0.50 to 0.70**

IN THE LESSON

- Let us try the anchors on Kestrel's bonds.
- Over a month, the chance that Ashford's bond return is at or below a chosen level is point eight, so u is point eight.
- For Brook the chance is point seven, so v is point seven.
- With independence, the joint chance is point eight times point seven, point five six.
- With perfect positive dependence, it is the smaller of the two, point seven.
- The lower bound is point eight plus point seven minus one, which is point five, and that is larger than zero.
- So whatever copula Kestrel picks, the joint probability lies between point five and point seven.
- The copula decides where in that range the answer falls.

8 Example 2: bounds for two defaults

One-year default probabilities: Q_A = 5% (Ashford), Q_B = 3% (Brook)

Independence	5% × 3%	**0.15%**
Upper bound	min(5%, 3%)	**3%**
Lower bound	max(5% + 3% − 1, 0)	**0**
Upper ÷ independence	3% ÷ 0.15%	**20 times**

IN THE LESSON

- Now the question from the risk committee.
- Ashford's one year default probability, written Q A, is five per cent, and Brook's, Q B, is three per cent.
- With independence, both default with probability five per cent times three per cent, point one five per cent.
- With perfect positive dependence, both default whenever Brook does, so the joint probability is the smaller figure, three per cent.
- The lower bound is five per cent plus three per cent minus one, a negative number, so the bound is zero.
- The upper bound, three per cent, divided by the independent answer, point one five per cent, is twenty, so the upper bound is twenty times the independent answer.
- For rare events like defaults, the choice of copula moves the answer across a huge range.

9 Mistakes to avoid with copulas

- Correlation alone does not pin down the joint
- Marginals and copula are separate choices
- Check every joint probability lies within the bounds
- The copula is chosen, not observed: model risk

IN THE LESSON

- Before we pick a copula, four traps.
- First, correlation alone does not pin down the joint distribution, unless the variables are jointly normal.
- Second, the marginals and the copula are separate choices.
- Getting each default probability right says nothing about whether the dependence is right.
- Third, check that every joint probability lies within the bounds.
- Fourth, the copula is chosen by the modeller, not observed in the data, especially for joint defaults, which are rare.
- That creates model risk: two copulas with the same marginals can give very different joint default probabilities, as the twenty fold range we just found shows.
- Now let us meet the copula used most in credit.

