CHARTA LEARNING · LESSON NOTES

Future Value and Exposure

FRM Part II

1 A story: the swap that might go either way

- Swap is worth zero on day one, then rates move
- Future value V can be positive or negative to us
- We lose only if the counterparty fails while V is positive
- Exposure = max(V, 0): the positive half of the value
- V is unknown, so we model many scenarios

IN THE LESSON

- Northgate Bank has agreed a five year interest rate swap with Corvane Energy.
- A swap is a contract to exchange interest payments for several years.
- On the day it is signed it is worth nothing to either side.
- But interest rates will move.
- Next year the swap might be worth ten million dollars to Northgate, or minus six million.
- If it is worth ten million to Northgate, Corvane owes Northgate, and if Corvane then fails, Northgate loses.
- If it is worth minus six million, Northgate owes Corvane, and a failure of Corvane brings no loss on that contract, because Northgate still has to pay what it owes.
- So credit exposure is the positive half of a future value that nobody can know today.
- This lesson gives that idea numbers.

2 Exposure measures at one date

EE = average of max(V, 0)
ENE = average of min(V, 0)
Expected MtM = EE + ENE

Expected exposure, EE	Mean of max(V, 0) across scenarios
Expected negative exposure	Mean of min(V, 0), negatives kept
Expected mark-to-market	EE + ENE: negatives not floored
Potential future exposure	High percentile: 97.5% or 99%

IN THE LESSON

- Let V be the value of the contract to us at a future date.
- Exposure keeps only the positive part: the larger of V and zero.
- Because V is uncertain, we simulate many scenarios.
- Expected exposure, E E, is the average of exposure across the scenarios.
- Expected negative exposure, E N E, is the average of the smaller of V and zero, so it is a negative number, or zero.
- Expected mark-to-market is the plain average of V, with the negatives left in, and it equals E E plus E N E.
- Potential future exposure, P F E, is different.
- It is a high percentile of the exposure distribution, such as the ninety seventh point five or the ninety ninth.
- So E E is a mean, and P F E is a quantile.

3 Example 1: measures at one date

Values V of the swap in 10 equal scenarios ($m): −6, −4, −3, −1, 0, 2, 3, 5, 7, 10

EE	(2+3+5+7+10) ÷ 10	**$2.7m**
ENE	(−6−4−3−1) ÷ 10	**−$1.4m**
Expected MtM	2.7 + (−1.4)	**$1.3m**
PFE, 90th percentile	9th smallest exposure	**$7m**
Wrong: average V as EE	includes negatives	**$1.3m**

IN THE LESSON

- Take one future date for the Northgate swap, and ten equally likely scenarios for its value, in millions of dollars: minus six, minus four, minus three, minus one, zero, two, three, five, seven and ten.
- Exposure keeps only the positive values, so it is zero five times, then two, three, five, seven and ten.
- These add to twenty seven, so expected exposure is two point seven.
- The negative values add to minus fourteen, so expected negative exposure is minus one point four.
- Expected mark-to-market is the sum, one point three, which is also the plain average of all ten values.
- For a ninetieth percentile P F E, take the ninth smallest exposure: seven million.
- The trap is averaging the negatives into E E, which would give one point three instead of two point seven.

4 Exposure measures through time

Four ways to summarise a profile

Measure	What it is
Expected positive exposure	Average of EE through time
Effective EE	EE forced never to fall (running max)
Effective EPE	Average effective EE, first year
Maximum exposure	Peak of the PFE profile

IN THE LESSON

- Now we move through time.
- Each future date has its own E E and its own P F E, and together the dates form profiles.
- Four measures summarise a profile.
- Expected positive exposure, E P E, averages the E E profile through time.
- Effective E E fixes a quirk: it forces the profile never to fall, by taking a running maximum.
- The thinking is that when trades mature and exposure shrinks, new trades are usually written in their place.
- Effective E P E then averages that effective profile over the first year.
- Maximum exposure is the peak of the P F E profile, the highest high-percentile value at any date.
- Averages describe typical exposure.
- The maximum describes the extreme.

