Estimating Default Probabilities

FRM Part II

1 A story: how likely is Harbour Mills to fail?

- Northgate may lend $5m to Harbour Mills, a maker
- Question: how likely is Harbour Mills to default?
- Ratings, scores, history, prices, models all answer
- Each gives a different kind of number: know which you hold

IN THE LESSON

- Northgate Bank is asked to lend five million dollars to Harbour Mills, a mid-size manufacturer.
- The loan officer asks the obvious question.
- How likely is it that Harbour Mills fails to pay?
- Nobody can see the future, so the bank estimates the chance, and there are many ways to do it.
- You can ask a rating agency for its opinion.
- You can build a score from the company's accounts.
- You can count how often similar firms failed in the past.
- You can read the price of insurance against its default, or model its shares.
- Each method gives a different kind of number, and a good risk manager knows which one is in hand.

2 External ratings versus internal ratings

Two rating systems, built for different jobs

	Agency	Internal
Covers	Public, issuer pays	Unrated, private
Information	Public	Plus private data
Style	Through the cycle	Point in time
Reaction	Slow, if lasting	Quick
Basel IRB	Obligor: PD	Facility: LGD

IN THE LESSON

- Start with ratings.
- An agency rating, from firms such as Standard and Poor's, Moody's and Fitch, is an opinion on creditworthiness.
- The issuer pays for it, and it covers mainly public debt issuers.
- Agencies aim for ratings that are stable through the cycle, so they change a rating slowly, and only when they judge the change to be lasting.
- Banks also build internal ratings.
- These cover the unrated mid-size and private borrowers that agencies ignore, use private information such as financial statements, account behaviour and covenant compliance, and are usually point in time, reacting quickly to conditions.
- Under the Basel internal ratings approach the obligor, which drives default probability, is rated separately from the facility, which drives loss given default.
- Never treat the two as interchangeable: in a recession through-the-cycle ratings barely move while point-in-time ratings migrate.

3 Altman's Z score

Z = 1.2·X1 + 1.4·X2 + 3.3·X3 + 0.6·X4 + 1.0·X5
Higher Z = lower default risk

X1, liquidity	Working capital ÷ total assets
X2, retained earnings	Retained earnings ÷ total assets
X3, operating return	EBIT ÷ total assets
X4, market cushion	Market equity ÷ book liabilities
X5, turnover	Sales ÷ total assets

IN THE LESSON

- Altman's Z score turns a company's accounts into a warning of failure.
- It is a weighted sum of ratios, called a linear discriminant function, fitted to separate bankrupt manufacturing firms from healthy ones.
- Five ratios go in.
- X one is working capital over total assets.
- X two is retained earnings over total assets.
- X three is earnings before interest and tax, over total assets.
- X four is the market value of equity over the book value of total liabilities.
- X five is sales over total assets.
- For public companies the weights are one point two, one point four, three point three, zero point six and one point zero.
- A higher score means a lower risk of default.

4 Example 1: Z score of Harbour Mills

Ratios: X1 0.15, X2 0.20, X3 0.08, X4 0.90, X5 1.10

Liquidity and history	1.2×0.15 + 1.4×0.20	**0.460**
Operating return	3.3 × 0.08	**0.264**
Market cushion	0.6 × 0.90	**0.540**
Turnover	1.0 × 1.10	**1.100**
Z score	0.460+0.264+0.540+1.100	**2.364**

IN THE LESSON

- Here are the ratios for Harbour Mills.
- Working capital is fifteen per cent of assets and retained earnings are twenty per cent.
- Earnings before interest and tax are eight per cent of assets.
- The market value of equity is ninety per cent of book liabilities, and sales are one point one times assets.
- Multiply each by its weight.
- One point two times point one five is point one eight, and one point four times point two is point two eight, together point four six.
- Three point three times point zero eight is point two six four.
- Point six times point nine is point five four.
- One times one point one is one point one.
- The total is two point three six four, so Z is about two point three six.

5 Sorting firms into classes

Cut-offs: about 2.99 and 1.81

Firm	Z and zone
Harbour Mills	2.36: grey
Firm B	3.40: safe
Firm C	1.50: distress
Cut-offs	>2.99 safe, <1.81 distress
Caution	Not a PD; for manufacturers

IN THE LESSON

- Now sort firms into credit quality classes.
- The usual cut-offs are about two point nine nine and one point eight one.
- Above two point nine nine is the safe zone.
- Below one point eight one is the distress zone.
- In between is a grey zone, where the score makes no firm call.
- Harbour Mills, at two point three six, is grey.
- Firm B scores three point four, safe, and firm C one point five, distress.
- Two warnings.
- First, Z is a score, not a probability.
- To turn it into a default probability, you need a mapping built from historical experience, such as the share of past firms in each score band that went bankrupt.
- Second, this version suits manufacturers; variants exist for private firms and non-manufacturers.