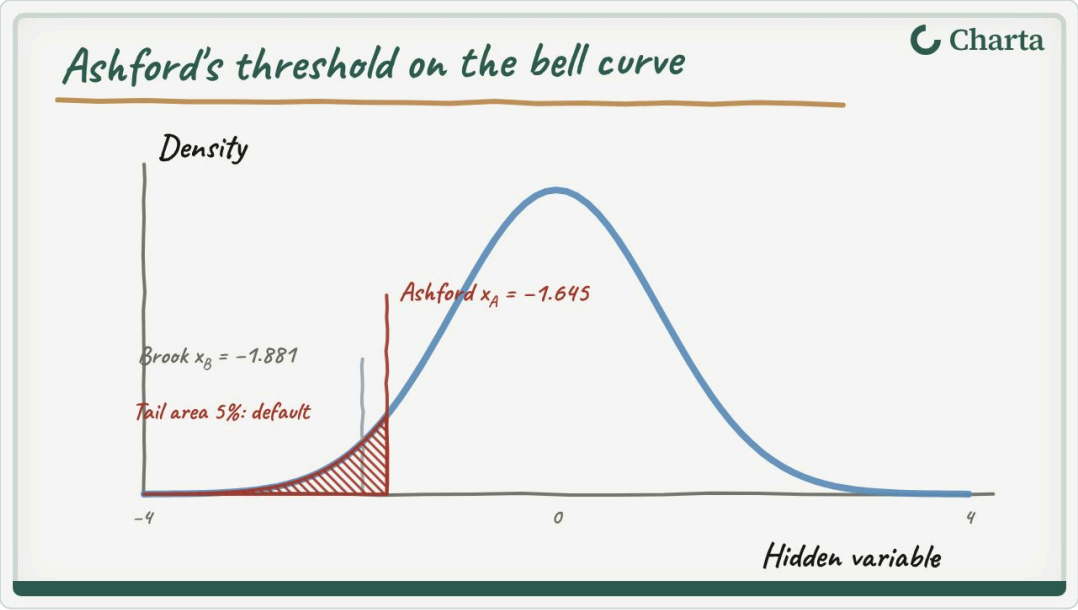
10 The Gaussian copula: the idea

- Give each name a hidden standard normal variable
- It defaults by time t if the variable falls below a threshold
- Threshold set so the chance of falling below equals Q
- Link the hidden variables: bivariate normal, correlation ρ
- ρ is the correlation of the hidden normals, not of defaults

IN THE LESSON

- The Gaussian copula links the marginals with a normal distribution.
- Here is the idea.
- Give each name a hidden, or latent, standard normal variable.
- Think of it as the firm's financial health.
- The name defaults by time t when its variable falls below a threshold.
- The threshold is chosen so that the chance of falling below it equals the name's cumulative default probability, Q.
- That keeps each name's own default probability exactly right.
- Then the hidden variables of the two names are linked by a bivariate normal distribution, meaning a joint normal distribution of two variables, with correlation rho.
- Keep one fact in mind from the start: rho is the correlation of the hidden normals, not the correlation of the defaults themselves.

11 Ashford's threshold on the bell curve



IN THE LESSON

- This chart shows the standard normal bell curve for Ashford's hidden variable.
- Along the bottom is the value of the variable, from minus four, through zero, to four.
- Up the side is the density, the height of the curve.
- The vertical line marks Ashford's threshold, minus one point six four five.
- The hatched tail to its left holds five per cent of the area under the curve, exactly Ashford's default probability.
- Any draw that lands in that tail is a default.
- For Brook, with three per cent, the threshold sits further left, at minus one point eight eight one, because a smaller tail needs a lower cut off.
- Lower default probability, lower threshold.

12 From default probability to joint default

x = N^{−1}(Q)
P_{AB} = M(x_A, x_B, ρ)

Default probability, Q	Cumulative default prob. by t
Normal threshold, x	Default if hidden variable below
Copula correlation, ρ	Correlation of hidden normals
Bivariate normal, M	Bivariate normal cumulative

IN THE LESSON

- Here are the two steps as formulas.
- First, map each name's cumulative default probability by time t, Q, to a normal threshold.
- The threshold x equals the inverse of the standard normal distribution function, written N inverse, at Q.
- Do this for each name, so Q A gives x A, and Q B gives x B.
- Second, the chance that both A and B default, written P A B, is the bivariate normal cumulative probability at those two thresholds, with correlation rho.
- It is written capital M of x A, x B and rho, where capital M is simply the name of that function.
- In words, it is the chance that both hidden variables fall below their thresholds at the same time.

13 Example 3: Ashford and Brook default together

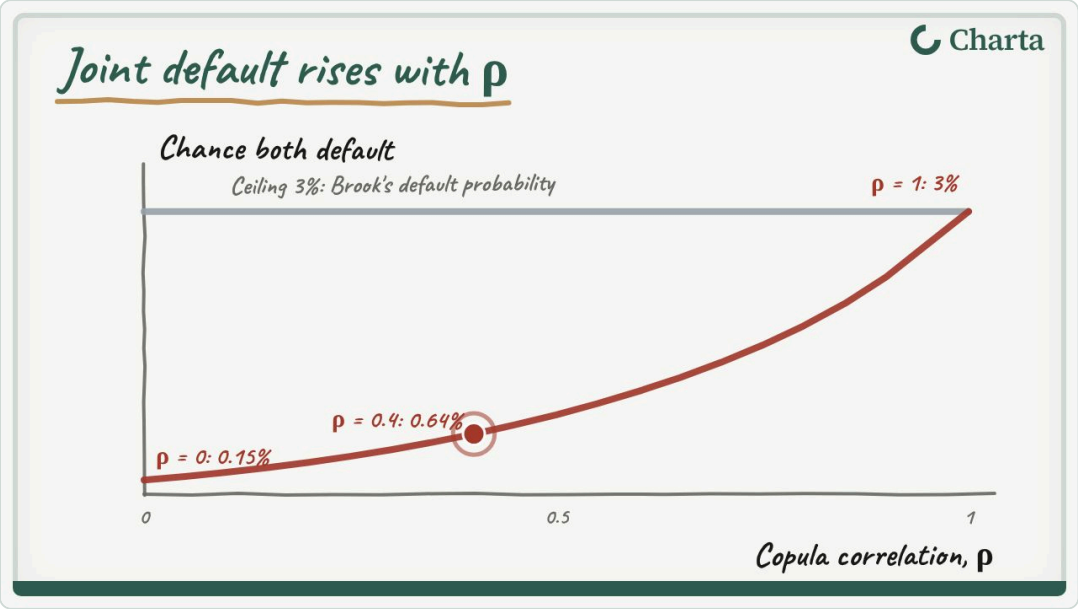
Q_A = 5%, Q_B = 3% by one year. Copula correlation ρ = 0.4