5 Example 2: a profile over the first year

Quarterly EE ($m): 2.0, 3.0, 2.6, 3.4. PFE: 4.5, 6.5, 6.0, 7.5

EPE	(2.0+3.0+2.6+3.4) ÷ 4	**$2.75m**
Effective EE	never falls: running max	**2.0, 3.0, 3.0, 3.4**
Effective EPE	(2.0+3.0+3.0+3.4) ÷ 4	**$2.85m**
Maximum exposure	peak of 4.5, 6.5, 6.0, 7.5	**$7.5m**

IN THE LESSON

- Here is the first year of the swap, at four quarter dates.
- Expected exposure is two, three, two point six and three point four million.
- E P E is the average: two plus three plus two point six plus three point four is eleven, divided by four, which is two point seven five.
- Effective E E is the running maximum.
- It starts at two, rises to three, and stays at three at the third date even though E E fell to two point six, then ends at three point four.
- Averaging gives an effective E P E of two point eight five, a little above E P E.
- The P F E at the four dates is four point five, six point five, six and seven point five, so the maximum exposure is seven point five million.
- Remember: E E is a mean and P F E is a quantile.

6 Exposure compared with value at risk

Both describe a loss distribution, built differently

	VaR	Exposure
Horizon	1 to 10 days	Every date to maturity
Counts	Gains and losses	max(V, 0) only
Needs	Today's positions	Model at each date, path
Calibration	Real-world data	Limits real, CVA risk-neutral

IN THE LESSON

- Value at risk and credit exposure both describe a distribution of possible outcomes, but they are built differently.
- Value at risk looks at the change in value of today's portfolio over a short horizon, typically one to ten days, and counts gains and losses alike.
- Exposure cannot stop at ten days, because a counterparty can default at any date until the last trade matures.
- So exposure is projected forward to every future date.
- And on each date it keeps only the positive side of the value.
- Exposure is the larger of V and zero, so a negative value counts as zero.
- The two measures also serve different jobs, and, as we see next, are calibrated differently.

7 Further considerations in measuring exposure

- Pricing model at every future date and path
- Cash flows paid before default must be modelled
- Netting and collateral applied to the right trades
- Limits: real-world data. CVA pricing: risk-neutral
- Scaling VaR by √time overstates long-dated exposure

IN THE LESSON

- Measuring exposure brings extra work that value at risk does not need.
- First, a pricing model is needed at each future date and on each simulated path, to revalue the trades in that scenario.
- Second, cash flows paid before the default must be modelled, since once paid they are no longer exposed.
- Third, netting and collateral must be applied to the right set of trades.
- Fourth, calibration.
- Exposure used for limits is usually calibrated to real-world data, because we want a realistic view of what could happen.
- Exposure used to price C V A, the credit valuation adjustment, is risk-neutral, consistent with market prices.
- Practice varies by institution.
- And a shortcut tempts people: scaling a short-horizon value at risk by the square root of time.
- The next example shows why that fails.

8 Example 3: the square-root-of-time trap

10-day VaR $0.8m. Swap has 5 years; look at year 4 (250 days a year)

Ten-day periods in 4 years	1,000 ÷ 10	**100**
Scaled VaR	0.8m × √100	**$8.0m**
Amortisation factor	(5 − 4) ÷ 5	**0.2**
Allowing for amortisation	8.0m × 0.2	**$1.6m**
Overstatement	8.0 ÷ 1.6	**5 times**

IN THE LESSON

- Here is the shortcut trap.
- Suppose the ten day value at risk of the Northgate swap is zero point eight million dollars.
- Someone scales it to four years by the square root of time.
- Taking two hundred and fifty business days a year, four years is one thousand days, which is one hundred periods of ten days each.
- The square root of one hundred is ten, so the scaled figure is eight million.
- But by year four the swap has only one year left of its five year life, so only one fifth of its sensitivity remains.
- The amortisation factor is five minus four, over five, which is point two.
- Adjusting gives one point six million.
- Scaling by the square root of time overstated the long-dated exposure five times.

9 What shapes an exposure profile

- Diffusion: uncertainty grows with time, so exposure rises
- No cash flows: SD of value grows with √time
- Amortisation: payments made, less left, so exposure falls
- Also: volatility, correlation, forward curve, netting
- Product shape = the balance of the two forces

IN THE LESSON

- What shapes the exposure profile of a trade?
- Two opposing forces.
- The first is diffusion: the uncertainty of the underlying risk factors, such as interest rates, grows as time passes.
- For a trade with no cash flows, the standard deviation of value grows with the square root of time.
- If it is two million dollars after one year, it is four million after four years, since the square root of four is two.
- So exposure rises.
- The second force is amortisation: as payments are made and maturity nears, the trade has less sensitivity left to market moves, so exposure falls.
- Each product's shape comes from the balance of the two.
- Volatility, correlation, the shape of the forward curve and netting also feed in.

