CHARTA LEARNING · LESSON NOTES

Credit Value at Risk

FRM Part II

1 A story: two desks at Kestrel Bank

- Same question: how much could we lose?
- Trading desk: liquid positions, look a few days ahead
- Lending desk: illiquid loans, look a full year ahead
- Different horizons lead to different tools

IN THE LESSON

- Kestrel Bank has two desks.
- The trading desk holds bonds and currencies that can be sold within hours.
- The lending desk has made loans to companies that will not be repaid for years.
- The chief risk officer asks both desks the same thing: with ninety nine per cent confidence, what is the most we could lose?
- The trading desk answers in a moment, looking just a day or two ahead, because it could close its positions if trouble came.
- The lending desk cannot close anything.
- If a borrower fails, the bank is stuck for months.
- So it looks a full year ahead.
- The question is the same, but the two answers are built in very different ways.

2 Market VaR versus credit VaR

Same idea, different construction

	Market VaR	Credit VaR
Horizon	1-10 days	About 1 year
Shape of P&L	Symmetric	Skewed, fat tail
Tools	History, var-covar, MC	PD, LGD, EAD, correlation
Confidence	Often 99%	Often 99.9%
Quoted	From today's value	Relative to EL

IN THE LESSON

- Here are the differences side by side.
- Market value at risk covers a short horizon, typically one to ten days, because positions are liquid and can be closed or hedged.
- Its profit and loss is roughly symmetric, and it is measured from observed prices by historical simulation, variance covariance or Monte Carlo.
- Credit value at risk covers about a year, because loans are illiquid and defaults are slow to resolve.
- Losses come from rare, lumpy defaults, so the distribution is skewed with a fat right tail.
- It must be modelled from default probabilities, recovery rates, exposures and default correlation, often at ninety nine point nine per cent.
- It is usually quoted relative to expected loss.

3 Example 1: normality fails for credit

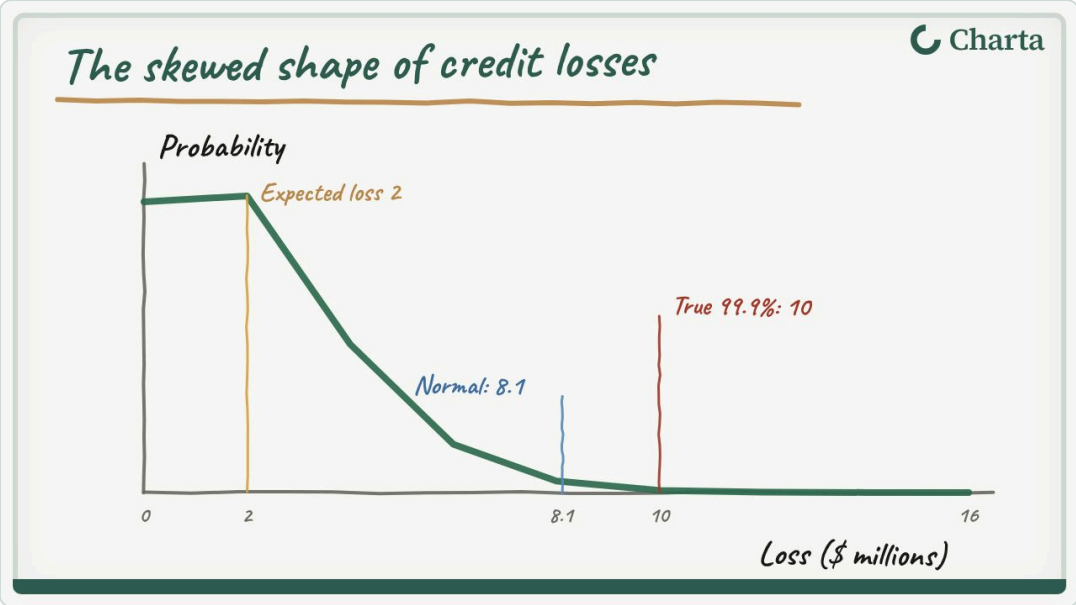
50 loans, PD 2%, loss $2m each, independent. $ millions