Ashford threshold x_A	N^{−1}(0.05)	**−1.6449**
Brook threshold x_B	N^{−1}(0.03)	**−1.8808**
Joint default, ρ = 0	5% × 3%	**0.15%**
Joint default, ρ = 1	min(5%, 3%)	**3%**
Joint default, ρ = 0.4	M(−1.6449, −1.8808, 0.4)	**0.639%**

IN THE LESSON

- Now compute for Kestrel.
- Q A is five per cent and Q B is three per cent, both by one year, and the fund uses a copula correlation, rho, of point four.
- Ashford's threshold is inverse N of point zero five, minus one point six four four nine.
- Brook's is inverse N of point zero three, minus one point eight eight zero eight.
- Now the joint default probability.
- Setting rho to zero gives five per cent times three per cent, point one five per cent, the independent case.
- Setting rho to one gives the smaller of the two, three per cent.
- At rho of point four, the bivariate normal at the two thresholds gives about point six three nine per cent: more than four times the independent figure, and still inside the bounds.

14 Joint default rises with ρ



IN THE LESSON

- This chart plots the joint default probability of Ashford and Brook as rho moves from zero to one.
- Along the bottom is rho, from zero, through point five, to one.
- Up the side is the chance that both default.
- The grey line at three per cent is the ceiling, Brook's own default probability, the upper bound.
- The red curve starts at point one five per cent when rho is zero, passes point six four per cent at rho of point four, and rises more and more steeply over most of the range, reaching three per cent at rho of one.
- Notice how slowly it climbs at first.
- Even a fairly high latent correlation leaves the joint default probability far below its ceiling.
- Most of the rise comes at high values of rho.

15 Default correlation

Corr = (P_{AB} − P_A × P_B) ÷ (s_A × s_B)
s_A = √(P_A − P_A^2)

Joint default probability	Chance both default
Default probabilities	Each name's default chance, Q
Indicator standard deviation	SD of a 0 or 1 indicator
Default correlation	Correlation of the defaults

IN THE LESSON

- Rho is not what people usually mean by default correlation.
- Default correlation is the correlation between two default indicators.
- A default indicator is a variable that equals one if the name defaults by time t, and zero if it does not.
- Its mean is the default probability, P, the same number as Q.
- The covariance of two indicators is the joint default probability minus the product of the two default probabilities: P A B minus P A times P B.
- Divide that by the product of the two standard deviations, s A times s B.
- Each indicator's standard deviation is the square root of P A minus P A squared, which is the same as P A times one minus P A, and s B likewise.