10 How collateral changes exposure

Exposure = max(V − C , 0)

Trade value, V	Value at close-out
Collateral held, C	Value at the last margin call
Exposure left	Value can still move in the gap

IN THE LESSON

- Collateral is an asset the counterparty hands over to protect us.
- Let V be the trade value and C the collateral we hold.
- Collateral reduces exposure by the amount held, so what is left is the larger of V minus C, and zero.
- Note what it does not do: it does not remove exposure.
- The collateral we hold reflects the value at the last margin call, the last request for collateral.
- Between that call and the moment we close out the position after a default, the value can still move.
- So collateralised exposure depends on the margin period of risk, not on the maturity of the trade.
- We define that period properly in a few minutes.

11 Example 4: collateral and the gap

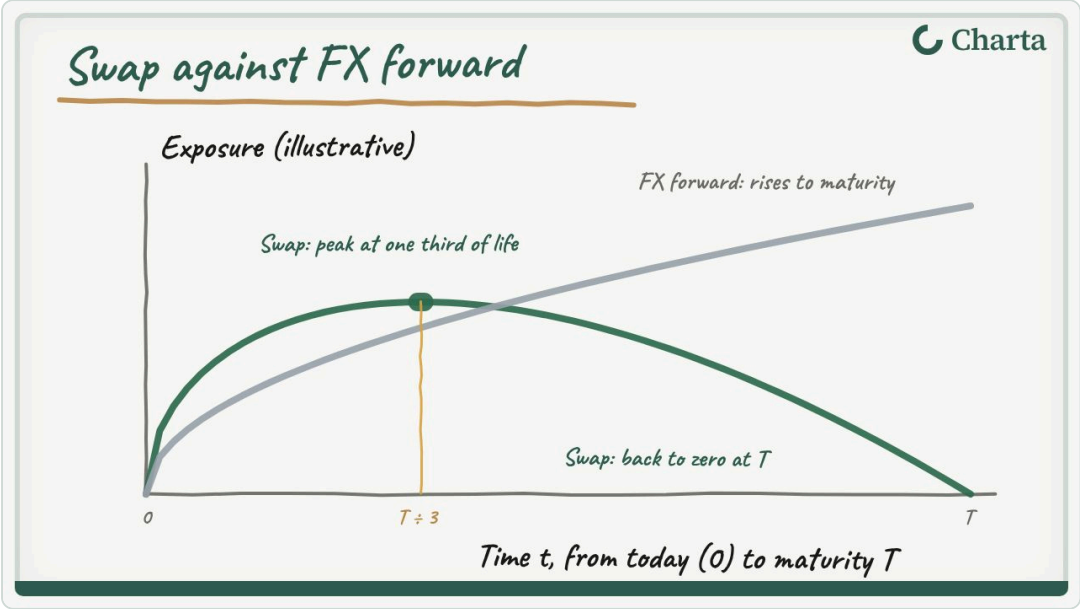
Collateral held $8m (value at last margin call). Corvane then defaults

Value at close-out	swap worth to us	**$11m**
Exposure with collateral	max(11 − 8, 0)	**$3m**
Exposure with none	max(11 − 0, 0)	**$11m**
If value had fallen to 6	max(6 − 8, 0)	**$0**

IN THE LESSON

- Suppose Northgate last called for margin when the swap was worth eight million dollars, and Corvane delivered eight million of collateral.
- Then Corvane defaults.
- By the time Northgate has closed out and replaced the trade, the swap is worth eleven million.
- Exposure is the larger of eleven minus eight and zero, which is three million.
- Without collateral it would have been eleven million.
- Now suppose instead the value had fallen to six million.
- Eight million is held, so six minus eight is negative and exposure is zero.
- Collateral covers most of the value, yet three million was still lost to value moving in the gap.
- That gap is the point of the margin period of risk.