Expected loss	50 × 2% × 2	**2.00**
SD of loss	2 × √(50×0.02×0.98)	**1.98**
Normal 99.9% loss	2 + 3.09 × 1.98	**8.12**
True 99.9% loss	5 defaults × 2	**10.00**
Understated by	10.00 − 8.12	**1.88**

IN THE LESSON

- Let us see why the shortcuts of market risk do not carry over.
- Kestrel has fifty small loans.
- Each has a two per cent default probability and loses two million dollars if it defaults.
- Defaults are independent.
- Expected loss is fifty times point zero two times two million, which is two million dollars.
- The standard deviation of loss is about one point nine eight million.
- A normal approximation at ninety nine point nine per cent adds three point zero nine standard deviations, giving eight point one two million.
- But the true distribution is skewed.
- Four or fewer defaults has probability only ninety nine point six eight per cent, so the true quantile needs five defaults, ten million dollars.
- Normality understates the tail.

4 The skewed shape of credit losses



IN THE LESSON

- Here is the same fifty loan book as a picture.
- Along the bottom is the number of defaults, turned into dollars, two million for each.
- Up the side is the probability of each outcome.
- Most years show nought, one or two defaults, so the line peaks on the left near the expected loss of two million dollars.
- But the right side does not fall away symmetrically.
- It has a long tail, because a few unlucky years bring five or six defaults.
- The red line marks the true ninety nine point nine per cent loss at ten million dollars, and the blue line marks where a normal curve would put it, at about eight million.
- Credit value at risk will be measured from the expected loss out to that red line.

5 What credit VaR is

Credit VaR = Loss quantile − Expected loss
EL = PD × LGD × EAD

Loss quantile	Loss at the confidence level
Expected loss	Average loss: priced, provisioned
Credit VaR	Unexpected loss: basis of capital

IN THE LESSON

- Now the definition.
- Credit value at risk is the loss that will not be exceeded at a chosen confidence level over a chosen horizon, measured relative to expected loss.
- You find the loss quantile, then subtract the mean.
- The result is the unexpected loss, and it is the basis for economic capital.
- Expected loss is covered by loan pricing and provisions, so capital is held only for the excess.
- For a single exposure, expected loss is the product of three things: the probability of default, the loss given default, which is the share of the exposure not recovered, and the exposure at default, the amount owed when the borrower fails.

6 Example 2: expected loss on one loan

Loan $20m, PD 1.5%, recovery 55%

Exposure at default	full amount drawn	**$20m**
Loss given default	1 − 55%	**45%**
Expected loss	1.5% × 45% × 20m	**$135,000**
Wrong: forget LGD	1.5% × 20m	**$300,000**

IN THE LESSON

- Start with the simplest piece, expected loss.
- Kestrel lends twenty million dollars to a manufacturer called Brightfield Foods.
- The probability of default over a year is one point five per cent.
- If Brightfield fails, lenders expect to recover fifty five per cent, so loss given default is forty five per cent.
- The exposure at default is the full twenty million dollars.
- Expected loss is point zero one five times point four five times twenty million, which is one hundred and thirty five thousand dollars.
- The classic slip is forgetting loss given default.
- That would give point zero one five times twenty million, three hundred thousand dollars, more than twice the right answer.

7 Example 3: a discrete loss distribution

Loss distribution of Kestrel's small book. $ millions

Loss 0	Probability 90.0%	**Cum. 90.0%**
Loss 2	Probability 6.0%	**Cum. 96.0%**
Loss 5	Probability 2.5%	**Cum. 98.5%**
Loss 10	Probability 1.0%	**Cum. 99.5%**
Loss 20	Probability 0.5%	**Cum. 100.0%**

IN THE LESSON

- Now a whole portfolio.
- Kestrel's small book has a loss distribution with just five possible outcomes, in millions of dollars.
- Zero loss has probability ninety per cent, two million six per cent, five million two point five per cent, ten million one per cent, and twenty million half of one per cent.
- The cumulative probability adds these up from the bottom: ninety, ninety six, ninety eight point five, ninety nine point five, and finally one hundred per cent.
- To read a quantile from a discrete table, you look down the cumulative column for the first loss whose cumulative probability reaches at least the confidence level.
- We apply this on the next slide.