16 Example 4: Ashford and Brook's default correlation

P_A = 5%, P_B = 3%, P_AB = 0.639% (from ρ = 0.4)

Covariance	0.639% − 0.15%	**0.489%**
s_A	√(0.05 − 0.0025)	**0.2179**
s_B	√(0.03 − 0.0009)	**0.1706**
s_A × s_B	0.2179 × 0.1706	**0.03718**
Default correlation	0.00489 ÷ 0.03718	**13.2%**

IN THE LESSON

- Now compute it for Kestrel's pair, with the joint default probability of point six three nine per cent from rho of point four.
- The covariance is point six three nine per cent minus point one five per cent, which is point four eight nine per cent.
- Ashford's standard deviation is the square root of point zero five minus point zero zero two five, about point two one seven nine.
- Brook's is the square root of point zero three minus point zero zero zero nine, about point one seven zero six.
- Their product is point zero three seven one eight.
- Dividing point zero zero four eight nine by point zero three seven one eight gives a default correlation of about thirteen point two per cent.
- A latent correlation of forty per cent gives a default correlation of only about thirteen.

17 Trap: ρ is not the default correlation

- ρ: correlation of the hidden normals (here 0.4)
- Default correlation: of the 0 or 1 indicators (here 13.2%)
- Small default probabilities: default correlation much lower
- Never use ρ where default correlation is needed
- Given ρ: go through M to P_AB first

IN THE LESSON

- This is the most common mistake in the topic.
- Rho is the correlation of the hidden normal variables, here point four.
- Default correlation is the correlation of the zero or one default indicators, here thirteen point two per cent.
- They are different numbers.
- For small default probabilities, the default correlation is much lower than rho, because defaults are rare events in the far tail, and even strongly linked hidden variables seldom land in both tails together.
- So never put rho into a formula that needs default correlation, and never quote rho as if it were the correlation of defaults.
- If a question gives you rho, go through the bivariate normal to the joint default probability first.

18 Example 5: how one default changes the other

Conditional probability of B given A = P_AB ÷ P_A

Brook alone	Q_B	**3.0%**
Brook, given Ashford defaults	0.639% ÷ 5%	**12.8%**
Ratio to Brook alone	12.8% ÷ 3.0%	**4.3×**
Given Ashford, if ρ = 0	0.15% ÷ 5%	**3.0%**

IN THE LESSON

- The joint probability also tells Kestrel how bad news spreads.
- The conditional probability of B given A, the chance Brook defaults given that Ashford has defaulted, is P A B divided by P A.
- On its own, Brook's default chance is three per cent.
- Given that Ashford has defaulted, it is point six three nine per cent divided by five per cent, about twelve point eight per cent.
- That is about four point three times Brook's stand alone figure.
- Compare the independent case, with rho at zero.
- Then the conditional probability is point one five per cent divided by five per cent, just three per cent again: Ashford's default tells us nothing about Brook.
- This is why correlated defaults cluster.

19 From two names to a whole portfolio

- Many names need a full correlation matrix
- 100 names: 100 × 99 ÷ 2 = 4,950 pairs
- Simulate: Cholesky of the matrix × independent normals
- One-factor version: one common factor, one ρ

IN THE LESSON

- Kestrel's real portfolio holds a hundred names, not two.
- To model them all at once, the Gaussian copula needs a correlation structure across all of them: a correlation matrix, with one entry for every pair.
- With a hundred names, that is a hundred times ninety nine divided by two, four thousand nine hundred and fifty pairs, each to be estimated.
- The usual simulation draws correlated standard normals by taking a Cholesky decomposition of the correlation matrix, a kind of square root of the matrix, and applying it to independent standard normals.
- That works, but it needs every pair.
- The one factor version is far simpler: every name shares one common factor, and one parameter, rho, sets the correlation of every pair.

20 The one-factor Gaussian copula

$X_i = \sqrt{\rho} \times M + \sqrt{(1 - \rho)} \times Z_i$

Hidden draw, X_i	Name i's hidden draw
Common factor, M	Shared by all: the economy
Own shock, Z_i	Name i's own shock
Copula correlation, ρ	Between 0 and 1

IN THE LESSON

- Here is the one factor model.
- For name i, the draw is written X sub i.
- X sub i equals the square root of rho times M, plus the square root of one minus rho, times Z sub i.
- Careful: this capital M is a new letter.
- It is the common systematic factor shared by every name, so think of it as the state of the economy.
- It is not the bivariate normal function from before.
- Z sub i is the independent idiosyncratic term, the shock that belongs to name i alone.
- M and every Z are independent standard normal draws, and rho is the copula correlation, between zero and one.
- A low M, a bad economy, pulls every name's draw down at once, which is exactly how defaults come to cluster.