6 From a rating to a default probability

q = (F(t) − F(t−1)) ÷ (1 − F(t−1))

F(t)	Cumulative PD by year t
F(t−1)	Cumulative PD a year earlier
q, conditional PD	Default in year t, given survival

IN THE LESSON

- A rating is a ranking.
- Double A is safer than A, which is safer than triple B, and so on.
- But a ranking is not a probability.
- It becomes one only through calibration to history: the agencies publish the fraction of issuers in each rating that defaulted within one, two, five or more years.
- Probabilities rise as the rating worsens, slowly at the top and steeply among speculative grades.
- There is also a term structure.
- Let F of t be the cumulative default probability by year t.
- The conditional probability of default in year t, given survival to the year before, is the increase in F divided by the survival probability.
- That is q equals F of t minus F of t minus one, over one minus F of t minus one.

7 Example 2: conditional PD, strong versus weak

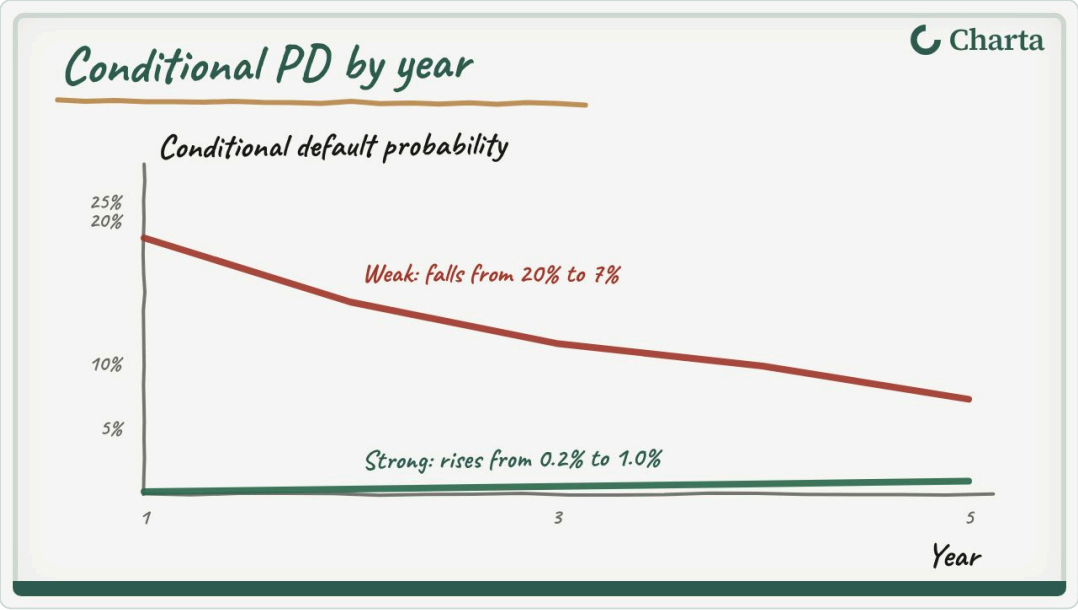
Cumulative PD, years 1 and 2: strong 0.2%, 0.6%. Weak (CCC) 20%, 32%

Strong, year 2 conditional	(0.006 − 0.002) ÷ (1 − 0.002)	**0.40%**
Weak, year 2 conditional	(0.32 − 0.20) ÷ (1 − 0.20)	**15.0%**
Strong, constant-hazard 2 years	$1 - 0.998^2$ (true 0.6%)	**0.40%**
Weak, constant-hazard 2 years	$1 - 0.80^2$ (true 32%)	**36%**

IN THE LESSON

- The strong issuer has cumulative default probabilities of point two per cent after one year and point six per cent after two.
- The conditional probability in year two is the increase, point four per cent, divided by the point nine nine eight who survived: point four zero per cent, double the first year.
- The weak issuer, a triple C name, has twenty per cent and thirty two per cent.
- The increase is twelve per cent, divided by the eighty per cent who survived, which is fifteen per cent, down from twenty.
- A constant hazard shortcut gets both wrong: point four per cent against a true point six, and thirty six against a true thirty two.

8 Conditional PD by year



IN THE LESSON

- This chart shows the term structure of conditional default probability for our two issuers over five years.
- Along the bottom is the year.
- Up the side is the probability of default in that year, given survival to its start.
- The weak issuer's line, in red, starts at twenty per cent and falls to about seven per cent by year five, because the firms still alive are the stronger ones.
- The strong issuer's line, in green, starts near zero and creeps upward, because today's strength may fade with time.
- One constant hazard rate describes neither.

