CHARTA LEARNING · LESSON NOTES

Credit Risk Modelling and Assessment

FRM Part II

1 Is this bank sound? The CAMEL system

- C Capital adequacy: can capital absorb losses?
- A Asset quality: how healthy is the loan book?
- M Management: governance, controls, strategy
- E Earnings: does it make money, and for how long?
- L Liquidity: can it pay when funds fall due?

IN THE LESSON

- Imagine a bank examiner arriving at Harbourline Bank.
- She cannot check every loan, so she uses a framework that supervisors have relied on for decades, called the CAMEL system.
- It rates a bank on five dimensions, one for each letter.
- C is capital adequacy: can capital absorb losses?
- A is asset quality: how healthy is the loan book?
- M is management: the people, their controls and their strategy.
- E is earnings: does the bank make money, and will it keep doing so?
- And L is liquidity: can it pay what it owes when funds fall due?
- The letters spell a word, but they carry no ranking.
- A weakness on any one of them is a warning sign.

2 What each letter looks at

Match every indicator to its letter

Letter	Typical indicators
Capital	Tier 1 and total capital ratios
Asset quality	NPLs ÷ loans, reserves ÷ NPLs, concentration
Management	Governance, risk controls, strategy
Earnings	ROA, net interest margin, cost-to-income
Liquidity	Loans ÷ deposits, funding, maturity gaps

IN THE LESSON

- Each letter has its own typical indicators, and matching them is exactly what exam items test.
- Under capital adequacy we look at capital relative to risk-weighted assets, using the tier one and total capital ratios.
- Under asset quality we look at the loan book: non-performing loans as a share of gross loans, reserves held against those non-performing loans, and concentrations in a single borrower or sector.
- Management is the most judgmental element, covering governance, risk controls and strategy, because there is no single ratio for it.
- Earnings are judged through return on assets, net interest margin and the cost-to-income ratio.
- Liquidity covers funding stability, loans to deposits, and gaps between when assets and liabilities mature.

3 Example 1: scoring Harbourline

$m: capital 150, assets 2,000, loans 1,200, deposits 1,000, NPLs 24, NI 22

Capital: total ratio	150 ÷ RWA (Ex. 6)	**13.0%**
Asset quality: NPL ratio	24 ÷ 1,200	**2.0%**
Earnings: ROA (1.1% as is)	(22 − 4) ÷ 2,000	**0.9%**
Management	governance review	**judgmental**
Liquidity: loans ÷ deposits	1,200 ÷ 1,000	**120%**

Strong on four letters can still fail on the fifth

IN THE LESSON

- Here is Harbourline in numbers, in millions of dollars.
- Capital is one hundred and fifty, and later we will show the total capital ratio is thirteen per cent.
- Non-performing loans are twenty four on gross loans of twelve hundred, a ratio of two per cent.
- Net income is twenty two on assets of two thousand, a return on assets of one point one per cent, but four of that was a one-off gain from selling a building.
- Strip it out, because sustainability matters as much as level, and the underlying return is zero point nine per cent.
- Management gets a judgmental review rather than a ratio.
- Now liquidity.
- Loans are twelve hundred but deposits only a thousand, so loans to deposits is one hundred and twenty per cent: the bank leans on other funding.
- Four letters look fine, the fifth is weak.

4 Credit risk: three inputs and a horizon

- EAD: amount owed at default, with expected draw
- LGD: share of EAD not recovered = 1 − recovery
- PD: chance of default over the horizon
- Horizon: usually one year, the same for all three
- Expected loss is priced; unexpected loss needs capital

IN THE LESSON

- Now to the heart of credit risk: putting a number on it.
- Harbourline has lent to Brightwater Foods, a made-up food distributor, through a credit line.
- Three questions decide the loss if things go wrong.
- First, how much will be owed at the moment of default?
- That is exposure at default, E A D.
- For a credit line it includes the part of the undrawn limit we expect the borrower to draw as trouble nears.
- Second, what share is not recovered?
- That is loss given default, L G D, equal to one minus the recovery rate.
- Third, how likely is default?
- That is the probability of default, P D.
- All three must refer to the same time horizon, usually one year, and that horizon must match the horizon of the loss we want to measure.

