CHARTA LEARNING · LESSON NOTES

Capital Structure in Banks

FRM Part II

1 A story: the bank and the bad year

- Most years: a few loans fail, and pricing covers it
- Rare bad year: many fail together, losses jump
- Capital is the bank's own cushion against that year
- Question: is Harbourline's cushion big enough?

IN THE LESSON

- Harbourline Bank has made a few thousand loans.
- In a typical year some borrowers fail, and the bank charges interest high enough to cover them.
- That is an ordinary cost of business.
- But once in a long while a recession hits and many borrowers fail together.
- Losses jump far above normal, and if the bank's own money, its capital, is too thin, the bank itself fails.
- Management therefore asks a precise question: how much capital would we need to survive all but one year in a thousand?
- That amount is the bank's economic capital.

2 Is economic capital adequate?

EC = Loss at chosen quantile − EL
Adequate if capital held ≥ EC

Loss at the chosen quantile, Q	Exceeded in 0.1% of years
Expected loss, EL	The average loss, priced in
Economic capital, EC	Capital the bank must hold

IN THE LESSON

- Here is the test.
- Economic capital is adequate when the capital the bank actually holds covers its unexpected credit loss at the confidence level management has chosen.
- That level is usually tied to a target credit rating, for example ninety nine point nine per cent over one year, or ninety nine point nine seven.
- Read the formula like this.
- Take the loss that is exceeded only with the remaining small probability, here one year in a thousand.
- Subtract expected loss, the average loss, written E L.
- The result is the economic capital requirement, written E C.
- Then compare the capital available with that requirement.

3 Example 1: testing Harbourline

Harbourline, one year, 99.9% level. Amounts in $ millions

Loss at 99.9% level	from the loss model	**410**
Expected loss	priced and provisioned	**60**
Requirement	410 − 60	**350**
Capital held	equity only	**300**
Surplus or shortfall	300 − 350	**−50**

IN THE LESSON

- Let us test Harbourline.
- Its loss model says that in all but one year in a thousand, the one-year credit loss is below four hundred and ten million dollars.
- The expected loss is sixty million.
- So the requirement is four hundred and ten minus sixty, which is three hundred and fifty million dollars.
- Now look at what the bank holds.
- Its equity available to absorb credit losses is three hundred million.
- That is fifty million short of the requirement, so its economic capital is not adequate.
- The sixty million of expected loss is not part of the requirement: loan pricing pays for it, and provisions, reserves set aside for loan losses, cover it.

4 What moves the requirement, and two traps

- Higher confidence, concentration, correlation: EC rises
- Reserves ($60m) are not capital: that double counts EL
- That would show a false surplus: 360 − 350 = $10m
- Economic capital is the bank's own view
- Regulatory minimum ($240m) met: no proof of adequacy

IN THE LESSON

- Three things push the requirement up.
- Raising the confidence level reaches further into the tail.
- Adding concentration, meaning a few very large exposures, makes the tail heavier.
- Raising correlation makes defaults cluster, which does the same.
- Now the two traps.
- First, Harbourline's reserves of sixty million are not capital against unexpected loss.
- If the bank counted them, its holding would look like three hundred and sixty million, a false surplus of ten million, because expected loss would be counted twice.
- Second, economic capital is the bank's own view, and it can differ from the regulatory minimum.
- Suppose that minimum is two hundred and forty million.
- Meeting it does not prove adequacy.

5 The three inputs

LGD depends on seniority, collateral and workout costs, not on the obligor

Input	What it is	Common error
EAD	Amount owed at default	Using the limit
PD	Chance over one year	Mixed horizons
LGD	1 − recovery rate	Using recovery

IN THE LESSON

- Economic capital for credit risk is built from three inputs.
- The first is exposure, or exposure at default, written E A D.
- It is the amount owed at the moment of default.
- The second is the probability of default, written P D.
- It is the chance that the borrower defaults over the horizon, normally one year.
- The third is the loss rate, called loss given default, written L G D.
- It is the share of exposure lost after recoveries, so L G D equals one minus the recovery rate.
- It depends on seniority, collateral and workout costs, not on the quality of the borrower.
- The right hand column shows the usual error for each input.