12 Swap against FX forward



IN THE LESSON

- Now we compare the shapes of exposure for different products.
- Along the bottom is time, from today to maturity, which we call capital T.
- Up the side is exposure, in made-up units.
- The green line is a plain interest rate swap.
- It starts at zero, rises to a peak early, at about one third of the way through its life, and then falls back to zero at maturity.
- The grey line is an F X forward, a contract to exchange currencies on a future date.
- It has no payments in between, so nothing amortises, and its exposure keeps rising right up to maturity.
- Both lines start at zero because both contracts are worth nothing on day one.

13 The swap profile in a simple model

Exposure ∝ √t × (T − t)
Slope zero: T − t = 2t
Peak at t = T ÷ 3

Diffusion, root t	Uncertainty grows with root time
Amortisation, T - t	Remaining life shrinks to zero
Peak time	One third of life: first half

IN THE LESSON

- Why does the swap peak at one third of its life?
- In a simple model the swap's exposure is proportional to the square root of t, times the remaining life, T minus t.
- The square root of t is diffusion: uncertainty about rates grows with time.
- T minus t is amortisation: the number of payments still to come shrinks towards zero.
- Early on, the first factor grows faster than the second shrinks, so exposure rises.
- Later the reverse is true.
- Setting the slope to zero gives T minus t equal to two t, so t equals T divided by three.
- For a five year swap that is one and two thirds years, in the first half of its life, not the middle.

14 Example 5: a five year swap

Index = √t × (5 − t) for a 5-year swap (made-up units)

Year 1	√1 × 4	**4.00**
Year 1.67: the peak	√1.667 × 3.333	**4.30**
Year 3	√3 × 2	**3.46**
Year 4	√4 × 1	**2.00**
Year 5: maturity	√5 × 0	**0**

IN THE LESSON

- Let us put numbers on the five year swap, using the index square root of t times five minus t.
- At year one the index is one times four, which is four.
- At year one and two thirds it is the square root of one point six six seven times three point three three three, about four point three, which is the peak.
- At year three it is one point seven three two times two, about three point four six.
- At year four it is two times one, which is two.
- At year five it is zero.
- Notice the shape: a gentle rise, a peak in the first half, then a fall that quickens towards maturity.
- Exposure at year four is under half the peak, even though uncertainty is still growing.

15 Profiles of the main contracts

Each product has a characteristic shape

Contract	Exposure profile
Interest rate swap	Hump in first half, zero at end
FX forward	Keeps rising to maturity
Cross-currency swap	Rises even more: final exchange
Option	Buyer: its value. Seller: none
Bond or loan	Lender: about notional. Borrower: none

IN THE LESSON

- Here is the catalogue of profiles.
- A plain interest rate swap rises to a hump in the first half of its life, then falls to zero.
- An F X forward has no interim payments, so exposure keeps rising to maturity.
- A cross-currency swap rises even more, because the notional amounts are exchanged again at the end, which adds a large final payment.
- A bought option has exposure equal to its value.
- The seller has no exposure once the premium has been paid, because the buyer has nothing left to owe.
- A bond or loan gives the lender exposure of about the notional, and the borrower none.
- Which side of the contract you are on matters.

16 Combining trades

- A combination blends the shapes of its trades
- Swap plus FX forward: hump on a rising line
- Offsetting positions cancel
- Pay-fixed +3m and receive-fixed −3m: net exposure 0
- Without netting the same pair shows exposure of 3m

IN THE LESSON

- Real counterparties trade many contracts, and the profile of a combination blends the individual shapes.
- Add an F X forward to a swap and you get the swap's hump on top of a rising line.
- Offsetting positions cancel.
- Suppose Northgate holds a pay-fixed swap worth plus three million dollars and a receive-fixed swap with identical terms, with the same counterparty, worth minus three million.
- The net value is zero, and so is the exposure, provided netting is enforceable.
- Without netting, the gross exposure would still be three million, because the positive trade counts and the negative one is floored at zero.
- This leads naturally to the effects of payment timing, next, and then netting in full.