9 The rating migration matrix

One-year migration probabilities, illustrative

From \ to	Good	Weak	Default
Good	90%	8%	2%
Weak	20%	60%	20%
Default	0%	0%	100%
Rules	Rows sum to 100%		Default absorbs

IN THE LESSON

- A rating migration matrix, also called a transition matrix, shows the probability that a borrower who starts the year in one rating ends it in each rating, including default.
- Ours has three states: good, weak and default.
- A good borrower has a ninety per cent chance of staying good, eight per cent of slipping to weak, and two per cent of defaulting.
- A weak borrower has a twenty per cent chance of recovering to good, sixty per cent of staying weak, and twenty per cent of defaulting.
- Every row sums to one hundred per cent.
- Default is an absorbing state: that row is zero except for a one hundred per cent chance of staying in default.
- The ordinary, one-year default probability is read from the default column.

10 Example 3: cumulative PD from the matrix

Good borrower, two years. One-year PD 2%

Good, then default	0.90 × 0.02	**1.8%**
Weak, then default	0.08 × 0.20	**1.6%**
Defaulted in year 1	0.02 × 1.00	**2.0%**
Cumulative PD by year 2, C(2)	1.8% + 1.6% + 2.0%	**5.4%**
Wrong: compound 1-year PD	$1 - 0.98^2$	**3.96%**

IN THE LESSON

- To look further ahead, multiply the matrix by itself.
- The two-year matrix is the one-year matrix squared, and its default column gives the cumulative probability of default by year two.
- Follow a good borrower.
- It can stay good in year one, probability point nine, and default in year two, point zero two: one point eight per cent.
- It can slip to weak with point zero eight and then default with point two, one point six per cent.
- Or it defaults at once, two per cent.
- The total, C of two, is five point four per cent.
- The classic error is to compound the one-year figure, one minus point nine eight squared, which gives only three point nine six per cent and ignores migration.

11 Example 4: marginal and conditional PD

C(1) 2.0%, C(2) 5.4%, C(3) 9.45% (matrix cubed)

Marginal, year 2	5.4% − 2.0%	**3.40%**
Conditional, year 2	3.40% ÷ (1 − 0.02)	**3.47%**
Cumulative, year 3	0.826×0.02 + 0.120×0.20 + 0.054	**9.452%**
Marginal, year 3	9.452% − 5.4%	**4.05%**
Conditional, year 3	4.052% ÷ (1 − 0.054)	**4.28%**

IN THE LESSON

- Write C of n for the cumulative probability by year n.
- The marginal, or unconditional, probability of default in year n is the increase in the cumulative figure.
- In year two that is five point four minus two, which is three point four per cent.
- The conditional probability divides the marginal figure by the survival probability to the start of the year.
- In year two, three point four over ninety eight is three point four seven per cent.
- Cube the matrix: point eight two six times point zero two, plus point one two times point two, plus point zero five four, is nine point four five two per cent.
- The marginal is four point zero five, and divided by the ninety four point six per cent who survived, the conditional is four point two eight per cent.

12 The hazard rate

Survival = e^(−λt)
Unconditional (t1 to t2) = e^(−λ·t1) − e^(−λ·t2)
Conditional (next Δt) = 1 − e^(−λ·Δt)

Hazard rate, λ	A rate, not a probability
Survival	Default by t is 1 minus this
Unconditional PD	Chance seen from today
Conditional PD	Chance given survival to t1

IN THE LESSON

- The hazard rate, or default intensity, is written lambda of t.
- Lambda times a small time step is the probability of default during that step, given no default before.
- It is a rate, not a probability.
- With a constant hazard rate, the survival probability to time t is e to the minus lambda t, and the default probability is one minus that.
- The unconditional probability of default between times t one and t two, as seen today, is e to the minus lambda t one, minus e to the minus lambda t two.
- The conditional probability of default in the next period of length delta t, given survival to t one, is one minus e to the minus lambda delta t.
- If the hazard varies, replace lambda times t by the cumulative hazard, the area under the curve.