6 Example 2: exposure and loss rate

Marlow Foods: limit $20m, drawn $12m, 40% of undrawn drawn, recovery 55%

Undrawn commitment	20 − 12	**$8m**
Exposure at default	12 + 40% × 8	**$15.2m**
Loss given default	1 − 55%	**45%**
Wrong: limit as exposure	ignores the drawn share	**$20m**
Wrong: recovery as LGD	55%, not 45%	**55%**

Amortising exposure: calculate period by period

IN THE LESSON

- Now numbers.
- Harbourline has given Marlow Foods a credit line with a limit of twenty million dollars.
- Marlow has drawn twelve million, so eight million is undrawn.
- Borrowers in trouble draw more, and the bank expects forty per cent of the undrawn part to be drawn before default.
- Exposure at default is therefore twelve million plus forty per cent of eight million, which is fifteen point two million dollars.
- The loan is senior and unsecured, and the bank expects to recover fifty five per cent.
- Loss given default is one minus the recovery rate, so forty five per cent.
- Two mistakes: using the twenty million limit as the exposure, and using the recovery rate where the loss rate is required.
- Where exposure amortises, a period by period calculation is needed.

7 Example 3: matching the horizon

Kestrel Retail: three-year cumulative PD of 27.1%. Find the annual PD

Three-year cumulative PD	from a rating study	**27.1%**
Three-year survival	1 − 27.1%	**72.9%**
Annual survival	72.9% to the power 1/3	**90%**
Annual PD	1 − 90%	**10%**
Wrong: divide by three	27.1% ÷ 3	**9.03%**

IN THE LESSON

- Probability of default is normally measured over one year, but rating studies often publish cumulative figures over several years.
- Suppose Harbourline's analysts quote a three year cumulative default probability of twenty seven point one per cent for Kestrel Retail, a weaker borrower.
- The chance of surviving three years is seventy two point nine per cent.
- Take the cube root to get annual survival of ninety per cent.
- So the annual default probability is ten per cent.
- The shortcut of dividing by three gives nine point zero three per cent, which is wrong, because survival compounds from year to year.
- All three inputs must share one horizon before they are combined.

8 Expected loss

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

Exposure at default	Amount owed at default
Probability of default	Chance of default in the year
Loss given default	Share lost after recoveries

IN THE LESSON

- With the three inputs in hand, we can compute expected loss, written E L.
- It is the average credit loss the bank anticipates over the horizon, the mean of the loss distribution.
- For one exposure it is exposure at default times probability of default times loss given default.
- For a portfolio, expected loss is simply the sum of the individual expected losses, whatever the correlations, because expectations always add.
- Expected loss is treated as a cost of doing business.
- It is built into the loan spread, which is the extra interest the bank charges, and it is covered by provisions or reserves.
- So it does not call for economic capital.

9 Example 4: expected loss of two loans

Marlow: EAD $15.2m, PD 2%, LGD 45%. Tidewater: EAD $6m, PD 1%, LGD 40%

Marlow EL	15.2m × 2% × 45%	**$136,800**
Tidewater EL	6m × 1% × 40%	**$24,000**
Two-loan EL	136,800 + 24,000	**$160,800**
Wrong: limit and recovery	20m × 2% × 55%	**$220,000**

IN THE LESSON

- Back to Harbourline.
- The Marlow Foods line has exposure of fifteen point two million, a default probability of two per cent and loss given default of forty five per cent.
- Expected loss is fifteen point two million times point oh two times point four five, which is one hundred and thirty six thousand eight hundred dollars.
- A second loan, to Tidewater Shipping, has exposure of six million dollars, a default probability of one per cent and loss given default of forty per cent.
- Its expected loss is twenty four thousand dollars.
- The two together have expected loss of one hundred and sixty thousand eight hundred.
- The common mistake, twenty million limit times two per cent times the fifty five per cent recovery rate, gives two hundred and twenty thousand, wrong on both counts.

10 Expected loss is a cost, not the risk

- EL = average loss: priced into spreads and provisions
- Portfolio EL = sum of ELs, whatever the correlations
- Trap: recovery rate used in place of LGD
- Trap: credit limit used in place of EAD
- EL is not the risk: the variability around it is

IN THE LESSON

- Three points to remember about expected loss.
- First, it is a cost, not the risk.
- A bank that loses exactly its expected amount every year is not in danger.
- Second, it adds across loans, so a portfolio's expected loss needs no correlation at all.
- Third, the typical mistakes are the two we saw: using the recovery rate for the loss rate, and multiplying by the credit limit instead of the exposure expected at default.
- The true risk is the variability around the average, the chance that losses are far above it.
- That variability is what the next outcomes measure, and it is what economic capital is for.