21 Why the square root of ρ

ρ = 0.4: √ρ = 0.6325 and √(1 − ρ) = 0.7746

Variance of X_i	0.6325^2 + 0.7746^2	**0.4 + 0.6 = 1**
Correlation of X_i and X_j	0.6325 × 0.6325	**0.4 = ρ**
Mean of X_i	both parts have mean 0	**0**

So each X_i is standard normal, and every pair has correlation ρ

IN THE LESSON

- Why the square roots?
- Take Kestrel's rho of point four.
- The square root of rho is point six three two five, and the square root of one minus rho is point seven seven four six.
- The variance of X sub i is the square of each weight, added up: point four plus point six, which is exactly one, because M and Z sub i are independent and each has a variance of one.
- Two names share only M, so the correlation of X sub i and X sub j is point six three two five times point six three two five, which is point four, rho itself.
- And the mean of X sub i is zero, since both parts have a mean of zero.
- So each X sub i is again standard normal, and every pair has correlation rho.

22 From a draw to a default time

U_i = N(X_i)
Q(t) = 1 − e^{−λt}
Time = −ln(1 − U_i) ÷ λ

Uniform number, U_i	Between 0 and 1
Default probability by t	Constant hazard rate case
Hazard rate, λ	Default rate a year
Default time	Low draw: early default

IN THE LESSON

- Next, each draw becomes a default time, in two moves.
- First, X sub i becomes a uniform number, U sub i, through the normal distribution function N.
- U sub i equals N of X sub i, and it lies between zero and one.
- Second, that uniform number is turned into a default time, by inverting the name's cumulative default distribution.
- Assume a constant hazard rate, lambda, the rate of default per year for a firm that has survived so far.
- Then the cumulative default probability by time t is one minus e to the minus lambda t.
- Setting that equal to U sub i and solving gives the default time: minus the natural log of one minus U sub i, divided by lambda.
- A low normal draw gives a small U sub i, and so an early default.

23 Example 6: hazard rates for Ashford and Brook

Match the one-year default probabilities: Q_A = 5%, Q_B = 3%

Ashford λ_A	−ln(1 − 0.05)	**5.13% a year**
Brook λ_B	−ln(1 − 0.03)	**3.05% a year**
Check: Ashford by 1 year	1 − e^{−0.0513}	**5.0%**

IN THE LESSON

- Kestrel calibrates each name's hazard rate so that the marginal matches the default probabilities we used before.
- Setting one minus e to the minus lambda equal to five per cent for Ashford gives lambda A as minus the natural log of one minus point zero five, about five point one three per cent a year.
- For Brook, minus the natural log of one minus point zero three gives about three point zero five per cent.
- Check Ashford: one minus e to the minus point zero five one three gives back five per cent by one year.
- This is the copula's two step split at work.
- The hazard rates fix each marginal, and rho, the same point four as before, fixes the dependence.

24 Example 7: one trial for Ashford

Trial draws: $M = -1.5$ and $Z_A = -1.2$, with $\rho = 0.4$

Common part	$0.6325 \times (-1.5)$	**−0.9487**
Own part	$0.7746 \times (-1.2)$	**−0.9295**
X_A	$-0.9487 - 0.9295$	**−1.8782**
$U_A = N(X_A)$	$N(-1.8782)$	**0.0302**
Default time	$-\ln(1 - 0.0302) \div 0.0513$	**0.60 years**

IN THE LESSON

- Now one simulated trial.
- The computer draws the common factor M as minus one point five, a bad year for the economy, and Ashford's own shock, Z A, as minus one point two.
- The common part is point six three two five times minus one point five, minus point nine four eight seven.
- The own part is point seven seven four six times minus one point two, minus point nine two nine five.
- Adding them, X A is minus one point eight seven eight two.
- The uniform number U A is N of that, about point zero three zero two.
- The default time is minus the natural log of one minus point zero three zero two, divided by point zero five one three, about point six of a year, roughly seven months.