13 Example 5: hazard rate of 3%

Constant hazard rate λ = 3% a year

Survival to year 3	e^(−0.03 × 3)	**91.39%**
Default by year 3	1 − 0.9139	**8.61%**
Unconditional, year 3	e^(−0.06) − e^(−0.09)	**2.78%**
Conditional, year 3	1 − e^(−0.03)	**2.96%**
Check	2.783% ÷ 94.18%	**2.955%**

IN THE LESSON

- Suppose Harbour Mills has a constant hazard rate of three per cent.
- Survival to year three is e to the minus point zero nine, which is ninety one point three nine per cent.
- So the cumulative default probability is eight point six one per cent.
- Now the year-three piece.
- The unconditional probability of default between year two and year three, seen from today, is e to the minus point zero six minus e to the minus point zero nine, which is two point seven eight per cent.
- The conditional probability, given survival to year two, is one minus e to the minus point zero three, which is two point nine six per cent.
- The conditional figure is larger, because it divides by the ninety four point one eight per cent survival.

14 Recovery rates and loss given default

- Recovery = bond price just after default, % of face value
- LGD = 1 − recovery. EL = EAD × PD × LGD
- Senior secured > senior unsecured > subordinated
- Absolute priority: each class paid in full before the next
- Recovery falls when defaults rise: fixed average understates

IN THE LESSON

- What does a lender get back after a default?
- The recovery rate is the share of a defaulted claim's face value that creditors recover, commonly measured as the bond's market price shortly after default, as a percentage of face value.
- Loss given default is one minus the recovery rate.
- Expected loss then combines exposure at default, the probability of default, and loss given default.
- Recovery depends on seniority and security.
- Senior secured claims recover more than senior unsecured ones, which recover more than subordinated claims, and under absolute priority each class is paid in full before the next receives anything.
- Empirically, recovery rates are negatively correlated with default rates: when many firms default, assets sell into a weak market.

15 Example 6: a bad year for recovery

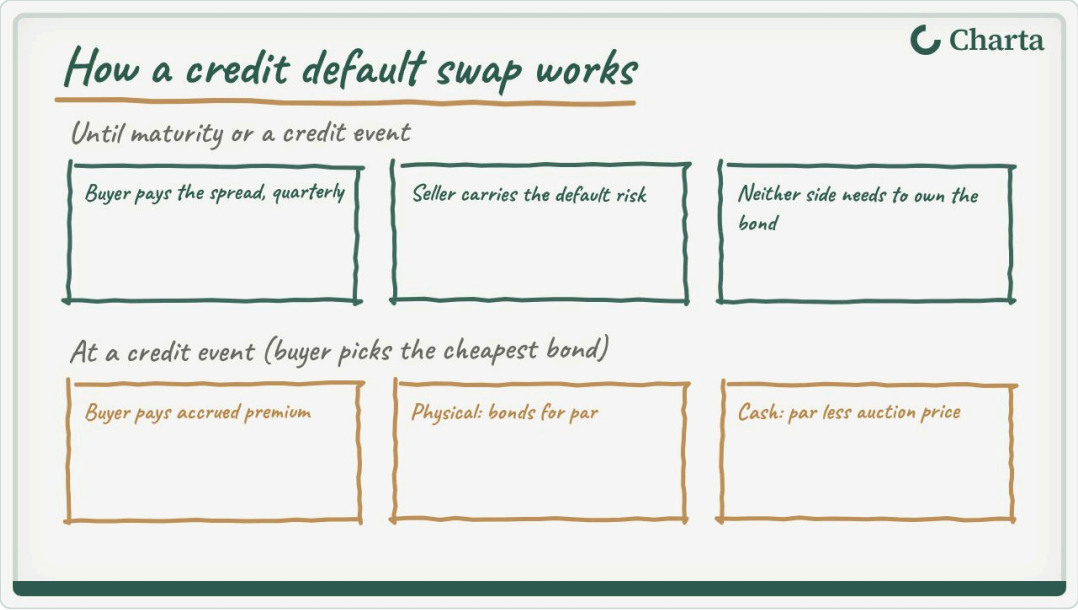
EAD $5m. Normal year: PD 3%, recovery 40%. Recession: PD 6%

Normal year EL	5m × 3% × 60%	**$90,000**
Recession, recovery fixed 40%	5m × 6% × 60%	**$180,000**
Recession, recovery falls to 30%	5m × 6% × 70%	**$210,000**
Understated by	210,000 − 180,000	**$30,000**

IN THE LESSON

- Northgate holds five million dollars of Harbour Mills debt.
- In a normal year the default probability is three per cent and recovery is forty per cent, so loss given default is sixty per cent.
- Expected loss is five million times point zero three times point six, which is ninety thousand dollars.
- Now a recession: defaults double to six per cent.
- If we wrongly keep recovery at forty per cent, expected loss is one hundred and eighty thousand.
- But in a recession recoveries fall too.
- At thirty per cent recovery, loss given default is seventy per cent, and expected loss is two hundred and ten thousand.
- Treating recovery as a fixed average, independent of defaults, understated the bad-year loss by thirty thousand dollars, and so understates tail credit risk.