11 Unexpected loss: default is random

UL = EAD × LGD × √(PD × (1 − PD))

Exposure at default	Amount owed at default
Loss given default	Loss rate, taken as fixed
Probability of default	Chance of default in the year

IN THE LESSON

- Now the risk itself.
- Unexpected loss, written U L, is the standard deviation of credit losses around the expected loss.
- It measures how far losses may stray from the average.
- Start with the simplest case: exposure and loss given default are fixed, and only default is random.
- A loan either defaults, losing exposure times loss given default, or it does not, losing zero.
- The standard deviation of that all or nothing outcome is exposure times loss given default, times the square root of probability of default times one minus probability of default.
- The square root term is the standard deviation of a yes or no event, and we derive it in a few minutes.

12 Unexpected loss: the loss rate is random too

UL = EAD × √(PD × σLGD² + LGD² × PD × (1 − PD))

Exposure at default	Amount owed at default
Spread of the loss rate	Standard deviation of LGD
Loss given default	Average loss rate
Probability of default	Chance of default in the year

IN THE LESSON

- In practice the loss rate is not known in advance either.
- Let it have its own standard deviation, sigma L G D.
- Then unexpected loss is the exposure times the square root of two terms.
- The first term is the default probability times sigma L G D squared.
- It covers uncertainty about how much is lost.
- The second is loss given default squared, times default probability times one minus default probability.
- It covers uncertainty about whether default occurs.
- Both terms sit inside the root, so the variance of the loss rate adds inside the root, not outside it.
- If sigma L G D is zero, the formula collapses to the previous one.
- It treats default and the loss rate as independent.

13 Example 5: unexpected loss on Marlow

Marlow: EAD $15.2m, LGD 45%, PD 2%. Case 1: fixed LGD. Case 2: σLGD = 20%

Fixed LGD: EAD × LGD	15.2m × 45%	**$6.84m**
Fixed LGD: UL	6.84m × √0.0196	**$957,600**
Random LGD: loss-rate term	PD × σ² = 0.02 × 0.04	**0.0008**
Random LGD: default term	0.45² × 0.0196	**0.003969**
Random LGD: UL	15.2m×√(0.0008+0.003969)	**$1,049,700**

IN THE LESSON

- Let us measure the unexpected loss on the Marlow Foods line.
- Exposure times loss given default is six point eight four million.
- The square root of point oh two times point nine eight is point one four.
- With a fixed loss rate, unexpected loss is nine hundred and fifty seven thousand six hundred dollars.
- Now let the loss rate be random, with a standard deviation of twenty per cent.
- The first term is point oh two times point oh four, which is point zero zero zero eight.
- The second is point four five squared times point zero one nine six, which is point zero zero three nine six nine.
- Add them, take the root, and multiply by the exposure.
- Unexpected loss rises to about one million forty nine thousand seven hundred dollars.

14 Unexpected loss: what to remember

- UL = standard deviation of loss around EL: a spread
- Two terms: whether default occurs, and how much is lost
- The LGD variance adds inside the root, not outside it
- Errors: no root, no (1 − PD), recovery rate used as LGD
- Assumes default and the loss rate are independent

IN THE LESSON

- Keep a few warnings in mind.
- Unexpected loss is not expected loss, and it is not a loss amount that can be booked: it is a standard deviation, a measure of spread.
- The typical errors are leaving out the square root, leaving out the one minus probability of default factor, and using the recovery rate in place of loss given default.
- Another is forgetting that the loss rate variance belongs inside the root.
- The formula also assumes default and the loss rate are independent.
- And this standard deviation is not itself economic capital, which reads a high quantile of the loss distribution; we connect the two shortly.

15 Deriving the variance of default

One obligor, one horizon: D = 1 on default, D = 0 otherwise

Mean of D	1 × PD + 0 × (1 − PD)	**PD**
D squared equals D	$0^2 = 0$ and $1^2 = 1$	**$E[D^2] = PD$**
Variance	$E[D^2] - (E[D])^2$	**$PD - PD^2$**
Factor it	PD × (1 − PD)	**PD(1 − PD)**
Standard deviation	square root of the variance	**√(PD(1−PD))**

IN THE LESSON

- Why does the square root of P D times one minus P D appear?
- We can derive it.
- Treat one obligor's default over the horizon as a binomial, or Bernoulli, trial.
- Define an indicator D that equals one, with probability P D, if the obligor defaults, and zero otherwise.
- Its mean is one times P D plus zero times one minus P D, which is just P D.
- Now the key trick.
- Because D is either zero or one, D squared equals D.
- So the expected value of D squared is also P D.
- The variance is the expected value of D squared minus the squared mean: P D minus P D squared.
- Factor it, and the variance is P D times one minus P D.
- The standard deviation is the square root of that.