8 Example 3 continued: reading the quantile

Same book, 99% confidence. $ millions

Expected loss	0.12 + 0.125 + 0.1 + 0.1	**0.445**
99% quantile	first cum. at least 99%	**10**
Credit VaR	10 − 0.445	**9.555**
Wrong: quantile alone	no EL deducted	**10**
Wrong: loss 5	cum. 98.5% is below 99%	**5**

IN THE LESSON

- First the mean.
- Multiply each loss by its probability and add: zero point one two, plus point one two five, plus point one, plus point one, which is point four four five million.
- Now the ninety nine per cent quantile.
- At a loss of five, the cumulative probability is only ninety eight point five per cent, below ninety nine.
- At ten it reaches ninety nine point five, so the quantile is ten million.
- Credit value at risk is ten minus point four four five, which is nine point five five five million dollars.
- Notice the two traps.
- Quoting ten alone forgets expected loss, and reading five reads the wrong point.
- At ninety nine point nine per cent, the quantile would move to twenty million and credit value at risk to nineteen point five five five.

9 Credit VaR: mistakes to avoid

- Quantile alone is not credit VaR: subtract EL
- Discrete table: first loss with cumulative at least the level
- EL = PD × LGD × EAD; do not forget LGD
- EL is covered by pricing and provisions
- Capital is held only for the unexpected loss

IN THE LESSON

- Let us collect the exam traps for this outcome.
- First, never quote the quantile as if it were credit value at risk.
- The measure is the quantile minus expected loss, because expected loss is already paid for by pricing and provisions.
- Second, on a discrete distribution take the first loss whose cumulative probability is at least the confidence level, not the one just below it, and not the one just above.
- Third, do not forget loss given default when computing expected loss: it is probability of default times loss given default times exposure at default.
- Finally, remember why the subtraction matters.
- Capital covers only the unexpected part, and that is why credit value at risk is the basis for economic capital.

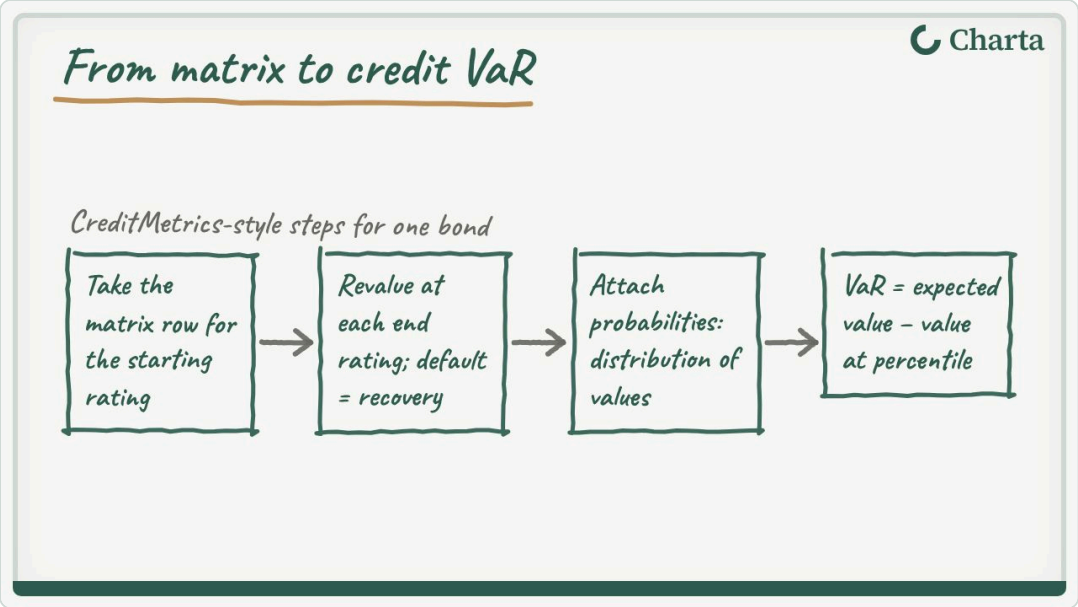
10 Rating transition matrices

- One row per starting rating
- Columns: probability of each end rating, with default
- Each row sums to 100%
- Default is the final, absorbing state

IN THE LESSON

- Kestrel owns a bond issued by Brightfield Foods, rated triple B.
- Over a year, Brightfield might be upgraded, stay the same, be downgraded, or default.
- A rating transition matrix records these chances.
- For each starting rating, it gives the probability of being in each rating class, including default, at the end of the horizon.
- Each row is one starting rating, and each row adds up to one hundred per cent, because the bond must end up somewhere.
- Default is the final state, since a defaulted firm does not climb back to a rating.
- This matrix is the raw material for a family of methods, usually called CreditMetrics style, that measure credit value at risk by tracking changes in rating.