17 Payment frequency and the zero-coupon swap

- More frequent net payments: value settled sooner
- Quarterly rather than annual: lower exposure
- Zero-coupon swap: one payment, at maturity
- No interim amortisation: exposure rises until near maturity
- Higher than a coupon swap

IN THE LESSON

- Payment frequency and exercise dates both change the amortisation that shapes exposure.
- Start with payment frequency.
- When net payments are exchanged more often, say quarterly instead of annually, value is settled sooner, so less of it piles up before a default.
- Exposure is lower.
- The extreme case is a zero-coupon swap, where there is just one payment, at maturity.
- It has no interim amortisation at all, so its exposure keeps rising until close to maturity, and ends higher than that of a coupon-paying swap.
- The next example makes the quarterly and annual comparison concrete.

18 Example 6: annual against quarterly payments

Net payment accrues at $2m a year (simplified illustration)

Annual: worst unpaid	2m × 1 year	**$2.0m**
Quarterly: worst unpaid	2m × 0.25 year	**$0.5m**
Annual: average unpaid	half of the worst	**$1.0m**
Quarterly: average unpaid	half of the worst	**$0.25m**
Ratio	2.0 ÷ 0.5	**4 times**

IN THE LESSON

- Imagine the net payment owed accrues at two million dollars a year, and is paid at each payment date.
- With annual payments, just before a payment date a full year of accrual is unpaid: two million dollars.
- With quarterly payments the most that can be unpaid is a quarter of a year: two times one quarter, half a million.
- On average, the unpaid amount is half the worst case: one million for annual payments and a quarter of a million for quarterly.
- Quarterly payments leave four times less value unsettled, so there is less to lose if the counterparty fails.
- This is a simplified view, which ignores the market moves of the swap, but it shows why frequency matters.

19 Exercise dates, and what exercise does

Exposure follows the value of the exercise rights

Contract	Exposure
European option	Value of one exercise right
Bermudan option	At least the European's: more dates
Physically settled swaption	Becomes a swap: continues
Cash-settled swaption	Ends at the cash settlement

IN THE LESSON

- For options, exposure follows the value of the exercise rights.
- A European option can be exercised on one date only.
- A Bermudan option can be exercised on several dates, so it is worth at least as much as a European option with the same final expiry, and its exposure is at least as large.
- Swaptions, options to enter a swap, show another difference.
- If a swaption is physically settled, exercise turns it into a real swap, so exposure continues after the expiry date, following the value of the swap.
- If it is cash settled, a cash amount is paid at settlement, and exposure ends then.
- Always ask what happens at exercise.

20 Example 7: Bermudan, and the two swaptions

Bought options, same final expiry. Amounts in $ millions

European value	one exercise date	**$2.1m**
Bermudan value	extra exercise dates	**$2.6m**
Bermudan ≥ European?	2.6 ≥ 2.1	**Yes**
Physical swaption after expiry	swap worth 4.0	**$4.0m, moving**
Cash-settled, once paid	cash received	**$0**

IN THE LESSON

- Take two options bought by Northgate with the same final expiry.
- The European one is worth two point one million dollars.
- The Bermudan one, with extra exercise dates, is worth two point six million, which is more than the European, as it must be, because it contains the European right plus others.
- Since Northgate is the buyer, exposure equals value: two point six million for the Bermudan.
- Now suppose a swaption expires in the money, and the swap it creates is worth four million dollars.
- If physically settled, Northgate now holds that swap, with four million of exposure that keeps moving.
- If cash settled, it receives four million in cash, and exposure falls to zero once it is paid.

21 Netting: add first, then floor

Netted: $\max(V_1 + V_2 + \dots , 0)$
No netting: $\max(V_1, 0) + \max(V_2, 0) + \dots$

Netted exposure	Add the values, then floor at zero
Exposure without netting	Floor each trade, then add
Netting benefit	Never above the gross sum

IN THE LESSON

- Now aggregation.
- A netting agreement says that if the counterparty defaults, all trades in the netting set are added up into a single amount.
- Under a legally enforceable agreement, the values of the trades are added before the zero floor is applied.
- Exposure for the whole set is the larger of the sum of the values, and zero.
- Without netting, we floor each trade separately and then add.
- Because the floor is convex, the netted exposure is never greater than the sum of the stand-alone exposures.
- Positive and negative values partly cancel.
- How much they cancel depends on how the trades move together, which we come to in a moment.