16 Many obligors: rate against count

Vrate = PD × (1 − PD) ÷ N
Vcount = N × PD × (1 − PD)
One obligor: SD is largest at PD = 50%

Probability of default	Chance one obligor defaults
Number of obligors, N	Independent, identical obligors
Variance of default rate	Shrinks as N grows
Variance of default count	Grows as N grows

IN THE LESSON

- Now extend this to many obligors.
- Take N independent, identical obligors.
- The default rate is the number of defaults divided by N.
- Its variance, V rate, is P D times one minus P D, divided by N.
- Diversification shrinks it: with more obligors, the rate becomes steadier.
- The number of defaults tells a different story.
- Its variance, V count, is N times P D times one minus P D, which grows with N.
- Both are correct.
- They measure different things, so do not confuse them.
- The single obligor standard deviation is largest at a P D of fifty per cent.
- Never confuse the variance with the P D itself, or with its square root.

17 Example 6: a hundred obligors, then four hundred

PD = 2% for every obligor, independent

One obligor, variance	0.02 × 0.98	**0.0196**
Rate variance, N = 100	0.0196 ÷ 100	**0.000196**
Rate SD, N = 100	√0.000196	**1.4%**
Count variance, N = 100	100 × 0.0196	**1.96**
N = 400: rate SD; count var.	√(0.0196 ÷ 400); 400 × 0.0196	**0.7%; 7.84**

One obligor: SD is largest at PD = 50% (variance 0.25, SD 0.5)

IN THE LESSON

- Take one hundred independent obligors, each with a default probability of two per cent.
- One obligor has variance point zero one nine six.
- The variance of the default rate is that divided by one hundred, and its standard deviation is one point four per cent.
- The variance of the default count is one hundred times point zero one nine six, which is one point nine six, so the count wobbles by about one point four defaults.
- Now quadruple the portfolio to four hundred.
- The standard deviation of the rate halves to point seven per cent, but the variance of the count quadruples to seven point eight four.
- Opposite effects, depending on whether we count in rates or in numbers.

18 Unexpected loss of a portfolio

$UL_i = EA_i \times LGD_i \times \sqrt{(\ PD_i \times (1 - PD_i)\ )}$
$UL_p = \sqrt{(\ \Sigma_i \Sigma_j\ \rho_{ij} \times UL_i \times UL_j\ )}$
Two assets: $\sqrt{(\ UL_1^2 + UL_2^2 + 2\rho\ UL_1\ UL_2\ )}$

Exposure of asset i	Amount lent to asset i
Standalone UL of asset i	Its own SD of loss
Correlation, rho ij	Pairwise default correlation
Portfolio UL	Below the sum unless all ρ = 1

IN THE LESSON

- Now to a portfolio.
- With the loss rate treated as fixed, each asset has its own standalone unexpected loss: its exposure, E A, times its loss given default, times the square root of its default probability times one minus that probability.
- Standalone figures cannot simply be added, because borrowers do not all fail together.
- Portfolio unexpected loss is the square root of the double sum, over every pair of assets, of the correlation times the two standalone figures.
- For two assets, this is the square root of U L one squared plus U L two squared plus two times rho times U L one times U L two.
- Portfolio U L is below the sum of the standalone figures unless every correlation equals one.
- That gap is the benefit of diversification.

19 Each asset's contribution

$ULC_i = UL_i \times (\ \Sigma_j\ \rho_{ij} \times UL_j\ ) \div UL_p$
$\Sigma_i\ ULC_i = UL_p$
Higher ρ raises ULC, not standalone UL

Standalone UL of asset i	Asset i's standalone UL
Link to the portfolio	Sum of correlation × UL, over j
Portfolio UL	Portfolio UL
Contribution of asset i	Adds up exactly to portfolio UL

IN THE LESSON

- Management also wants to know which assets use the capital.
- Asset i's contribution equals its standalone U L times its correlation with the portfolio.
- In symbols, U L C i equals U L i times the sum over j of rho i j times U L j, all divided by portfolio U L.
- The contributions add up exactly to the portfolio U L, so nothing is lost or counted twice.
- A larger exposure, a higher loss given default, or a higher default probability, as long as it is below fifty per cent, all raise the asset's standalone U L, and usually its contribution.
- Higher correlations raise only the contribution.
- Allocating in proportion to standalone U L is the common error, because it ignores correlation.