16 How a credit default swap works



IN THE LESSON

- A credit default swap, or C D S, is a contract that insures the buyer against default by a reference entity.
- In the top row, the protection buyer pays a premium, the C D S spread, until maturity or a credit event.
- It is quoted in basis points per year of the notional and usually paid quarterly.
- The protection seller carries the default risk.
- At a credit event the buyer also pays the premium accrued since the last payment date.
- Settlement is physical or in cash.
- In physical settlement the buyer delivers bonds with face value equal to the notional and receives par.
- In cash settlement the seller pays par less a post-default market price, usually set by an auction.
- Neither party needs to own the bond, and the buyer chooses which deliverable bond to hand over, the cheapest-to-deliver option.

17 Example 7: a ten million dollar CDS

Notional $10m, spread 120 bp, auction price 40 (recovery 40%)

Annual premium	10m × 1.20%	**$120,000**
Quarterly payment	120,000 ÷ 4	**$30,000**
Payout at credit event	10m × (1 − 0.40)	**$6,000,000**
Accrued premium, 1 month	120,000 × 1/12	**$10,000**

IN THE LESSON

- Suppose Northgate buys protection on ten million dollars of Harbour Mills debt, at a spread of one hundred and twenty basis points.
- The annual premium is one point two per cent of ten million, which is one hundred and twenty thousand dollars, paid in quarterly instalments of thirty thousand.
- Later a credit event occurs and the auction price is forty, so recovery is forty per cent.
- The seller pays the notional less the recovery value, ten million times one minus point four, which is six million dollars.
- Northgate must also pay the premium accrued since its last payment date.
- If one month has passed, that is one twelfth of the annual premium, or ten thousand dollars.

18 The credit triangle

Spread = λ × (1 − R)
λ = Spread ÷ (1 − R)
PD over T years = 1 − e^(−λT)

CDS spread, s	Annual premium on the notional
Recovery, R	Share recovered after default
Hazard rate, λ	Average over the swap's life
PD over T years	Cumulative, constant hazard

IN THE LESSON

- A C D S spread is the annual premium, as a fraction of the notional, that makes the swap worth zero at the start.
- It reflects the expected annual loss, which gives the credit triangle.
- Let lambda be the average hazard rate over the life of the swap, and R the recovery rate.
- The spread is roughly the hazard rate times loss given default, lambda times one minus R.
- Rearranged, the average hazard rate is the spread divided by one minus R.
- Once you have the hazard rate, the cumulative default probability over a horizon of T years, with a constant hazard rate, is one minus e to the minus lambda T.

19 Example 8: hazard from a spread

5-year spread 120 bp, recovery 40%

Hazard rate	0.012 ÷ (1 − 0.40)	**2.0%**
5-year default probability	1 − e^(−0.02 × 5)	**9.52%**
Wrong: divide by R	0.012 ÷ 0.40	**3.0%**
Wrong: spread as PD	1.2% read as the PD	**1.2%**

IN THE LESSON

- A one hundred and twenty basis point spread, with forty per cent recovery, implies a hazard rate of point zero one two divided by point six, which is two per cent.
- For a five-year swap the cumulative default probability is one minus e to the minus point one, nine point five two per cent.
- Two mistakes are common.
- One is to divide by the recovery rate instead of one minus the recovery rate.
- That would give three per cent, which is wrong.
- The other is to read the spread itself, one point two per cent, as the default probability.

20 Example 9: forward hazard rates

3-year spread 100 bp, 5-year spread 130 bp, recovery 40%

3-year average hazard	1.00% ÷ 0.60	**1.667%**
5-year average hazard	1.30% ÷ 0.60	**2.167%**
Forward hazard, years 3 to 5	(2.167×5 − 1.667×3) ÷ 2	**2.917%**

IN THE LESSON

- Spreads of different maturities bootstrap into forward hazard rates.
- Say the three-year spread for Harbour Mills is one hundred basis points and the five-year spread is one hundred and thirty, with forty per cent recovery.
- The average hazard to year three is point zero one over point six, one point six seven per cent.
- To year five it is point zero one three over point six, two point one seven per cent.
- The forward hazard between years three and five is lambda two times five, minus lambda one times three, divided by the two years between.
- That is ten point eight three minus five, over two, or two point nine two per cent.