5 Expected loss and unexpected loss

EL = PD × LGD × EAD
UL = EAD × LGD × √(PD × (1 − PD))

Probability of default	Chance of default in the horizon
Loss given default	Share not recovered
Exposure at default	Amount owed at default
Expected loss	Average loss: priced, provisioned
Unexpected loss	Spread around EL: held in capital

IN THE LESSON

- Put the three together and you get expected loss, the average loss a bank should expect over the horizon.
- It is probability of default times loss given default times exposure at default.
- This is the loss that Harbourline prices into its interest rate and provisions for, because it is a normal cost of lending.
- But averages hide risk.
- Unexpected loss is the volatility of loss around that average, and it is what capital is held to cover.
- For one exposure with a fixed loss given default there are only two outcomes, default or no default.
- Then unexpected loss is exposure times loss given default times the square root of P D times one minus P D.
- Remember that this version is for a single exposure, not for a whole portfolio.

6 Example 2: expected loss on Brightwater's line

Limit $5m, $3m drawn. 60% of undrawn expected drawn. Recovery 55%, PD 1.5%

Expected draw	60% × 2m	**$1.2m**
Exposure at default	3m + 1.2m	**$4.2m**
Loss given default	1 − 55%	**45%**
Expected loss	1.5% × 45% × 4.2m	**$28,350**
Wrong: recovery for LGD	1.5% × 55% × 4.2m	**$34,650**

IN THE LESSON

- Let us compute expected loss for Brightwater.
- The credit line has a limit of five million dollars, of which three million is drawn and two million is undrawn.
- Harbourline expects sixty per cent of the undrawn part to be drawn by the time of default, which is one point two million.
- Exposure at default is therefore three plus one point two, or four point two million dollars, not the three million drawn and not the five million limit.
- Recovery is expected to be fifty five per cent, so loss given default is forty five per cent.
- With a default probability of one point five per cent, expected loss is point zero one five times point four five times four point two million, which is twenty eight thousand three hundred and fifty dollars.
- The classic slip is to use the recovery rate in place of loss given default, which would give thirty four thousand six hundred and fifty.

7 Example 3: unexpected loss on the same line

Same line: EAD $4.2m, LGD 45%, PD 1.5%

Loss if default	4.2m × 45%	**$1.89m**
Spread of a yes/no event	√(0.015 × 0.985)	**0.12155**
Unexpected loss	1.89m × 0.12155	**$229,734**
UL compared with EL	229,734 ÷ 28,350	**8.1 times**

Pricing covers EL; capital covers UL

IN THE LESSON

- Now the unexpected loss on the same line.
- If Brightwater defaults, Harbourline loses exposure times loss given default, which is four point two million times forty five per cent, or one point eight nine million dollars.
- The yes or no event has a standard deviation of the square root of point zero one five times point nine eight five, which is the square root of point zero one four seven seven five, about point one two one five five.
- Multiply, and unexpected loss is about two hundred and twenty nine thousand seven hundred dollars.
- That is over eight times the expected loss.
- This is why the two are treated differently.
- The small expected loss is a routine cost that pricing recovers.
- The much larger unexpected loss is the uncertainty that the bank's own capital has to be ready to absorb.

8 Example 4: the horizon and cumulative PD

Constant annual PD p = 1.5%. Cumulative PD over T years

Wrong: T × p	5 × 1.5%	**7.5%**
Survive five years	0.985^5	**92.72%**
CPD = 1 − (1 − p)^T	1 − 0.9272	**7.28%**
Three-year CPD	1 − 0.985^3	**4.43%**

Match PD, EAD and LGD to the same horizon

IN THE LESSON

- What if we want default over more than one year?
- Suppose Brightwater has a constant annual default probability of one point five per cent.
- A tempting shortcut multiplies by the number of years, giving seven point five per cent over five years.
- That is wrong, because a firm can only default once, and it must first survive each earlier year.
- The chance of surviving one year is one minus p.
- Surviving five years in a row is that figure to the power five, which is point nine two seven two.
- The cumulative probability of default is one minus that, or seven point two eight per cent.
- Over three years it is four point four three per cent, again below three times one point five.
- Whatever horizon you choose, the probability of default, the exposure and the loss given default must all refer to it.