20 Example 7: three loans, one portfolio

Loans M, T, K. Correlations: M-T 0.3, M-K 0.2, T-K 0.1

UL of Marlow (M)	15.2m × 45% × √0.0196	**$957,600**
UL of Tidewater (T)	6m × 40% × √0.0099	**$238,797**
UL of Kestrel (K)	4m × 50% × √0.09	**$600,000**
Sum of standalone ULs	added, wrongly, as if ρ = 1	**$1,796,397**
Portfolio UL	√(1.334 + 0.396) trillion	**$1,315,182**

IN THE LESSON

- Add a third loan to Harbourline's book.
- Kestrel Retail has exposure of four million dollars, loss given default of fifty per cent and default probability of ten per cent, so its standalone unexpected loss is two million times the square root of point oh nine, which is six hundred thousand dollars.
- Tidewater Shipping's is about two hundred and thirty eight thousand eight hundred, and Marlow's is nine hundred and fifty seven thousand six hundred.
- The standalone figures sum to about one million eight hundred thousand.
- Using the correlations shown, the portfolio figure is the square root of the double sum.
- The squares add to about one point three trillion and the cross terms to about point four, so it comes to about one million three hundred and fifteen thousand dollars.
- Diversification has saved nearly half a million.

21 Example 8: sharing out the portfolio UL

$ thousands. ULC = UL × (Σ ρ × UL) ÷ P, P = portfolio UL

ULC of Marlow	957.6 × (957.6+71.6+120) ÷ P	**$836,775**
ULC of Tidewater	238.8 × (287.3+238.8+60) ÷ P	**$106,414**
ULC of Kestrel	600 × (191.5+23.9+600) ÷ P	**$371,994**
Sum of contributions	equals the portfolio UL	**$1,315,182**
Wrong: Kestrel by SD share	600 ÷ 1,796 × 1,315k	**$439,273**

IN THE LESSON

- Now share the portfolio figure out.
- Marlow's contribution is its standalone figure times its weighted link to all three loans, itself plus the correlated shares of the others, divided by the portfolio figure.
- That gives about eight hundred and thirty seven thousand dollars.
- Tidewater's contribution is only about one hundred and six thousand, and Kestrel's is about three hundred and seventy two thousand.
- They add up to one million three hundred and fifteen thousand, the portfolio figure, exactly.
- Notice that Tidewater has a standalone figure of two hundred and thirty nine thousand but contributes just one hundred and six thousand, because it diversifies the rest.
- Allocating by standalone share would charge Kestrel about four hundred and thirty nine thousand, too much by about sixty seven thousand.

22 Example 9: changing the mix

Same three loans. Change the size of T, then change every correlation

Base portfolio UL	three loans as before	**$1,315,182**
T doubled to $12m	UL of T becomes $477,594	**$1,437,580**
T's contribution	base, then doubled	**$106k to $274k**
All correlations = 1	the sum of standalone ULs	**$1,796,397**
All correlations = 0	√(sum of squared ULs)	**$1,154,999**

IN THE LESSON

- What if the mix changes?
- Suppose Harbourline doubles its lending to Tidewater, to twelve million.
- Tidewater's standalone figure doubles to about four hundred and seventy eight thousand.
- The portfolio figure rises only to about one million four hundred and thirty eight thousand, an increase of about one hundred and twenty two thousand, less than the standalone increase.
- Tidewater's own contribution climbs from about one hundred and six thousand to about two hundred and seventy four thousand.
- Now change the correlations instead.
- If every correlation were one, the portfolio figure would equal the plain sum, about one million eight hundred thousand.
- If every correlation were zero, it would fall to about one million one hundred and fifty five thousand.