21 The CDS-bond basis

Basis = CDS spread − (bond yield − risk-free rate)
Carry = bond yield − CDS spread − funding cost

Basis	Zero if both markets agree
Bond spread	Bond yield − risk-free rate
Funding cost	Cost of paying for the bond
Carry	Net return, before other costs

IN THE LESSON

- The C D S-bond basis compares two prices for the same default risk.
- It is the C D S spread minus the bond's yield spread over the risk-free benchmark.
- If both markets priced the same risk, owning the bond and buying C D S protection would be a risk-free combination earning the risk-free rate, and the basis would be zero.
- A negative basis means protection is cheap relative to the bond.
- A negative basis trade buys the bond and buys protection.
- Its net annual return, before other costs, is the bond yield minus the C D S spread minus the funding cost of buying the bond.
- With funding at the risk-free rate, carry equals minus the basis.

22 Example 10: a negative basis trade

Bond yield 5.0%, risk-free 3.0%, CDS spread 1.6%

Bond spread	5.0% − 3.0%	**200 bp**
Basis	160 − 200	**−40 bp**
Carry, funding at 3.0%	5.0 − 1.6 − 3.0	**0.40%**
Carry, funding at 3.4%	5.0 − 1.6 − 3.4	**0.00%**
Carry, funding at 3.5%	5.0 − 1.6 − 3.5	**−0.10%**

IN THE LESSON

- Harbour Mills has a bond yielding five per cent, and the risk-free rate is three per cent, so the bond spread is two hundred basis points.
- The C D S spread is one point six per cent, or one hundred and sixty.
- The basis is one hundred and sixty minus two hundred, minus forty basis points.
- Northgate buys the bond and the protection.
- If it can fund the bond at the risk-free three per cent, carry is five minus one point six minus three, which is point four per cent, the negative of the basis.
- But the trade is not riskless when funding costs exceed the risk-free rate.
- At three point four per cent funding, carry vanishes.
- At three point five, it is a loss of point one per cent.

23 Why the basis is not zero

Forces that move the basis away from zero

Force	Effect on CDS spread
Counterparty risk in CDS	Pushes it down
Cheapest-to-deliver option	Pushes it up
Broad credit event terms	Pushes it up
Illiquidity, funding, benchmark	Also move the basis
Negative basis trade	Buy bond + buy CDS

IN THE LESSON

- Several forces push the basis away from zero.
- Counterparty risk in the C D S pushes the C D S spread down, because the protection buyer might not be paid if the seller fails.
- The cheapest-to-deliver option works the other way, because the buyer's right to choose the bond makes protection more valuable.
- Broad credit event definitions also push the spread up.
- Other drivers are bond illiquidity, funding costs and the choice of risk-free benchmark, which also matter.
- The negative basis trade buys the bond and the protection, but can lose money when funding costs exceed the risk-free rate.

24 Historical versus spread-implied PDs

Spread 120 bp, recovery 40%. Historical default rate 0.8% (illustrative)

Spread-implied (risk-neutral)	1.2% ÷ (1 − 0.40)	**2.0%**
Historical (real-world)	counted default rate	**0.8%**
Ratio	2.0 ÷ 0.8	**2.5 times**
Expected-loss part of spread	0.8% × 0.60	**0.48%**
Rest of the spread	1.20% − 0.48%	**0.72%**

Spreads also pay for systematic risk, illiquidity and tax

IN THE LESSON

- There are two ways to estimate default probabilities.
- Historical estimates count how often issuers of a given rating defaulted, which gives real-world probabilities.
- Spread-based estimates infer a default intensity from bond or C D S spreads, which gives risk-neutral probabilities.
- For Harbour Mills, a one hundred and twenty basis point spread with forty per cent recovery implies a hazard rate of two per cent.
- Suppose issuers of its rating historically defaulted at point eight per cent a year.
- The spread-implied figure is two and a half times larger.
- The gap arises because spreads also pay for systematic risk, illiquidity and tax.
- Only point eight times point six, point four eight per cent, is expected loss.
- The other point seven two is compensation for other things.

25 Real-world versus risk-neutral: which to use

Rule: risk-neutral to value, real-world to forecast

Task	Probability
Value a credit derivative	Risk-neutral
Market value of a loan, CVA	Risk-neutral
Expected loss, provisions	Real-world
Credit VaR, stress, capital	Real-world
Merton: only the drift differs	r versus μ

IN THE LESSON

- Real-world default probabilities are the actual chances that defaults occur, estimated from history or from models of the firm.
- Risk-neutral probabilities are implied by market prices, and they are higher because they embed the premium investors demand for bearing default risk.
- The rule: risk-neutral to value things, real-world to estimate what will happen.
- Valuing a credit derivative, pricing a loan at market value, or computing a C V A charge for accounting purposes needs risk-neutral probabilities, because the price must be consistent with other traded instruments.
- Forecasting losses, provisioning, credit value at risk, stress testing and economic or regulatory capital need real-world probabilities.
- In the Merton model, the two differ only in the asset drift.