11 From matrix to credit VaR



IN THE LESSON

- Follow the boxes from left to right.
- First, take the row of the matrix for the bond's starting rating.
- Second, for each possible end rating, revalue the bond.
- Use the forward credit spread curve for that rating to discount the remaining cash flows, so a downgraded bond is worth less even though it has not defaulted.
- If the bond defaults, use its recovery value instead.
- Third, attach the transition probability to each value.
- That gives a distribution of bond values.
- The expected value is the probability weighted sum of those values.
- Finally, the credit measure is the expected value minus the value at the chosen percentile of that distribution.

12 Example 4: revaluing the bond

Bond: face 100, coupon 6%, 2 payments left. Risk-free 3%

Ends A: 3% + 1.0% = 4.0%	$6/1.04 + 106/1.04^2$	**103.77**
Ends BBB: 3% + 1.5% = 4.5%	$6/1.045 + 106/1.045^2$	**102.81**
Ends BB: 3% + 3.5% = 6.5%	$6/1.065 + 106/1.065^2$	**99.09**
Ends in default	recovery value	**40.00**

IN THE LESSON

- Here are made up numbers.
- The bond has a face value of one hundred, a six per cent annual coupon, and two payments left after the horizon.
- The risk free rate is three per cent.
- To keep it simple, each rating has one spread: one per cent for single A, one and a half per cent for triple B, and three and a half per cent for double B.
- Discount six at the first date and one hundred and six at the second date.
- At single A the bond is worth about one hundred and three point seven seven.
- At triple B, one hundred and two point eight one.
- At double B, ninety nine point oh nine.
- If the bond defaults, we use the recovery value of forty.

13 Example 4 continued: the credit measure

Matrix row for BBB: A 6%, BBB 86%, BB 5%, default 3%

Check on probabilities	6 + 86 + 5 + 3	**100%**
Expected value	6.23+88.42+4.95+1.20	**100.80**
5th percentile value	worst 3% + 2% of BB	**99.09**
Credit measure, 95%	100.80 − 99.09	**1.71**

IN THE LESSON

- Now attach probabilities from the matrix row: six per cent to end at single A, eighty six per cent to stay triple B, five per cent to fall to double B, and three per cent to default.
- They add to one hundred.
- The expected value is six per cent of one hundred and three point seven seven, plus eighty six per cent of one hundred and two point eight one, plus five per cent of ninety nine point oh nine, plus three per cent of forty.
- That is six point two three, plus eighty eight point four two, plus four point nine five, plus one point two oh, or about one hundred point eight oh.
- Take the fifth percentile, for illustration.
- The worst three per cent is default, and the next two per cent comes from double B, so the fifth percentile value is ninety nine point oh nine.
- The measure is one hundred point eight oh minus ninety nine point oh nine, or one point seven one.

14 Example 5: more than one year

States: Good, Weak, Default. Two-year matrix = matrix × matrix

Good to default, year 1	from the matrix	**1.00%**
Good to default, 2 years	0.94×0.01 + 0.05×0.10 + 0.01×1	**2.44%**
Good to good, 2 years	$0.94^2 + 0.05 \times 0.10$	**88.86%**
Ignoring migration	$1 - 0.99^2$	**1.99%**

The matrix alone says nothing about correlation between obligors

IN THE LESSON

- What about several years?
- Multiply the matrix by itself.
- Take a tiny three state matrix: good, weak and default.
- A good borrower has a ninety four per cent chance of staying good, five per cent of turning weak, and one per cent of defaulting.
- A weak borrower has ten per cent, eighty per cent and ten per cent.
- Default is absorbing, with a one in its own row.
- For two years, the chance that a good borrower ends in default is point nine four times point oh one, plus point oh five times point one, plus point oh one times one, which is two point four four per cent.
- Staying good is point nine four squared plus point oh five times point one, or eighty eight point eight six.
- Compounding one per cent while ignoring migration gives only one point nine nine, so it understates risk.
- The matrix alone says nothing about correlation between obligors.
- It describes one borrower at a time.