23 Deriving economic capital

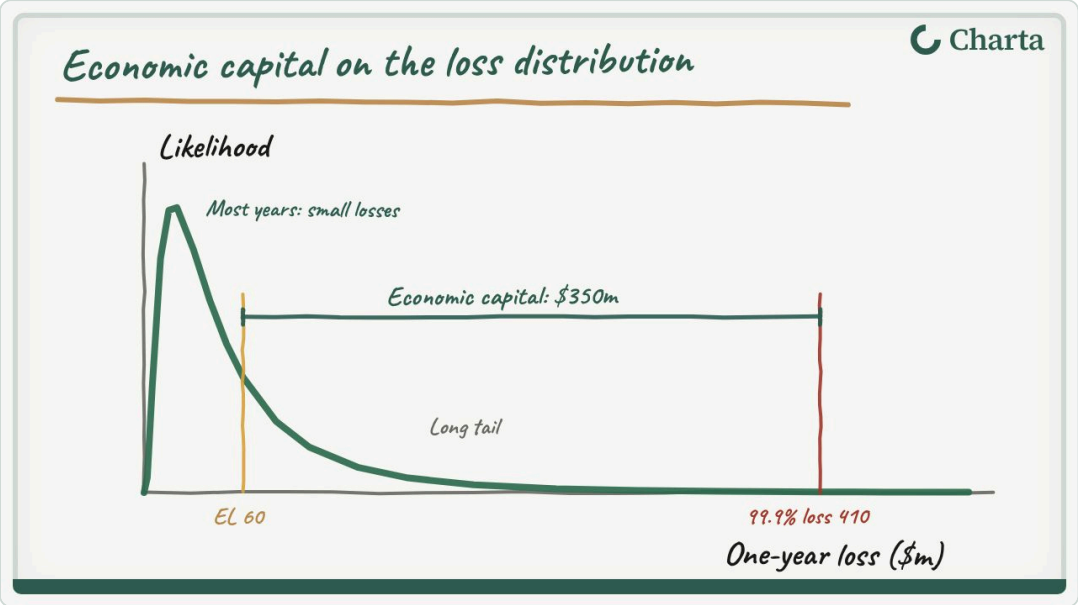
EC = Q − EL
EC = k × UL

Loss at the chosen quantile, Q	Loss exceeded in 0.1% of years
Expected loss, EL	Covered by pricing, provisions
Capital multiplier, k	About 3.09 if normal; often more
Unexpected loss, UL	Standard deviation of loss

IN THE LESSON

- How is economic capital derived?
- Start from the portfolio loss distribution, all the possible losses over one year and how likely each is.
- First pick a confidence level tied to the target rating, say ninety nine point nine per cent.
- Then read off the loss that is exceeded only with the remaining probability, one in a thousand.
- Call it Q.
- Finally subtract expected loss, because pricing and provisions cover that.
- Economic capital equals Q minus E L.
- Banks often quote it as a multiple of unexpected loss: E C equals k times U L, where k is the capital multiplier.
- If losses were normal, k at ninety nine point nine per cent would be about three point oh nine, from a one-sided normal quantile.
- But credit losses are skewed with a fat tail, so k is often substantially larger.

24 Economic capital on the loss distribution



IN THE LESSON

- Here is that derivation as a picture, not drawn to scale.
- Along the bottom is the one-year loss.
- The curve shows how likely each loss is.
- Most of the weight sits at small losses on the left, and a long tail stretches to the right.
- The first vertical line marks expected loss, the average, which sits to the right of the peak because the tail pulls the mean up.
- The second vertical line marks the ninety nine point nine per cent loss, with only one year in a thousand beyond it.
- The bracket between the two lines is economic capital.
- Capital covers the stretch from the average out to the chosen quantile.

25 Example 10: the multiplier and the allocation

Q = $410m, EL = $60m, UL = SD of loss = $80m. Amounts in $ millions

Requirement	410 − 60	**350**
Implied multiplier k	350 ÷ 80	**4.375**
Normal-table multiplier	99.9%, one-sided	**3.09**
Normal-based capital	3.09 × 80	**247.2**
Shortfall if used	350 − 247.2	**102.8**

Allocation: contribution 16 of 80 = one fifth, so 70 of the 350

IN THE LESSON

- Let us put numbers to the multiplier.
- Recall the requirement of three hundred and fifty million.
- Suppose Harbourline's portfolio unexpected loss, the standard deviation of loss, is eighty million.
- The implied multiplier is three hundred and fifty over eighty, which is four point three seven five.
- A normal table would give only three point oh nine, and so capital of two hundred and forty seven point two million, a shortfall of one hundred and two point eight million against the true figure.
- So the multiplier must come from the actual distribution.
- Finally, capital is allocated to assets through their U L contributions.
- An asset contributing sixteen million of the eighty gets one fifth, which is seventy million of the three hundred and fifty.