22 Example 8: netting two trades

Values (V_1, V_2) in 4 equal scenarios ($m): (6, −2), (−3, 5), (4, 4), (−5, −6)

Without netting	6, 5, 8, 0	**avg 4.75**
With netting	max(sum, 0): 4, 2, 8, 0	**avg 3.5**
Netting benefit	4.75 − 3.5	**$1.25m**
As a share of gross	1.25 ÷ 4.75	**about 26%**

IN THE LESSON

- Four equally likely scenarios for two trades with one counterparty.
- In scenario one the values are six and minus two.
- Scenario two: minus three and five.
- Scenario three: four and four.
- Scenario four: minus five and minus six.
- Without netting, the exposures are six, five, eight and zero, and the average is four point seven five.
- With netting, the sums of the values are four, two, eight and minus eleven, so the exposures are four, two, eight and zero, averaging three and a half.
- Netting cut expected exposure by one point two five million dollars, about a quarter.
- The benefit came in the scenarios where one trade was positive and the other negative.

23 Correlation and the size of the benefit

EE ≈ 0.399 × σ (zero-mean normal)
σ(pair) = √($s_1^2 + s_2^2 + 2\rho\, s_1\, s_2$)

Expected exposure	About 0.399 times the SD
Std deviation, s	Spread of each trade's value
Correlation, rho	How the two trades move together
Portfolio standard deviation	Equals $s_1 + s_2$ only if ρ=1

IN THE LESSON

- Can we put a number on the benefit?
- Use the neat fact that for a zero-mean normal value, expected exposure is about zero point three nine nine times the standard deviation.
- For two trades, let s one and s two be their standard deviations and rho their correlation.
- The standard deviation of the combined value is the square root of s one squared, plus s two squared, plus two rho s one s two.
- Expected exposure of the netted pair is then zero point three nine nine times that.
- With perfect positive correlation, the portfolio standard deviation is s one plus s two, so there is no benefit.
- Zero correlation gives a partial benefit, and negative correlation the most.

24 Example 9: netting benefit and correlation

s_1 = $10m, s_2 = $6m. Stand-alone EE: 3.99 + 2.39 = 6.38

ρ = +1: σ = 16, no gain	0.399 × 16	**6.38**
ρ = 0: σ = √136 = 11.66	0.399 × 11.66	**4.65**
ρ = −0.5: σ = √76 = 8.72	0.399 × 8.72	**3.48**
ρ = −1: σ = √16 = 4	0.399 × 4	**1.60**

IN THE LESSON

- Take two trades with standard deviations of ten and six million dollars.
- Stand-alone expected exposures are zero point three nine nine times ten, which is three point nine nine, and zero point three nine nine times six, two point three nine four, together six point three eight.
- At a correlation of one, the combined standard deviation is sixteen, and expected exposure is six point three eight: no benefit.
- At zero correlation, the square root of one hundred thirty six is eleven point six six, so expected exposure is four point six five.
- At minus a half, the root of seventy six is eight point seven two, giving three point four eight.
- At minus one the combined standard deviation is four, and expected exposure one point six.
- The lower the correlation, the larger the saving.

25 Netting: three points to keep

- Correlation +1: no benefit. Zero: some. Negative: most
- Netted exposure ≤ sum of stand-alone exposures
- Without an enforceable agreement, exposures just add
- Check legal enforceability before any netting credit

IN THE LESSON

- Three points to keep.
- First, the benefit depends on correlation.
- At perfect positive correlation there is none, at zero correlation it is partial, and at negative correlation it is largest.
- Second, netted exposure is never greater than the sum of the stand-alone exposures, so netting can only help, when it applies.
- Third, and this is the exam trap, netting credit is only allowed when the agreement is legally enforceable.
- Without an enforceable agreement, the positive exposures simply add, trade by trade.
- So before applying any netting credit, check enforceability.
- A bank that assumes netting, and then finds a court will not honour it, has understated its exposure.

26 Credit exposure against funding exposure

Opposite sides of the same derivative

	Credit	Funding
Question	Loss if they fail	Cash we must find
When	Our MtM positive	Our MtM negative
Net of	Collateral we hold	Posted, less received
Kind of risk	Default loss	Liquidity and cost

IN THE LESSON

- Next, funding exposure.
- Credit exposure and funding exposure look at opposite sides of a derivative.
- Credit exposure is the loss if the counterparty defaults when the contract has positive value to us.
- It relies on the positive mark-to-market, net of the collateral we hold.
- Funding exposure is the cash or securities we must find, and pay for, when we have to post collateral or pay out.
- That is typically when our own mark-to-market is negative.
- The first is a default loss, a credit risk.
- The second is a liquidity and cost issue, and nobody has defaulted.
- Which one matters depends on the sign of the value, as the example shows.