25 Example 8: the same trial for Brook

Same $M = -1.5$. Brook's own shock $Z_B = +0.5$

Common part	same as Ashford	**−0.9487**
Own part	0.7746×0.5	**+0.3873**
X_B	$-0.9487 + 0.3873$	**−0.5614**
$U_B = N(X_B)$	$N(-0.5614)$	**0.2873**
Default time	$-\ln(1 - 0.2873) \div 0.0305$	**11.1 years**

IN THE LESSON

- Brook lives in the same economy, so in the same trial it gets the same common factor, minus one point five, and its common part is again minus point nine four eight seven.
- But its own shock is plus point five, good news for Brook alone.
- The own part is point seven seven four six times point five, plus point three eight seven three.
- So X B is minus point five six one four.
- U B is N of that, about point two eight seven three.
- The default time is minus the natural log of one minus point two eight seven three, divided by point zero three zero five, about eleven point one years.
- The bad economy hurt both names, but Brook's own luck carried it well past the year.

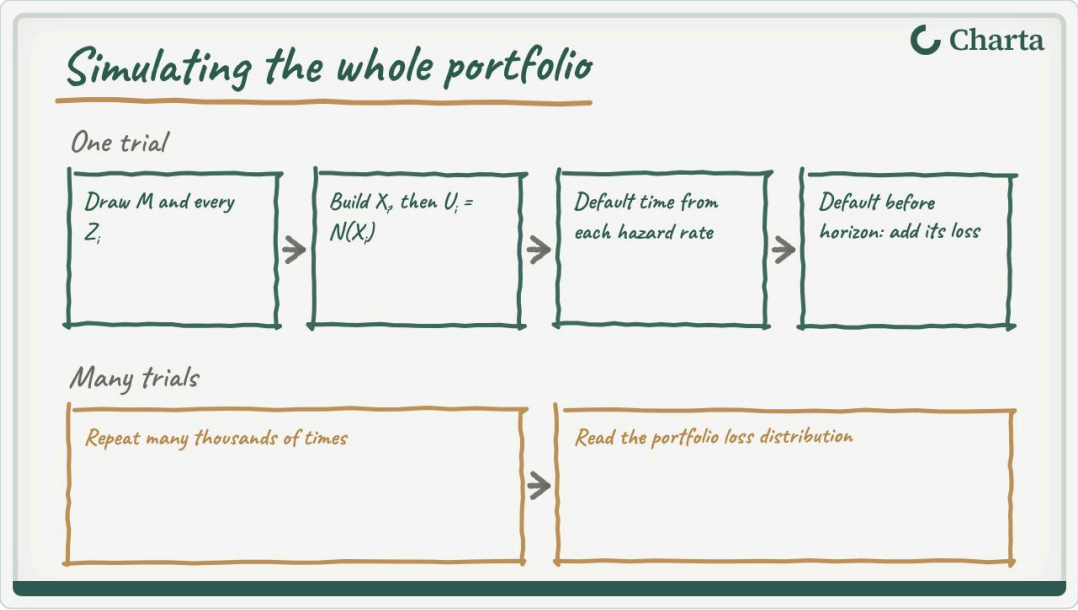
26 The threshold and the default time agree

- Default within 1 year: U < Q, the same as X < N^{−1}(Q)
- Ashford: −1.878 < −1.645, default at 0.60 years
- Brook: −0.561 > −1.881, survives to 11.1 years
- Low draw = early default
- The default time says when, not just whether

IN THE LESSON

- Notice how this links back to the thresholds.
- A name's default time falls within one year exactly when its uniform number is below its one year default probability, Q, which is the same as its draw X falling below the threshold, inverse N of Q.
- Ashford's draw, minus one point eight seven eight, is below its threshold of minus one point six four five, so it defaults, at point six of a year.
- Brook's draw, minus point five six one, is above its threshold of minus one point eight eight one, so it survives, until about eleven point one years.
- A low draw means an early default.
- The default time carries more than the yes or no of the threshold: it says when, not just whether.