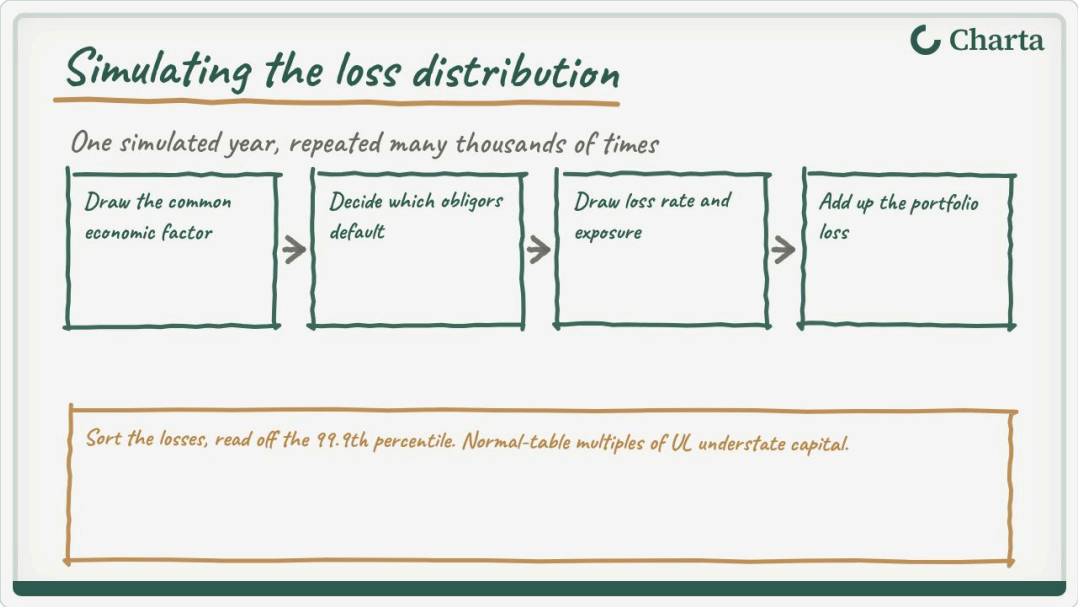
26 How credit losses are modelled

- Bounded at zero: many small losses, a long right tail
- Mean (EL) sits far below the tail quantiles of capital
- Three random pieces: default (PD), loss rate, exposure
- Loss rate often fitted by a beta distribution (0 to 1)
- Correlation and big single exposures fatten the tail

IN THE LESSON

- How are credit losses modelled?
- Not as a normal distribution.
- A bank cannot lose less than zero, so the distribution is bounded at zero.
- There is a high probability of small losses, and a long right tail from clustered defaults.
- As a result, the mean, expected loss, sits well below the tail quantiles that drive capital.
- The loss is built from three random pieces.
- Default is a binomial, or Bernoulli, event with probability P D.
- The loss rate is random, and since it lives between zero and one it is often fitted with a beta distribution.
- And exposure may vary as borrowers draw on their lines.
- Correlation among obligors, mainly through common economic factors, fattens the tail, and so does concentration in a few large exposures.

27 Simulating the loss distribution



IN THE LESSON

- The three pieces combine without a closed form, so banks use Monte Carlo simulation or analytic approximations.
- In each simulated year, draw the common economic factor, since it links the borrowers.
- Decide which obligors default, using their default probabilities and that factor.
- For each defaulter, draw a loss rate, perhaps from a beta distribution, and the exposure at default.
- Add up the losses to get one portfolio loss.
- Repeat many thousands of times.
- Then sort the simulated losses and read the required quantile directly, the ninety nine point ninth percentile for our target.
- The exam point is skewness: multiples of unexpected loss taken from a normal table understate capital.

28 Why credit risk is hard to quantify

- Defaults are rare and clustered: little data
- A 99.9% loss is a 1-in-1,000-year event: no validation
- Loans are illiquid, not marked to market: no price history
- Default correlations are small, unstable, hard to estimate
- PD and LGD rise together in downturns: averages mislead

IN THE LESSON

- Why is credit risk so hard to quantify?
- First, defaults are rare and they cluster in bad times, so there is little data.
- A ninety nine point nine per cent one-year loss is a one in a thousand year event, which no history can validate directly.
- Second, most loans are illiquid and are not marked to market, so there is no price history to learn from.
- Third, default correlations are small, unstable and hard to estimate, yet they drive the tail.
- And fourth, default probability and loss given default tend to rise together in downturns, so models that use averages understate the losses in exactly the years that matter.