27 Example 10: a hedged book that needs funding

Client trade has no collateral. Hedge is under a margin agreement. Funding cost 4%

Client trade value to us	uncollateralised	**+$6m**
Hedge value to us	exactly offsets	**−$6m**
Margin we post on hedge	no margin received	**$6m**
Funding need	mirrors client exposure	**$6m**
Yearly funding cost	4% × $6m	**$0.24m**

IN THE LESSON

- Now a funding need in a perfectly hedged book.
- Northgate has an uncollateralised trade with a client, Corvane.
- To hedge it, Northgate trades the opposite position with a dealer under a margin agreement.
- Suppose the client trade moves to be worth plus six million dollars to Northgate.
- The hedge is then worth minus six, so the book is flat.
- But Northgate must post six million of margin to the dealer, and receives nothing from the client, because the client trade is uncollateralised.
- It must find six million in funding.
- At a funding cost of four per cent, that costs two hundred and forty thousand dollars a year.
- The funding need mirrors the positive exposure of the client trade.
- It is a cost, not a default loss.

28 Margin period of risk, threshold and MTA

Exposure = max(V − C , 0)
Rough guide: threshold + MTA + move over the margin period

Threshold	Unsecured: no call below it
Minimum transfer amount	Smaller calls are not made
Margin period of risk	Remargin + call + close-out time
Value move in the period	Longer period, larger move
Exposure left	Initial margin cuts it further

IN THE LESSON

- Now collateral in more detail.
- With a margin agreement, close-out exposure is the larger of V minus C, and zero, where V is the value at close-out and C is the collateral actually held, which reflects the value at the last successful margin call.
- The gap between them arises over the margin period of risk.
- This is the remargining interval, plus the time to make the call, any dispute or grace period, and the time to close out and replace the position.
- It is often taken as about ten business days.
- Two terms leave exposure uncovered.
- The threshold is an unsecured amount, so no collateral is called below it.
- The minimum transfer amount, or M T A, means a call smaller than the M T A is not made.
- A rough guide to the exposure left is the threshold, plus the M T A, plus the value move over the margin period.

29 Example 11: sizing the collateralised exposure

Daily SD $0.8m, 99% (z = 2.326), threshold $2.0m, MTA $0.5m. Amounts in $ millions

10-day SD of value	0.8 × √10	**2.53**
99% move over 10 days	2.326 × 2.53	**5.88**
Exposure left, 10 days	2.0 + 0.5 + 5.88	**8.38**
Exposure left, 20 days	2.0 + 0.5 + 2.326 × 0.8 × √20	**10.82**

IN THE LESSON

- Let us size it for the Northgate netting set.
- Suppose the daily standard deviation of its value is zero point eight million dollars.
- Over a margin period of ten business days it is zero point eight times the square root of ten, which is two point five three million.
- At the ninety nine per cent level, multiply by two point three two six: five point eight eight million.
- Now the agreement terms.
- A threshold of two million and a minimum transfer amount of half a million add two and a half million.
- The rough guide to the exposure left is eight point three eight million.
- If the margin period doubled to twenty days, the move would be eight point three two million, and exposure ten point eight two million.
- Longer remargining periods raise exposure.

30 What each term does to the risk

- Longer remargining period: more unseen move
- Disputes and slow close-out lengthen the period
- Threshold: that much is uncollateralised from the start
- MTA: small gaps stay uncollateralised, value goes stale
- Initial margin held reduces exposure further

IN THE LESSON

- Judging the risk each term introduces.
- The remargining period: the longer between margin calls, the more value can move unseen, and the longer the margin period of risk, so exposure rises.
- Disputes and slow close-out lengthen it further.
- The threshold is a deliberate unsecured amount.
- Whatever the threshold is, that much exposure is simply uncollateralised, before any value move.
- The minimum transfer amount tolerates small gaps between value and collateral, so some stale uncollateralised value remains at every date.
- Initial margin held is the counterweight.
- It is collateral posted up front, in addition to the margin that tracks value, to cover exactly this gap, so it reduces exposure further.