9 Mistakes to avoid with credit risk inputs

- A one-year PD needs a one-year EAD and LGD
- Cumulative PD is 1 − (1 − p)^T, not T × p
- LGD is 1 − recovery, not the recovery rate
- EAD includes the expected draw on undrawn limits
- EL is priced; UL (the volatility) is held in capital

IN THE LESSON

- Let us gather the traps.
- First, horizon mismatch: a one-year default probability cannot be combined with an exposure or a loss measured over some other period.
- Second, cumulative default probability is not T times p; it is one minus the survival probability compounded.
- Third, loss given default is one minus the recovery rate, so a recovery of forty per cent means a loss given default of sixty.
- Fourth, for a credit line, exposure at default includes the expected draw on the undrawn part, not just what is drawn today.
- Fifth, keep the roles apart.
- Expected loss is the average, which pricing and provisions cover.
- Unexpected loss is the volatility around it, and capital covers that.
- Exam questions often hide one of these slips inside a calculation.

10 The capital adequacy ratio

CAR = Capital ÷ RWA
RWA = exposures × risk weights + 12.5 × charges

Regulatory capital	Regulatory capital held
Risk-weighted assets	Exposures scaled by credit risk
Supervisory minimum	E.g. 8% total, plus buffers
Off-balance-sheet items	Convert with a CCF, then weight

IN THE LESSON

- Back to Harbourline's own health.
- Its capital is one hundred and fifty million dollars against two thousand million of assets, but a dollar of cash and a dollar lent to a struggling company are not equally risky.
- So regulators scale each exposure by a risk weight and add the results into risk-weighted assets.
- The capital adequacy ratio is regulatory capital divided by risk-weighted assets.
- Supervisors set a minimum; the Basel framework gives eight per cent for total capital as an example, plus buffers.
- Guarantees and other off-balance-sheet items are first turned into an on-balance-sheet equivalent using a credit conversion factor, then weighted.
- Market and operational risk capital charges join in too, multiplied by twelve point five, the reciprocal of eight per cent.
- Note that the denominator is risk-weighted assets, not total assets.

11 Illustrative risk weights

Illustrative weights, as used in this lesson

Exposure	Risk weight
Cash and sovereign claims	0%
Interbank exposures	20%
Residential mortgages	35%
Corporate loans	100%
Guarantee (off balance sheet)	CCF first, then weight

IN THE LESSON

- For illustration, here are the risk weights we will use.
- Cash and claims on highly rated sovereigns carry zero per cent, because we treat them as essentially safe.
- Exposures to other banks carry twenty per cent.
- Residential mortgages carry thirty five per cent, since they are backed by property.
- Corporate loans carry one hundred per cent, the full weight.
- These are illustrative; real weights depend on the rules and ratings in force.
- A guarantee that Harbourline has given on behalf of a client is not on the balance sheet, so we first multiply its face amount by a credit conversion factor, and then by the risk weight of the party it stands behind.
- Do not skip the conversion step: a question often gives the face amount and expects you to apply both.

12 Example 5: Harbourline's risk-weighted assets

$m: cash/sovereign 600, interbank 200, mortgages 400, corporate 800; guarantee 100

Cash and sovereign	600 × 0%	**0**
Interbank	200 × 20%	**40**
Mortgages	400 × 35%	**140**
Corporate loans	800 × 100%	**800**
Guarantee, CCF 50%	100 × 50% × 100%	**50**

On-balance RWA 980 (assets 2,000) + 50 = credit RWA 1,030

IN THE LESSON

- Here is Harbourline's balance sheet, in millions of dollars.
- Cash and sovereign claims are six hundred, at zero per cent, so they add nothing.
- Interbank exposures of two hundred at twenty per cent add forty.
- Residential mortgages of four hundred at thirty five per cent add one hundred and forty.
- Corporate loans of eight hundred at one hundred per cent add eight hundred.
- Total assets are two thousand, but risk-weighted assets on the balance sheet are only nine hundred and eighty.
- Then there is a guarantee of one hundred given for a corporate client.
- With a conversion factor of fifty per cent, its on-balance-sheet equivalent is fifty, and at a one hundred per cent weight it adds fifty.
- Credit risk-weighted assets are therefore one thousand and thirty.