29 More difficulties and their consequences

Each difficulty has a consequence for the capital number

Difficulty	Consequence
Skewed, fat-tailed losses	SD alone misleads
Uncertain parameters	Risk diversification cannot remove
Estimates with wide error	High model risk
99.9% estimate, 25 years of data	Backtesting nearly impossible

IN THE LESSON

- There are more difficulties.
- The loss distribution is skewed and fat tailed, so a standard deviation alone misleads about the tail.
- The parameters themselves are uncertain: a default probability or a correlation is an estimate with error.
- That adds systematic risk that diversification cannot remove, because adding more loans does not cure a wrong input.
- The result is high model risk.
- And backtesting a high confidence estimate is nearly impossible with realistic history.
- With twenty five years of data, a model that is badly wrong will probably show no breach at all, just as a correct one will.
- The table pairs each difficulty with its consequence.

30 Mistakes to avoid

- Counting reserves as capital: EL is already covered
- Limit used as EAD, or recovery rate used as LGD
- UL without the square root or the (1 − PD) factor
- Adding standalone ULs, or allocating by standalone share
- Normal multiples (3.09); regulatory minimum as proof

IN THE LESSON

- Here are the mistakes this topic invites.
- Counting reserves as capital: expected loss is already covered by pricing and provisions, so counting them double counts.
- Using the credit limit as exposure at default, or the recovery rate where the loss rate is required.
- In unexpected loss, dropping the square root or the one minus default probability factor.
- Adding standalone unexpected losses instead of combining them with correlations, or allocating capital by standalone share, which ignores correlation.
- Using a normal table multiplier such as three point oh nine for a skewed credit loss distribution.
- And treating the regulatory minimum as proof of adequacy.

31 Recap: the main ideas

- Adequate if capital held ≥ Q − EL; reserves are not capital
- Inputs: EAD, PD, LGD = 1 − recovery, on one horizon
- EL = EAD × PD × LGD (adds); UL = SD, √(PD(1 − PD))
- Portfolio UL uses correlations; contributions sum to it
- EC = Q − EL = k × UL, k from the skewed distribution

IN THE LESSON

- Let us pull it together.
- Economic capital is adequate when the capital held covers the loss at the chosen quantile minus expected loss.
- Its inputs are exposure at default, probability of default and loss given default, all on one horizon.
- Expected loss is exposure times probability times loss rate; it adds up and is priced.
- Unexpected loss is a standard deviation, built from the variance of a default, probability times one minus probability.
- Portfolio unexpected loss uses correlations, and the contributions add up to the total.
- Economic capital equals the quantile minus expected loss, or k times unexpected loss, with k taken from the real, skewed distribution, often simulated by Monte Carlo.
- And all of it is uncertain.

32 Try three quick questions

1. Limit $5m, $3m drawn, half of undrawn drawn at default, recovery 60%, PD 4%. EL?
   Answer: EAD $4m, LGD 40%: 4m × 4% × 40% = $64,000.
2. EAD $8m, fixed LGD 50%, PD 10%. UL?
   Answer: 4m × √(0.1 × 0.9) = 4m × 0.3 = $1.2m.
3. 99.9% loss $275m, EL $45m, capital held $220m. Adequate?
   Answer: Requirement 275 − 45 = $230m. No: $10m short.

IN THE LESSON

- Now try three quick questions.
- Pause, think, then listen.
- Question one.
- A credit line has a limit of five million dollars, three million drawn, and half of the undrawn part is expected to be drawn.
- Recovery is sixty per cent and the default probability is four per cent.
- Expected loss?
- Exposure is four million and loss given default is forty per cent, so expected loss is sixty four thousand dollars.
- Question two.
- A loan has exposure of eight million dollars, a fixed loss given default of fifty per cent and a default probability of ten per cent.
- Unexpected loss?
- Four million times the square root of point oh nine is one point two million dollars.
- Question three.
- A bank's ninety nine point nine per cent loss is two hundred and seventy five million, expected loss is forty five million, and it holds two hundred and twenty million of capital.
- Adequate?
- The requirement is two hundred and thirty million, so no, ten million short.