15 Vasicek: one bad year for everybody

- One systematic factor: the state of the economy
- Plus an idiosyncratic factor for each borrower
- Asset correlation ρ; large, granular portfolio
- Bad factor outcome: default rate becomes the worst case
- Basel II IRB uses 99.9% over one year

IN THE LESSON

- Now a different route to credit value at risk, one that regulators use.
- Imagine Kestrel lends to thousands of small firms.
- In a calm year, about one per cent of them fail.
- In a deep recession, many more fail together, because they share the same economy.
- Vasicek's model captures this with two ingredients.
- Each borrower's asset value is driven by one systematic factor, the state of the economy, plus an idiosyncratic factor of its own.
- The weight on the common factor is set by the asset correlation, rho.
- The portfolio is large and granular, meaning no single loan matters much, so the own shocks average out.
- Then a bad economy fixes the default rate at a worst case figure.

16 The worst-case default rate

$WCDR = N\left[\left(N^{-1}(PD) + \sqrt{\rho} \times N^{-1}(0.999) \right) \div \sqrt{(1 - \rho)} \right]$

Probability of default	Average default probability
Asset correlation	Common-factor weight is $\sqrt{\rho}$
Confidence level	99.9% over one year in IRB
Normal functions	N: normal cdf. N^{-1}: its inverse
Worst-case default rate	Default rate in a bad year

IN THE LESSON

- Here is the formula.
- The worst case default rate, W C D R, is N of the quantity: inverse N of the default probability, plus the square root of rho times inverse N of nought point nine nine nine, all divided by the square root of one minus rho.
- N is the standard normal cumulative distribution function, and inverse N is its inverse.
- The Basel two internal ratings based approach applies it at ninety nine point nine per cent over one year.
- Look at the pieces.
- With rho of zero the formula returns the default probability itself, since there is no common factor.
- The higher the correlation, the further the worst case sits above the average.

17 From worst-case rate to capital

$K = LGD \times (WCDR - PD)$
$RWA = K \times 12.5 \times EAD$

Capital per unit	Capital per $1 of exposure
PD deducted	PD deducted: EL is provisioned
Maturity adjustment	Multiplies K where required
Risk-weighted assets	$RWA = K \times 12.5 \times EAD$
Why 12.5	12.5 = 1 ÷ 8% minimum ratio

IN THE LESSON

- Next, turn the worst case rate into capital.
- Expected loss, probability of default times loss given default, is handled by provisions and pricing.
- That is why the default probability is deducted from the worst case rate.
- Capital per unit of exposure, before any maturity adjustment, is K equals loss given default times the worst case rate minus the default probability.
- Where the approach requires it, a maturity adjustment multiplies K.
- Risk weighted assets are then K times twelve point five times the exposure at default.
- Twelve point five is the reciprocal of eight per cent, the minimum capital ratio, so eight per cent of risk weighted assets equals the capital K times exposure.

18 Example 6: IRB capital for a loan book

EAD $200m, PD 1.5%, LGD 45%, ρ = 0.15, 99.9%

Numerator	−2.170 + 0.387 × 3.090	**−0.973**
Divide by √0.85	−0.973 ÷ 0.922	**−1.056**
WCDR	N(−1.056)	**14.56%**
K	45% × (14.56% − 1.5%)	**5.88%**
RWA	5.876% × 12.5 × 200m	**$146.9m**

IN THE LESSON

- Let us compute.
- Kestrel's small business book has an exposure of two hundred million dollars.
- The default probability is one point five per cent, loss given default is forty five per cent, and rho is point one five.
- Inverse N of point oh one five is minus two point one seven.
- Root point one five is point three eight seven, times three point oh nine, which adds one point one nine seven.
- The numerator is minus point nine seven three.
- Root point eight five is point nine two two, so the ratio is minus one point oh five six, and N of that is fourteen point five six per cent.
- K is point four five times fourteen point five six minus one point five, which is five point eight eight per cent.
- Risk weighted assets are about one hundred and forty six point nine million dollars.