13 Example 6: the ratio itself

Capital $150m. Market charge 4, operational charge 6 ($m). Minimum 8%

Charges turned into RWA	(4 + 6) × 12.5	**125**
Total RWA	1,030 + 125	**1,155**
Capital adequacy ratio	150 ÷ 1,155	**13.0%**
Capital needed at 8%	8% × 1,155	**92.4**
Surplus over the minimum	150 − 92.4	**57.6**

IN THE LESSON

- Now add market and operational risk.
- Harbourline's market risk capital charge is four million and its operational risk charge is six million, ten in all.
- To turn a capital charge into risk-weighted assets, multiply by twelve point five.
- That gives one hundred and twenty five.
- Total risk-weighted assets are one thousand and thirty plus one hundred and twenty five, or one thousand one hundred and fifty five.
- With capital of one hundred and fifty, the capital adequacy ratio is one hundred and fifty divided by one thousand one hundred and fifty five, which is thirteen point zero per cent.
- At the eight per cent minimum, Harbourline needs ninety two point four, so it has a surplus of fifty seven point six million dollars.
- Notice that the ratio is judged against risk, not against the size of the balance sheet.

14 Example 7: a write-off, and the wrong ratio

Harbourline writes off a $30m corporate loan (no reserve) and removes it

Capital after the loss	150 − 30	**120**
RWA after removal	1,155 − 30	**1,125**
CAR after write-off	120 ÷ 1,125	**10.7%**
If RWA had not fallen	120 ÷ 1,155	**10.4%**
Wrong: ÷ total assets	150 ÷ 2,000	**7.5%**

Capital ÷ total assets is a leverage-type ratio, not CAR

IN THE LESSON

- Now a bad event.
- Harbourline writes off a thirty million dollar corporate loan, with no reserve held, and removes it from the books.
- Capital falls by the loss, from one hundred and fifty to one hundred and twenty.
- Because the asset has left the balance sheet, risk-weighted assets fall too, by thirty at a one hundred per cent weight, to one thousand one hundred and twenty five.
- The new ratio is one hundred and twenty over one thousand one hundred and twenty five, ten point seven per cent.
- Had risk-weighted assets stayed the same, it would be ten point four per cent, so removing the asset softens the blow a little but does not cancel it.
- And one wrong ratio is tempting: capital divided by total assets is seven point five per cent.
- That is a leverage-type ratio, a different measure, not the capital adequacy ratio.

15 Three families of default prediction

- Judgmental: credit officers, expert systems, five C's
- Empirical: statistics fitted to defaulters and survivors
- Financial (structural): asset value, volatility, debt
- Each suits different data and has its own weakness
- Exam: match approach to data, name the limitation

IN THE LESSON

- A bank must also predict which borrowers will fail.
- There are three families of approach.
- Judgmental approaches rely on the experience of credit officers and expert systems, for example the five C's of credit: character, capacity, capital, collateral and conditions.
- They work with small samples and with private firms, but they are subjective, and two analysts can reach different conclusions.
- Empirical models are statistical: they are fitted to historical data on defaulting and surviving firms, using accounting ratios as inputs.
- Financial, or structural, models derive default from the firm's own asset value, volatility and liabilities, using market data.
- Think of them as three desks at Harbourline, each suited to different data.
- An exam question usually asks you to match the approach to the data available, and to name its main limitation.

16 Strengths and limits of each family

What each family needs, and where it falls short

Family	Needs	Main limitation
Judgmental	Experience	Subjective
Empirical	Past defaults	Backward looking
Financial	Traded equity	Public firms only

IN THE LESSON

- Here is the comparison.
- Judgmental methods need only experience, and can cope with a small sample or a private firm, but they are subjective and can differ between analysts.
- Empirical models, such as discriminant analysis like the Altman Z score, or logit and probit models, are objective and use accounting ratios.
- But they look backward, and they are only as good as the sample they were built on.
- A model always fits its own sample better than new data, so in-sample fit overstates out-of-sample accuracy.
- Financial models such as Merton use market data, but they need traded equity, so they cannot be applied to a private firm.
- Each strength is the mirror image of another approach's weakness, which is why a bank often uses more than one.

17 Example 8: matching approach to data

Which approach fits each case?