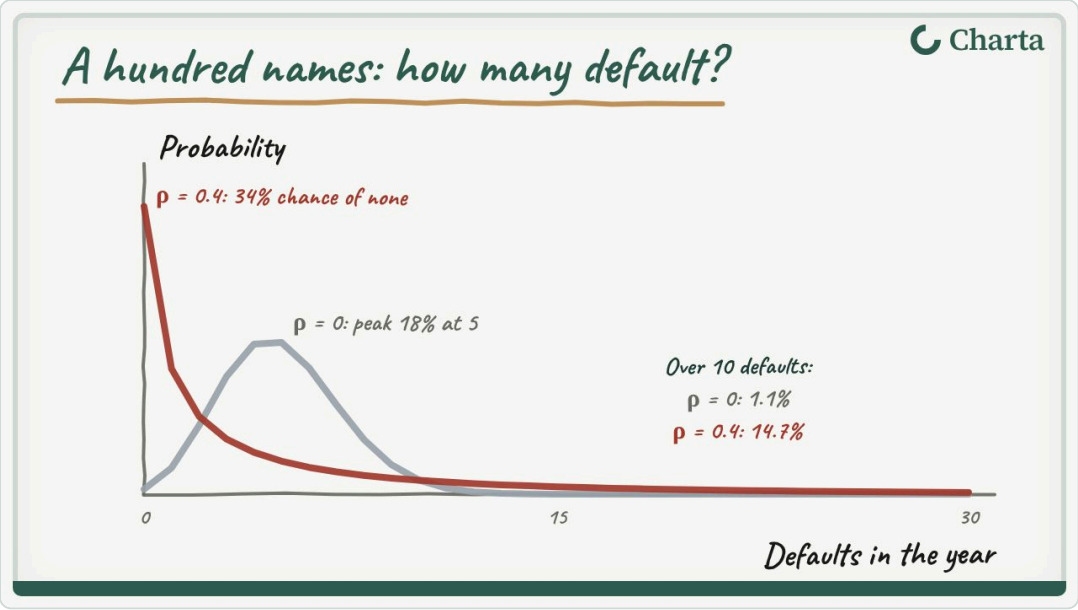
27 Simulating the whole portfolio



IN THE LESSON

- Here is the whole simulation.
- Look at the boxes.
- In one trial, first draw one common factor M, and an own shock Z sub i for every name.
- Second, build each X sub i, and then each uniform number U sub i as N of X sub i.
- Third, turn each into a default time with that name's hazard rate.
- Fourth, every name whose default time falls before the horizon has defaulted, and its loss, the exposure times the loss given default, is added to the portfolio loss for that trial.
- Now the many trials.
- Repeat the trial many thousands of times.
- Finally, read off the portfolio loss distribution: the expected loss, a high quantile such as the ninety nine per cent loss, or the losses that fall on each C D O tranche.

28 A hundred names: how many default?



IN THE LESSON

- This chart shows what the simulation produces for a hundred names, each with Ashford's one year default probability of five per cent.
- Along the bottom is the number of defaults in the year, from zero, through fifteen, to thirty.
- Up the side is the probability.
- The grey curve is rho of zero, independence.
- It is a tidy hump, peaking at about eighteen per cent for five defaults.
- The red curve is rho of point four.
- There is about a thirty four per cent chance of no defaults at all, yet a long right tail.
- More than ten defaults has a chance of about one point one per cent under independence, but about fourteen point seven per cent at rho of point four.
- Both curves average five defaults, yet the tail is far fatter.