19 Vasicek and IRB: common errors

- Use √ρ in the formula, not ρ
- Deduct PD: capital covers only the unexpected part
- Maturity adjustment multiplies K where required
- RWA = K × 12.5 × EAD; capital is 8% of RWA
- Needs a large, granular portfolio

IN THE LESSON

- The same example gives a useful check.
- Capital itself is K times exposure, about eleven point seven five million dollars, and that equals eight per cent of the risk weighted assets.
- Now the exam traps.
- The first is using rho where the formula needs its square root.
- Rho is the correlation, but the weight on the common factor is the square root of rho.
- The second is omitting the deduction of the default probability, which would charge capital for expected loss that provisions already cover.
- Third, remember that the maturity adjustment multiplies K where the approach requires it.
- And do not apply Vasicek to a portfolio dominated by a few big loans, because the model needs a large, granular portfolio.

20 Three models, three loss distributions

One question, three constructions

	Vasicek	CR+	CreditMetrics
Type	Factor	Actuarial	Structural
Events	Default	Default	Migration
Correlation	Common ρ	Shared factors	Asset returns
Solved by	Closed form	Recursion	Monte Carlo

IN THE LESSON

- Three famous models build the distribution of default losses in different ways.
- Vasicek's model uses one normally distributed systematic factor and a common asset correlation.
- CreditRisk plus is actuarial: it models the number of defaults directly, with no asset values.
- CreditMetrics is structural and multi state: correlated asset returns drive rating migrations, and bonds are revalued at every outcome.
- Each also handles default correlation differently, as the table shows.
- In Vasicek, the correlation is the single common asset correlation.
- In CreditRisk plus, it comes from shared background factors.
- In CreditMetrics, it comes from correlated asset returns.
- We take each model in turn.

21 Vasicek's model: one factor, closed form

- One normally distributed systematic factor
- Common asset correlation ρ for all borrowers
- Large homogeneous portfolio: own shocks average out
- Closed-form distribution of the default rate

IN THE LESSON

- In Vasicek's model, each borrower's asset value has a systematic part and an idiosyncratic part.
- The systematic factor is normally distributed and common to everyone, which is how default correlation enters.
- All borrowers share one asset correlation, rho.
- For a large homogeneous portfolio, meaning many loans of equal size and equal default probability, the idiosyncratic parts average out.
- The default rate then depends only on the systematic factor, and the whole distribution of the default rate has a closed form.
- There is no simulation needed.
- The worst case default rate we used for Basel is simply one point of that distribution, the point at ninety nine point nine per cent.
- The price of simplicity is the assumption of a single factor and one common correlation.

22 CreditRisk+: an actuarial model

- Number of defaults is Poisson
- Mean default rates are stochastic
- Shared background factors move intensities together
- Exposures banded; analytic recursion gives the losses
- Default is the only event

IN THE LESSON

- CreditRisk plus borrows from insurance.
- It does not model asset values at all.
- The number of defaults in a year is taken to be Poisson, which is the standard distribution for rare, independent events, but with a twist: the mean default rates are stochastic, meaning they themselves vary from year to year.
- Obligors share background factors, such as the state of the economy in a sector, that move their default intensities together.
- That is how default correlation arises.
- Exposures are grouped into bands, and the loss distribution then comes from an analytic recursion, a repeated calculation with no simulation.
- Default is the only event considered: there is no migration, and no revaluation of bonds.

23 Example 7: shared factors fatten the tail

Expected defaults 3 per year. Chance of 8 or more

Fixed mean 3	Poisson tail	**1.19%**
Mean 1 (half of years)	Poisson tail	**0.001%**
Mean 5 (half of years)	Poisson tail	**13.34%**
Mixture of the two	½ × 0.001% + ½ × 13.34%	**6.67%**
Ratio to fixed mean	6.67 ÷ 1.19	**5.6 times**

IN THE LESSON

- Let us see why stochastic means matter.
- Suppose Kestrel's book expects three defaults a year.
- If the count is Poisson with a fixed mean of three, the chance of eight or more defaults is about one point one nine per cent.
- Now let the background conditions vary.
- In half of years the mean is one, and in the other half it is five, so the average is still three.
- The chance of eight or more is almost nothing in the calm years, and about thirteen point three four per cent in the stressed years.
- Averaging gives six point six seven per cent, more than five times larger.
- Same average, but a much fatter tail.
- That is correlation at work.