Private firm, thin records	no share price	**Judgmental**
5,000 past loans with ratios	defaulters and survivors	**Empirical**
Listed firm, liquid shares	asset value from equity	**Financial**
Fits its own sample well	in-sample fit flatters	**Test out-of-sample**

IN THE LESSON

- Let us practise matching.
- Harbourline's first case is a small private family firm with thin records and no share price.
- There is no statistical sample and no market data, so the answer is a judgmental approach, with its subjectivity as the limit.
- The second case is a book of five thousand past loans with known accounting ratios and a record of which ones failed.
- That suits an empirical model, such as logit, and the limit is that it looks backward.
- The third is a listed company with liquid shares.
- Equity prices let us infer asset value and volatility, so a structural model fits, and the requirement is traded equity.
- Finally, suppose a model fits its own training data very well.
- Do not take that as proof of accuracy, because in-sample fit overstates out-of-sample accuracy.

18 Merton: default as an option on assets

$d2 = [\ \ln(V \div D) + (r - \sigma^2/2) \times T\ ] \div (\ \sigma \times \sqrt{T}\ )$
Default probability = $N(-d2)$
Real-world version: replace r with expected return μ

Asset value, V	Market value of the firm's assets
Debt face value, D	Zero-coupon debt due at T
Asset volatility, σ	Volatility of asset value
Risk-free rate, r	Risk-free rate
Distance to default	d2: SDs above the default point

IN THE LESSON

- Now the classic structural model, Merton.
- It imagines Brightwater financed by equity and one zero-coupon debt of face value D, due at time T.
- At T, shareholders keep whatever is left after the debt is repaid, so equity is a European call option on the firm's assets V with strike D.
- Default happens only if assets are below D at T.
- With asset volatility sigma and risk-free rate r, the standardised term d two is the log of V over D, plus r minus half sigma squared times T, all divided by sigma times the square root of T.
- The risk-neutral default probability is N of minus d two.
- Distance to default is d two: how many standard deviations the log asset value sits above the default point.
- The real-world version replaces r with the asset's expected return, mu.

19 Example 9: Merton for Brightwater

V = 150, D = 100, σ = 25%, r = 3%, T = 1 year ($m)

ln(V ÷ D)	ln 1.5	**0.4055**
Drift term	0.03 − 0.25²/2	**−0.00125**
d2 = distance to default	(0.4055 − 0.00125) ÷ 0.25	**1.617**
Default probability	N(−1.617)	**5.3%**
Trap: N(−d1) instead	d1 = 1.867	**3.1%**

IN THE LESSON

- Brightwater has assets worth one hundred and fifty million dollars and a single zero-coupon debt with face value one hundred million, due in one year.
- Asset volatility is twenty five per cent and the risk-free rate is three per cent.
- The log of one hundred and fifty over one hundred is point four zero five five.
- The drift term is point zero three minus half of point two five squared, which is minus point zero zero one two five.
- Add them to get point four zero four two, and divide by point two five: d two is one point six one seven.
- The default probability is N of minus one point six one seven, about five point three per cent.
- The common slip is to use d one, which is d two plus sigma root T, one point eight six seven.
- That gives three point one per cent and understates the risk.

20 Example 10: what higher volatility does

Same firm: asset volatility rises from 25% to 35%

Distance to default	1.617 → 1.069	**falls**
Default probability	5.3% → 14.2%	**rises**
Equity (the call)	$53.5m → $55.1m	**rises**
Risky debt = V − call	$96.5m → $94.9m	**falls**
Credit spread	0.51% → 2.23%	**widens**

Call = V·N(d1) − D·e^(−rT)·N(d2); spread = ln(D ÷ debt) − r

IN THE LESSON

- Now suppose Brightwater's business becomes riskier, with asset volatility rising from twenty five to thirty five per cent.
- Distance to default shrinks from one point six one seven to one point zero six nine, and the default probability rises from five point three to fourteen point two per cent.
- By the option formula, equity becomes more valuable, from fifty three point five to fifty five point one million dollars, because shareholders own a call option: they gain from the upside while the downside is capped at zero.
- Risky debt equals asset value minus the call, so it falls from ninety six point five to ninety four point nine.
- A lower debt value for the same promised payment means a higher yield, and the credit spread, the yield less the risk-free rate, widens from about half a per cent to two point two per cent.
- Higher volatility helps equity and hurts lenders.