26 The Merton model: equity and debt

E0 = V0·N(d1) − D·e^(−rT)·N(d2)
d1 = [ln(V0÷D) + (r + σ²÷2)T] ÷ (σ√T)
d2 = d1 − σ√T, D0 = V0 − E0

Equity, E0	A call on the firm's assets
Debt, D0	What is left: V0 − E0
d1 and d2	Normal-curve inputs
Asset volatility, σ	Volatility of firm value

IN THE LESSON

- The Merton model is structural.
- Assets follow geometric Brownian motion, and debt is a single zero-coupon claim with face value D, maturing at T.
- Equity holders hold a call option on the assets with strike D, so equity is V zero times N of d one, minus D times e to the minus rT times N of d two.
- Here d one is the log of V zero over D, plus the quantity r plus sigma squared over two, times T, all over sigma root T.
- And d two is d one minus sigma root T.
- Debt is what is left, V zero minus equity, which equals the present value of D at the risk-free rate minus a put option on the assets.

27 Example 11: valuing equity and debt

V0 = $120m, D = $100m, r = 4%, σ = 25%, T = 1 year. $ millions

d1, d2 (N: 0.8448, 0.7777)	(ln 1.2 + 0.07125) ÷ 0.25	**1.0143; 0.7643**
PV of the debt face value	100 × e^(−0.04)	**96.08**
Equity E0	120×0.8448 − 96.08×0.7777	**26.66**
Debt D0	120 − 26.66	**93.34**
Put on assets	96.08 − 93.34	**2.74**

IN THE LESSON

- Harbour Mills has assets of one hundred and twenty million dollars and zero-coupon debt of face value one hundred million, due in a year.
- The risk-free rate is four per cent and asset volatility twenty five per cent.
- D one is the log of one point two, plus point zero seven one two five, over point two five, which is one point zero one four.
- D two is point seven six four.
- N of d one is point eight four four eight and N of d two is point seven seven seven seven.
- Equity is one hundred and twenty times point eight four four eight, minus ninety six point oh eight times point seven seven seven seven: twenty six point six six million.
- Debt is the rest, ninety three point three four million.
- The put is worth two point seven four.

28 Example 12: volatility of equity

Link: σE × E0 = N(d1) × σV × V0

Right side	0.8448 × 25% × 120	**25.34**
Equity value	from Example 11	**26.66**
Equity volatility σE	25.34 ÷ 26.66	**95.1%**
Asset volatility σV	given	**25%**

Errors: undiscounted D; no N(d1); σE as σV

IN THE LESSON

- The link is that equity volatility times equity value equals N of d one, times asset volatility, times asset value.
- With Harbour Mills, N of d one is point eight four four eight, asset volatility is twenty five per cent and assets are one hundred and twenty.
- The right side is point eight four four eight times point two five times one hundred and twenty, which is twenty five point three four.
- Divide by equity of twenty six point six six, and equity volatility is about ninety five per cent.
- Leverage makes equity far more volatile than the assets.
- The usual errors are forgetting to discount the debt, omitting N of d one in the link, and using equity volatility as asset volatility.

29 Distance to default

DD = [ln(V0÷D) + (μ − σ²÷2)T] ÷ (σ√T)
PD = N(−DD)
Risk-neutral: μ = r, so DD = d2

Distance to default	Standard deviations above default
Drift, μ	Actual return; r if risk-neutral
Default probability	Larger DD, lower PD

IN THE LESSON

- In the Merton model default happens only at maturity, when the asset value falls below the face value of the debt.
- Distance to default is the number of standard deviations by which the expected log asset value exceeds the default point.
- Under the risk-neutral measure mu equals r, and the distance becomes d two.
- The probability of default is the normal probability N of minus distance to default, so a larger distance means a lower default probability.
- Distance rises with lower leverage, lower asset volatility and a higher expected return.
- A real-world probability uses the actual expected return, a risk-neutral one uses r.

30 Example 13: distance to default

Same firm. Real-world μ = 10%; risk-neutral μ = r = 4%

Real-world DD	(0.1823 + 0.10 − 0.03125) ÷ 0.25	**1.004**
Real-world PD	N(−1.004)	**15.8%**
Risk-neutral DD = d2	μ = r = 4%	**0.764**
Risk-neutral PD	N(−0.764)	**22.2%**
Wrong: N(DD)	N(1.004)	**84.2%**

IN THE LESSON

- Suppose the real-world expected asset return is ten per cent.
- The distance to default is the log of one point two, plus ten per cent minus point zero three one two five, divided by point two five.
- That is one point zero zero four.
- The default probability is N of minus one point zero zero four, which is fifteen point eight per cent.
- Now the risk-neutral case, with mu equal to the four per cent risk-free rate.
- The distance is d two, point seven six four, and the default probability is twenty two point two per cent.
- Because mu is above r, the real-world probability is lower.
- Two errors to avoid: reading N of the distance, eighty four per cent, as the default probability, and using d one instead of d two.