24 CreditMetrics: structural and multi-state

- Structural, multi-state: migration and default
- Correlated asset returns, typically from equity data
- Returns drive rating changes via thresholds
- Bonds revalued at each outcome
- Distribution found by Monte Carlo simulation

IN THE LESSON

- CreditMetrics is the third approach.
- It is structural, meaning default and migration are linked to the value of the firm's assets.
- Each obligor has a standard normal asset return.
- Returns are correlated, typically estimated from equity data, and that is how default correlation enters.
- Each return is compared with thresholds taken from the rating transition matrix, so a very low return means default and a moderately low one means a downgrade.
- Every bond is then revalued at its new rating, or at its recovery value if it defaulted.
- With many obligors the distribution of portfolio value is usually found by Monte Carlo simulation, repeating thousands of trials.
- Unlike CreditRisk plus, migration losses are captured.

25 Credit spread risk

- Spread risk: spreads change, value moves, no default
- Migration models fix the forward spreads
- So they leave spread risk out
- Credit VaR is understated when spreads are volatile
- Fix: simulate spreads, or add a spread VaR

IN THE LESSON

- There is one more risk we have not counted.
- A credit spread is the extra yield a risky bond pays over a risk free bond.
- Credit spread risk is the risk that spreads change and so move the market value of a credit risky position, even though no default or rating change occurs.
- A bond can lose value simply because investors demand more compensation for risk.
- Migration based models such as CreditMetrics revalue each rating state with fixed forward spreads.
- They therefore leave this risk out, and understate credit value at risk when spreads are volatile.
- Including it, by simulating spreads or adding a spread value at risk, widens the distribution.

26 Measuring a spread move

Loss = Market value × Spread duration × Change in spread
Implied PD = Spread ÷ LGD

Market value	Market value of the position
Spread duration	% price move per 1% spread
Change in spread	Widening, as a decimal
Implied PD	Overstates real-world PD

IN THE LESSON

- For a small change in spreads, a quick approximation works.
- The loss equals the market value, times the spread duration, times the change in spread.
- Spread duration measures how many per cent a position's price falls for each one per cent rise in its credit spread.
- There is a second use of spreads.
- A spread contains compensation for expected loss, so with no risk premium, the spread equals probability of default times loss given default.
- Rearranged, the default probability implied by the spread is the spread divided by loss given default.
- But spreads also include a risk premium, so a default probability implied from spreads overstates the real world default probability.

27 Example 8: a spread move

Bond book $500m, spread duration 4.5. No default, no downgrade

Spread widens	40 basis points	**0.40%**
Loss	500m × 4.5 × 0.40%	**$9.0m**
Per 10 basis points	500m × 4.5 × 0.10%	**$2.25m**
Fixed-spread model shows	no spread move	**$0**

IN THE LESSON

- Let us put numbers on it.
- Kestrel's bond book has a market value of five hundred million dollars and a spread duration of four point five.
- Suppose spreads widen by forty basis points, which is point four per cent.
- Nobody defaults and no rating changes.
- The loss is five hundred million times four point five times point oh oh four, which is nine million dollars.
- Each ten basis points costs two point two five million.
- A migration model with fixed forward spreads would show none of this.
- If spreads are volatile, that is a real risk the model does not see, so the measure of credit value at risk is too low.

28 Example 9: adding spread risk

Migration credit VaR $30m, spread VaR $12m. $ millions

Independent	$\sqrt{(30^2 + 12^2)}$	**32.3**
Correlation 0.6	$\sqrt{(900 + 144 + 2\times0.6\times30\times12)}$	**38.4**
Perfect correlation	30 + 12	**42.0**
Spread risk left out	38.4 − 30	**8.4 short**

IN THE LESSON

- How should the two pieces be combined?
- Suppose a migration model gives a credit value at risk of thirty million dollars, and a separate spread value at risk is twelve million.
- Spreads tend to widen when defaults rise, so the two components are positively correlated.
- If they were independent, the total would be the square root of thirty squared plus twelve squared, which is about thirty two point three.
- With a correlation of point six, add two times point six times thirty times twelve, and the total is about thirty eight point four.
- With perfect correlation it would be forty two.
- Treating them as independent understates the total, and leaving spread risk out understates it by more than eight million dollars.