21 Where the Merton model falls short

- Default only at maturity, never before
- One zero-coupon debt; real firms have many
- Constant volatility and normal asset returns
- Asset value and volatility are inferred, not seen
- PD is N(−d2), not N(−d1); more vol, less distance

IN THE LESSON

- Merton is elegant, but it has well-known shortcomings.
- Default can happen only at maturity, so a firm cannot fail earlier, though real firms default whenever they miss a payment.
- There is a single zero-coupon liability, whereas real firms have many debts with different maturities and coupons.
- Volatility is assumed constant and asset returns normal, but real returns have fat tails.
- And asset value and asset volatility cannot be observed.
- They must be inferred from equity prices, which adds estimation error.
- Finally, keep the formula straight: the default probability is N of minus d two, not N of minus d one.
- And remember that higher asset volatility raises equity value but lowers distance to default.

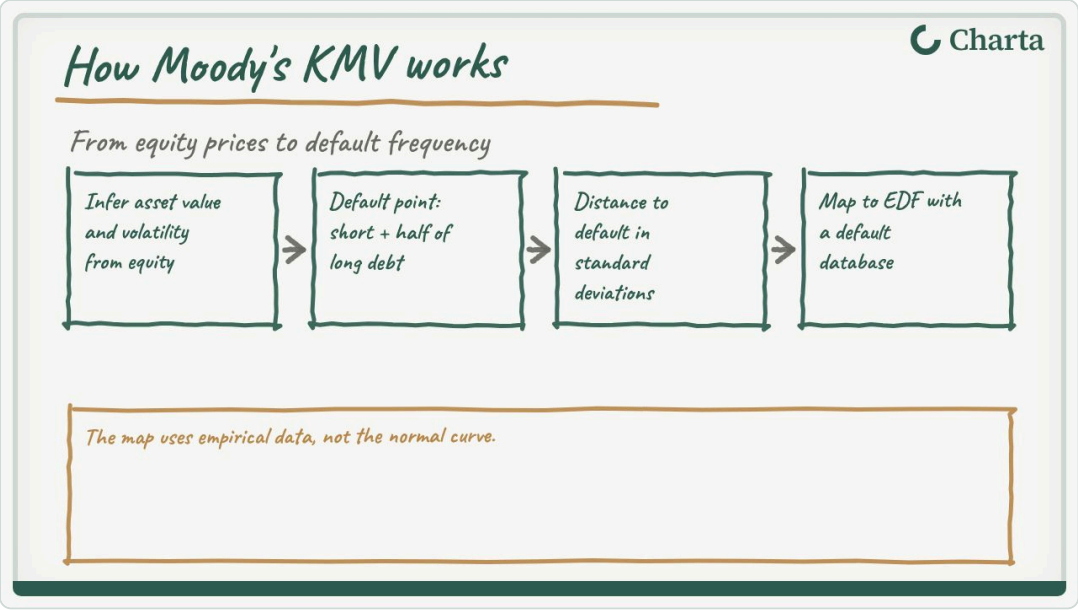
22 Four models, four questions

- Merton: equity is a call; default if assets < debt at T
- Moody's KMV: Merton idea plus an empirical default map
- CreditMetrics: portfolio model of rating migration
- CreditRisk+: actuarial, defaults as a Poisson count
- Merton, KMV, CreditMetrics are asset-value driven

IN THE LESSON

- Four named models answer different questions, so exam items often ask you to tell them apart.
- Merton treats equity as a call option on firm assets, and defaults when assets fall below debt at maturity.
- Moody's K M V builds on it.
- CreditMetrics is a portfolio model based on rating migration: it asks how the value of each loan changes if the borrower is upgraded, downgraded or defaults.
- CreditRisk plus is actuarial: it models the number of defaults as a Poisson count whose mean rate is itself uncertain, and it says nothing about why firms default.
- The first three are driven by asset values, and the fourth is not.
- Hold that split in mind as we look at each in turn.

23 How Moody's KMV works



IN THE LESSON

- Moody's K M V starts where Merton stops.
- Follow the boxes.
- First, it infers the firm's asset value and asset volatility from its equity price and volatility, because assets are not traded.
- Second, it sets a default point: short-term liabilities plus half of long-term liabilities.
- The idea is that a firm need not fail the moment assets dip below total debt, since long-term debt is not yet due.
- Third, it computes distance to default, the gap between asset value and the default point measured in standard deviations of asset value.
- Fourth, instead of using the normal curve, it maps distance to default to an expected default frequency, using a large empirical database of defaults.
- So K M V mixes a structural idea with empirical calibration.