31 Segregated or rehypothecated collateral

Who holds it decides who bears the risk

	Segregated	Rehypothecated
Held by	Custodian	Receiver, free to re-use
If receiver fails	Poster protected	Poster is unsecured
Receiver's funding	Cannot use it	Cheaper funding
Main cost	Receiver may charge	Poster's excess exposed

IN THE LESSON

- Last, how collateral is held decides who bears the risk.
- Segregated collateral is held by a custodian in the poster's name, kept apart from the receiver's own assets.
- If the receiver defaults, it is protected, so the poster has no exposure on it.
- The cost is that the receiver cannot use it, so it loses the funding benefit, and may charge for that.
- If collateral is not segregated and can be rehypothecated, meaning re-used by the receiver, the receiver gains cheaper funding.
- But the poster becomes an unsecured creditor for its return.
- So segregation trades lower counterparty risk for higher funding cost, and rehypothecation does the opposite.

32 Example 12: who loses when Corvane fails

Northgate owes $6m, posted $10m. Corvane defaults. Recovery on unsecured claims 40%

Excess collateral posted	10 − 6	**$4m**
Segregated: loss	4 back in full	**$0**
Rehypothecated: recovered	40% × 4	**$1.6m**
Rehypothecated: loss	60% × 4	**$2.4m**

IN THE LESSON

- Suppose Northgate owes Corvane six million dollars on its trades, and has posted ten million of collateral: six million of variation margin and four million of initial margin.
- Corvane defaults.
- If the collateral was segregated, six million settles what Northgate owes, the four million excess comes back in full, and it loses nothing.
- If it was rehypothecated, six million is used to settle what Northgate owes.
- The excess, four million, is an unsecured claim.
- At a recovery of forty per cent, Northgate gets back one point six million, and loses two point four million.
- So over-collateralisation, including initial margin, creates counterparty risk for the poster.
- That is the price of the cheaper funding that rehypothecation gave Corvane.

33 Recap: the main ideas

- Exposure = max(V, 0). EE is a mean, PFE is a quantile
- EPE averages EE. Max exposure = peak PFE
- Diffusion against amortisation shapes the profile
- Collateral: max(V − C, 0); margin period of risk remains
- Netting needs enforceability. Segregation protects poster

IN THE LESSON

- Let us pull it together.
- Exposure is the positive part of future value.
- At a date, E E is the mean of exposure, E N E the mean of the negative part, expected mark-to-market their sum, and P F E a high percentile.
- Over time, E P E averages E E, effective E P E uses a running maximum over the first year, and maximum exposure is the peak P F E.
- Exposure differs from value at risk by being projected to every date.
- Diffusion and amortisation shape the profiles.
- Collateral cuts exposure but the margin period of risk remains.
- Netting is never worse than no netting, if enforceable.
- Funding is a cost issue, and segregation protects the poster.

34 Try three quick questions

1. V equally likely −2, 0, 4, 10 ($m). EE and expected MtM?
   Answer: EE = 14 ÷ 4 = 3.5. ENE = −0.5, so MtM = 3.0.
2. Worth $9m at close-out, $5m collateral held. Exposure?
   Answer: max(9 − 5, 0) = $4m.
3. SDs $5m and $12m, correlation 0. Netted EE?
   Answer: σ = √169 = 13. EE ≈ 0.399 × 13 = $5.19m (gross 6.78).

IN THE LESSON

- Now try three quick questions.
- Pause the video, think, then listen to the answers.
- Question one.
- A future value is equally likely to be minus two, zero, four or ten million dollars.
- What are expected exposure and expected mark-to-market?
- Expected exposure is fourteen divided by four, three point five million.
- Expected negative exposure is minus half a million, so expected mark-to-market is three million.
- Question two.
- A trade is worth nine million dollars at close-out and five million of collateral is held.
- What is the exposure?
- The larger of nine minus five, and zero: four million.
- Question three.
- Two trades have standard deviations of five and twelve million and zero correlation.
- What is the netted expected exposure?
- The portfolio standard deviation is thirteen, so about five point one nine million, below the six point seven eight million without netting.