29 Example 10: a PD implied by a spread

Spread 1.8%, LGD 60%. Assume real-world PD is 2%

Implied PD	1.8% ÷ 60%	**3.0%**
Real expected loss	2% × 60%	**1.2%**
Risk premium in spread	1.8% − 1.2%	**0.6%**
Implied PD versus real PD	3.0% versus 2.0%	**overstated**

IN THE LESSON

- Finally, the implied default probability.
- Suppose a Brightfield bond trades at a spread of one point eight per cent, and loss given default is sixty per cent.
- Ignoring the risk premium, the implied default probability is one point eight divided by sixty, which is three per cent.
- Now suppose the real world default probability is only two per cent.
- The real expected loss is two per cent times sixty per cent, one point two per cent.
- The remaining point six per cent of spread is the risk premium that investors demand for bearing the risk.
- So a default probability implied from spreads overstates the real world one, and should not be used as a forecast.

30 Mistakes to avoid

- No √time scaling or normality for credit VaR
- A transition matrix gives no default correlation
- CreditRisk+: default only. CreditMetrics: migration too
- Migration models with fixed spreads omit spread risk
- Spread-implied PD overstates the real-world PD

IN THE LESSON

- Here are the traps that appear across this lesson.
- First, do not carry over market risk habits.
- Square root of time scaling and normality do not work for credit, where losses are skewed with a fat tail and the horizon is a year.
- Second, do not treat the transition matrix as a model of correlation: it describes one borrower at a time, and the correlation must come from elsewhere.
- Third, remember that CreditRisk plus considers default only, while CreditMetrics also captures migration.
- Fourth, a migration based figure leaves out credit spread risk.
- And finally, a default probability read from a spread includes a risk premium and overstates the real one.

31 Recap: the main ideas

- Credit VaR = loss quantile − EL; 1 year, skewed, 99.9%
- Matrix: revalue at each end rating; row sums to 100%
- IRB: K = LGD × (WCDR − PD); RWA = K × 12.5 × EAD
- Vasicek: factor. CR+: Poisson. CreditMetrics: migration
- Spread risk widens the distribution

IN THE LESSON

- Let us pull it together.
- Market and credit value at risk ask the same question, but credit covers a year, is skewed with a fat right tail, and is built from default probabilities, recoveries, exposures and correlation.
- Credit value at risk is the loss quantile minus expected loss, which is the basis for economic capital.
- Transition matrices give probabilities for each end rating, and bonds are revalued at each.
- Vasicek's worst case default rate gives K and risk weighted assets under the internal ratings based approach.
- The three models differ in how they build the distribution and the correlation, and spread risk widens the distribution.

32 Try three quick questions

1. Loss at 99.9% is $45m; EL is $12m. Credit VaR?
   Answer: 45 − 12 = $33 million.
2. EAD $50m, LGD 40%, PD 1%, WCDR 10%. K and RWA?
   Answer: K = 0.4 × 9% = 3.6%. RWA = 3.6% × 12.5 × 50m = $22.5m.
3. Book $200m, spread duration 3, spreads +50 bp. Loss?
   Answer: 200m × 3 × 0.5% = $3 million.

IN THE LESSON

- Now try three quick questions.
- Pause the video, think, then listen to the answers.
- Question one.
- A portfolio's loss at the ninety nine point nine per cent level is forty five million dollars, and its expected loss is twelve million dollars.
- What is the credit value at risk?
- Forty five minus twelve, thirty three million dollars.
- Question two.
- A loan has an exposure of fifty million dollars, a loss given default of forty per cent, a default probability of one per cent, and a worst case default rate of ten per cent.
- What are K and the risk weighted assets?
- K is point four times point oh nine, or three point six per cent.
- Risk weighted assets are that times twelve point five times fifty million, twenty two point five million dollars.
- Question three.
- A bond book of two hundred million dollars has a spread duration of three.
- Spreads widen by fifty basis points.
- What is the loss?
- Two hundred million times three times point oh oh five, three million dollars.