29 Trap: ρ instead of √ρ

Wrong: $X = \rho \times M + \sqrt{(1 - \rho)} \times Z$, same draws $M = -1.5$, $Z_A = -1.2$

Variance of X, not 1	0.4^2 + 0.6	**0.76**
Covariance of two names	0.4 × 0.4	**0.16**
Correlation, not 0.4	0.16 ÷ 0.76	**0.21**
Ashford's draw X_A	0.4 × (−1.5) − 0.9295	**−1.5295**
Against threshold −1.6449	−1.5295 is above it	**No default**

IN THE LESSON

- Now the frequent error: putting rho, instead of the square root of rho, on the common factor.
- With rho of point four, the variance of X becomes point four squared plus point six, which is point seven six, not one.
- So X is no longer standard normal, and the thresholds no longer match the default probabilities.
- Two names now have a covariance of point four times point four, point one six, and a correlation of point one six divided by point seven six, about point two one, not point four.
- Rerun Ashford's trial with the same draws.
- Its common part becomes point four times minus one point five, which is minus point six.
- Adding the own part of minus point nine two nine five gives a draw of minus one point five two nine five.
- That is above the threshold, so the wrong model misses Ashford's default.

30 Why it matters, and the model's limits

- Joint defaults drive portfolio tails and tranche losses
- Changing ρ reshapes the tail a lot
- Gaussian copula: thin joint tails, may understate clusters
- One ρ for every pair is a strong simplification

IN THE LESSON

- Why does all this matter?
- Joint defaults drive the tail of a credit portfolio's loss distribution and the losses on C D O tranches, so the copula and the value of rho move the capital and the prices that depend on them.
- Changing rho reshapes the tail a lot, as the chart showed.
- Two limits are worth knowing.
- The Gaussian copula has thin joint tails: extreme moves in two hidden variables at the same time are rarer than in many real crises, so it can understate the chance of many defaults together.
- And using one rho for every pair is a strong simplification, convenient but rarely true.
- The model is a tool for reasoning about dependence, not a fact about the world.

31 Recap: the main ideas

- Copula joins marginals: Joint = C(u, v), by Sklar
- Bounds: max(u + v − 1, 0) to min(u, v); independence u × v
- Gaussian: x = N^{−1}(Q); P_AB = M(x_A, x_B, ρ)
- Default correlation is not ρ, and is much smaller
- X_i = √ρ M + √(1−ρ) Z_i; U_i = N(X_i); Time = −ln(1−U_i)/λ

IN THE LESSON

- Let us pull it together.
- A copula joins separate marginal distributions into one joint distribution and carries all of the dependence: by Sklar's theorem, the joint probability is C of u and v.
- Every copula lies between the larger of u plus v minus one and zero, and the smaller of u and v, with independence at u times v.
- The Gaussian copula maps each default probability to a threshold, inverse N of Q, and gives the joint default probability as capital M, the bivariate normal, at the two thresholds.
- Default correlation is not rho.
- It is the covariance of the indicators over the product of their standard deviations, and much smaller.
- For a portfolio, the one factor draw puts root rho on the common factor M, becomes a uniform number through N, and then a default time through the hazard rate.

32 Try three quick questions

1. u = 0.9, v = 0.6. Independence? Lower and upper bounds?
   Answer: 0.54. Lower 0.9 + 0.6 − 1 = 0.5; upper 0.6.
2. P_A = 4%, P_B = 2%, P_AB = 0.5%. Default correlation?
   Answer: (0.5% − 0.08%) ÷ (0.196 × 0.14) ≈ 15.3%.
3. ρ = 0.25, M = −2, Z = 0, λ = 2%. Default time?
   Answer: X = −1, U = 0.1587; −ln(0.8413) ÷ 0.02 ≈ 8.6 years.

IN THE LESSON

- Now try three quick questions.
- Pause the video, think, then listen to the answers.
- Question one.
- Two marginal cumulative probabilities are point nine and point six.
- What is the joint probability under independence, and what are the bounds?
- Independence gives point five four.
- The lower bound is point nine plus point six minus one, point five, and the upper bound is point six.
- Question two.
- Two names have default probabilities of four and two per cent, and a joint default probability of point five per cent.
- What is the default correlation?
- The covariance is point four two per cent.
- The standard deviations are about point one nine six and point one four, so the answer is about fifteen point three per cent.
- Question three.
- Rho is point two five, M is minus two, Z is zero and the hazard rate is two per cent.
- What is the default time?
- X is minus one, and U is about point one five eight seven, so the time is about eight point six years.