31 Merton, KMV and Kamakura

Three models, three structures

Model	Structure and PD map
Merton	Structural. PD = N(−DD)
Moody's KMV	Structural DD, EDF from history
KMV default point	Short debt + half long debt
Kamakura	Reduced-form hazard model
Merton weakness	Low short-horizon PD

IN THE LESSON

- Three well-known default models differ in structure.
- The Merton model is structural: it treats equity as a call on the firm's assets and backs the asset value and volatility out of equity prices.
- The default probability is the normal probability of minus distance to default.
- It assumes a single zero-coupon debt, default only at maturity and constant volatility, and its short-horizon probabilities for healthy firms are usually too low.
- Moody's K M V keeps the structural distance to default, but uses a default point of short-term debt plus half of long-term debt, and maps distance to an expected default frequency using a large empirical default history instead of the normal curve.
- Kamakura is reduced-form: a statistical hazard or logistic model on firm, market and macro variables, including a distance to default.

32 Example 14: judging a model

1,000 firms, predicted PD 1%. 25 default, 20 in top decile

Predicted defaults	1,000 × 1%	**10**
Actual defaults	observed	**25**
Calibration: actual ÷ predicted	25 ÷ 10	**2.5 times**
Defaulters in riskiest decile	20 ÷ 25	**80%**
Share by pure chance	1 ÷ 10	**10%**

IN THE LESSON

- There are two tests of quality.
- Discrimination asks whether the model ranks defaulters above survivors.
- Calibration asks whether predicted probabilities match realised default rates.
- Suppose a model scores one thousand firms and predicts an average default probability of one per cent, so ten defaults.
- In the next year twenty five firms default, two and a half times as many.
- The calibration is poor: it understated risk.
- But twenty of those twenty five defaulters sat in the model's riskiest ten per cent of firms, eighty per cent, where chance would give ten per cent.
- So discrimination is good.
- A model can rank well and still be badly calibrated, and for provisioning and capital the calibration matters.

33 Recap: the main ideas

- Ratings: agency through the cycle; internal point in time
- Z score ranks; history, matrices, hazards give PDs
- Spread ÷ (1 − R) = hazard; basis = CDS − bond spread
- Risk-neutral to value, real-world to forecast
- Merton: equity is a call; PD = N(−DD)

IN THE LESSON

- Agency ratings are stable through the cycle, internal ratings are usually point in time.
- Altman's Z score is a ranking tool, not a probability.
- A rating becomes a default probability only through history, and conditional probabilities tend to rise for strong names, fall for weak.
- A migration matrix raised to a power gives cumulative probabilities, and marginal and conditional follow.
- The hazard rate turns into survival through e to the minus lambda t.
- Recovery falls when defaults rise.
- A C D S spread divided by one minus recovery gives a hazard rate, and the basis compares C D S with bonds.
- Risk-neutral probabilities value, real-world ones forecast.
- In the Merton model, equity is a call and default probability is N of minus the distance to default.

34 Try three quick questions

1. CDS spread 175 bp, recovery 30%. Implied hazard rate?
   Answer: 0.0175 ÷ 0.70 = 2.5%.
2. Cumulative PD: 3.0% by year 3, 4.5% by year 4. Conditional PD in year 4?
   Answer: (4.5 − 3.0) ÷ (100 − 3.0) = 1.55%.
3. Merton distance to default is 2. Default probability?
   Answer: N(−2) = 2.28%.

IN THE LESSON

- Now try three quick questions.
- Pause the video, think, then listen to the answers.
- Question one.
- A C D S spread is one hundred and seventy five basis points and the recovery rate is thirty per cent.
- What hazard rate does the credit triangle imply?
- Divide point zero one seven five by one minus point three, which is point seven.
- The hazard rate is two point five per cent.
- Question two.
- The cumulative default probabilities are three per cent after three years and four point five per cent after four.
- What is the conditional probability in year four?
- The increase is one point five, divided by ninety seven, which is one point five five per cent.
- Question three.
- A firm has a distance to default of two.
- What is the default probability in the Merton model?
- It is N of minus two, about two point two eight per cent.