24 Example 11: KMV distance to default

V = 150, short-term debt 40, long-term debt 100, σ = 25% ($m)

Default point	40 + 0.5 × 100	**90**
Distance in dollars	150 − 90	**60**
One SD of asset value	25% × 150	**37.5**
Distance to default	60 ÷ 37.5	**1.6 SDs**
Then EDF	from default database	**not N(−DD)**

IN THE LESSON

- Apply it to Brightwater, now with different liabilities.
- Asset value is one hundred and fifty million dollars, short-term liabilities are forty, long-term liabilities are one hundred, and asset volatility is twenty five per cent.
- The default point is forty plus half of one hundred, which is ninety.
- The distance in dollars is one hundred and fifty minus ninety, sixty million.
- One standard deviation of asset value is twenty five per cent of one hundred and fifty, thirty seven point five million.
- So distance to default is sixty divided by thirty seven point five, one point six standard deviations.
- Now the key step: K M V does not read the probability from the normal curve.
- It looks at firms in its database that once had a similar distance to default, and records the share that actually defaulted.
- That share is the expected default frequency.

25 The four models side by side

Similar goal, different engines

Model	Asset values?	Core idea
Merton	Yes	Equity is a call on assets
Moody's KMV	Yes	DD mapped to EDF
CreditMetrics	Yes	Rating migration, correlations
CreditRisk+	No	Poisson default counts

IN THE LESSON

- Now put the four side by side.
- Merton, K M V and CreditMetrics are all asset-value driven.
- Merton and K M V look at a firm's assets against its debt, and K M V adds the default point and the empirical map.
- CreditMetrics is the portfolio model: it revalues each loan under every possible end-of-horizon rating, including default, using transition probabilities, and it links borrowers through correlations of asset or equity returns.
- CreditRisk plus is the odd one out.
- It needs no asset values, treats defaults as events with an uncertain Poisson mean rate, and so cannot say why a firm defaults.
- The similarity is that all four are ways of putting numbers on default risk over a horizon.
- The difference is where the numbers come from.

26 RAROC: does the loan earn its capital?

RAROC = (Revenues − Costs − EL) ÷ Economic capital
Creates value if RAROC > hurdle rate

Revenues	Interest and fees, net of funding
Costs	Operating costs
Expected loss	Expected loss: an expected cost
Economic capital	Held against UL, not the loan size
Hurdle rate	Shareholders' required return

IN THE LESSON

- Finally, is a loan worth making?
- Harbourline compares the profit from a loan with the capital that the loan consumes.
- That is risk-adjusted return on capital, R A R O C.
- The numerator is revenues, meaning interest and fees net of funding cost, minus operating costs, minus expected loss, which is subtracted because it is an expected cost of lending.
- The denominator is economic capital, the capital held against unexpected loss.
- It is not the loan amount.
- Some versions add the return earned on the economic capital, for instance at the risk-free rate, to the numerator, and a question will say so if it applies.
- A loan creates value if its R A R O C exceeds the hurdle rate, typically the shareholders' required return on equity.

27 Example 12: a loan that clears the hurdle

Marlow Freight loan $20m. $m: net revenue 0.90, costs 0.15, EL 0.20, capital 2.5. Hurdle 15%

Profit after EL	0.90 − 0.15 − 0.20	**0.55**
RAROC	0.55 ÷ 2.5	**22%**
Compare with hurdle	22% > 15%	**creates value**
Wrong: divide by loan size	0.55 ÷ 20	**2.75%**

IN THE LESSON

- First loan: twenty million dollars to Marlow Freight, a made-up trucking firm.
- In millions, net revenue is point nine, operating costs are point one five, and expected loss is point two.
- Profit after expected loss is point five five.
- Economic capital is two point five.
- R A R O C is point five five over two point five, which is twenty two per cent.
- Harbourline's hurdle rate is fifteen per cent, so the loan clears it and creates value.
- Look at the wrong answer: dividing by the twenty million loan size gives only two point seven five per cent, a profit measured against the wrong base.
- The denominator is the capital the loan consumes, not the loan amount.

28 Example 13: bigger profit, worse loan

Loan B $100m. $m: net revenue 3.00, costs 0.50, EL 1.00, capital 12. Hurdle 15%

RAROC	(3.0 − 0.5 − 1.0) ÷ 12	**12.5%**
Verdict	12.5% < 15%	**destroys value**
Wrong: forget EL	(3.0 − 0.5) ÷ 12	**20.8%**
If 4% return on capital added	(1.5 + 0.48) ÷ 12	**16.5%**

Only if the question says so: higher profit does not mean higher RAROC

IN THE LESSON

- Now a bigger loan: one hundred million dollars.
- Net revenue is three, costs are point five, expected loss is one, and economic capital is twelve.
- Profit after expected loss is one point five, more dollars than the first loan.
- But R A R O C is one point five over twelve, twelve point five per cent, below the fifteen per cent hurdle, so this loan destroys value.
- Comparing R A R O C with the hurdle rate, rather than ranking loans by absolute profit, avoids favouring large, capital-hungry loans.
- Two more traps.
- Forgetting expected loss gives two point five over twelve, twenty point eight per cent, which flatters the loan.
- And if a question adds a four per cent risk-free return on the capital, that adds point four eight to the numerator and gives sixteen point five per cent.
- The question will say if so.

29 Recap: the first half

- CAMEL: capital, assets, management, earnings, liquidity
- EL = PD × LGD × EAD; UL is the spread around it
- CPD = 1 − (1 − p)^T, never T × p; keep one horizon
- CAR = capital ÷ RWA (charges × 12.5), not ÷ total assets
- A write-off cuts capital and, if removed, RWA too

IN THE LESSON

- Let us pull it together.
- The CAMEL system rates a bank on capital adequacy, asset quality, management, earnings and liquidity, and a bank can look strong on four letters and fail on the fifth.
- Credit risk is quantified through exposure at default, loss given default, probability of default and a common horizon.
- Expected loss is the product of the three and is priced.
- Unexpected loss is the volatility around it, and capital covers it.
- A cumulative default probability is one minus the survival probability compounded, never T times p.
- The capital adequacy ratio is capital divided by risk-weighted assets, with market and operational charges multiplied by twelve point five, and not capital divided by total assets.

30 Recap: the second half

- Judgmental, empirical, financial: match approach to data
- Merton: PD = N(−d2); more vol, less distance to default
- KMV: DD to EDF by database. CreditMetrics: migration
- CreditRisk+: Poisson counts, no asset values
- RAROC = (rev − costs − EL) ÷ capital, versus hurdle

IN THE LESSON

- Three families predict default: judgmental, empirical and financial, and each suits different data.
- In the Merton model equity is a call on assets, the default probability is N of minus d two, and higher volatility lowers distance to default while raising equity value.
- Moody's K M V maps distance to default to an expected default frequency using a database.
- CreditMetrics revalues loans under every rating.
- CreditRisk plus counts defaults with a Poisson model and uses no asset values.
- Finally, R A R O C is revenue minus costs minus expected loss, divided by economic capital, and a loan creates value only if it beats the hurdle rate.

31 Try three quick questions

1. Limit $8m, $5m drawn, 50% of undrawn drawn. Recovery 60%, PD 2%. EL?
   Answer: EAD 6.5m, LGD 40%. 2% × 40% × 6.5m = $52,000.
2. Capital $90m. Corporate $700m (100%), mortgages $300m (35%), cash $200m (0%). CAR?
   Answer: RWA = 700 + 105 = 805. 90 ÷ 805 = 11.2%.
3. Revenue $2.0m, costs $0.4m, EL $0.3m, capital $8m, hurdle 14%. Value created?
   Answer: RAROC = 1.3 ÷ 8 = 16.25% > 14%. Yes.

IN THE LESSON

- Now try three practice questions.
- Pause, think, then listen.
- Question one.
- A credit line has a limit of eight million dollars, with five million drawn.
- Half of the undrawn part is expected to be drawn by default.
- Recovery is sixty per cent and the default probability is two per cent.
- What is the expected loss?
- Exposure at default is five plus one point five, which is six point five million.
- Loss given default is forty per cent.
- So expected loss is fifty two thousand dollars.
- Question two.
- Capital is ninety million.
- Corporate loans of seven hundred million carry a full weight, mortgages of three hundred million thirty five per cent, and cash of two hundred million zero.
- What is the ratio?
- Risk-weighted assets are eight hundred and five, so it is eleven point two per cent.
- Question three.
- Net revenue is two million, costs point four, expected loss point three, economic capital eight million, hurdle fourteen per cent.
- Value created?
- R A R O C is sixteen point two five per cent, so yes.