CHARTA LEARNING · LESSON NOTES

Credit Scoring and Retail Credit Risk

FRM Part II

1 A story: Larkfield Bank

- Retail: mortgages, cards, auto and personal loans
- Very many small, similar positions, not a few big names
- Larkfield Bank: 200,000 customers
- Losses come from shared shocks, not one bad borrower

IN THE LESSON

- Meet Larkfield Bank.
- It has two hundred thousand customers, and almost all of its lending is small: a mortgage here, a credit card there, a car loan, a personal loan, an overdraft.
- But the bank has a different worry.
- If unemployment rises or house prices fall, thousands of customers struggle at the same time.
- And credit losses are not the only thing that can go wrong.
- The bank funds long mortgages with short deposits, customers repay early, fraudsters attack, and staff sometimes sell people the wrong product.
- Retail banking means a very large number of small, similar positions, with several kinds of risk wrapped around them.

2 Exposure and expected loss

EAD = Drawn balance + CCF × Undrawn limit
EL = PD × LGD × EAD

Credit conversion factor	Share of undrawn limit used
Probability of default	Chance the customer defaults
Loss given default	Share not recovered
Exposure at default	Amount owed at default
Expected loss	Average loss to plan for

IN THE LESSON

- Start with how much is at risk.
- Expected loss, or E L, is the average loss we should plan for over a year.
- It is the product of three parts.
- P D is the probability of default.
- L G D is the loss given default, the share of the exposure we do not get back.
- And E A D is the exposure at default, the amount owed when the customer fails.
- A credit card has a limit, and the customer may not have used all of it.
- But a customer in trouble often draws the rest before failing.
- So we add a share of the undrawn limit, called the credit conversion factor, or C C F.
- Exposure at default equals the drawn balance plus C C F times the undrawn limit.

3 Example 1: a credit card

Limit $10,000, drawn $3,000, CCF 40%, PD 3%, LGD 80%

Exposure at default	3,000 + 40% × 7,000	**$5,800**
Expected loss	3% × 80% × 5,800	**$139.20**
Wrong: drawn balance only	3% × 80% × 3,000	**$72**
Fifty thousand such cards	139.20 × 50,000	**$6.96m**

IN THE LESSON

- Larkfield has a credit card customer with a ten thousand dollar limit.
- She has drawn three thousand, so seven thousand is undrawn.
- Customers like her draw forty per cent of the undrawn part before default.
- Forty per cent of seven thousand is two thousand eight hundred, so exposure at default is three thousand plus two thousand eight hundred, which is five thousand eight hundred dollars.
- Cards like this default with probability three per cent, and the loss given default is eighty per cent.
- Expected loss is point zero three times point eight times five thousand eight hundred, which is one hundred and thirty nine dollars and twenty cents.
- Leave out the undrawn limit and you get seventy two dollars, roughly half.
- Across fifty thousand such cards, expected loss is six point nine six million dollars.

4 The other risks in retail banking

Credit risk is one of several; they interact

Interest rate, liquidity	**Long fixed loans, short deposits**
Prepayment	Early repayment changes cash flows
Operational	Fraud and process failures
Conduct, legal, reputation	Mis-selling, unfair practice
Interaction	Rate rise, payment shock, credit loss

IN THE LESSON

- Credit is not the only risk a retail bank carries.
- Funding long fixed-rate mortgages with short deposits creates interest rate risk and liquidity risk.
- If rates rise, deposits cost more while mortgage income is fixed.
- Prepayment is another surprise: borrowers repay early, usually when rates fall, and the cash flows change.
- Fraud and process failures are operational risk.
- Mis-selling or unfair practice creates conduct, legal and reputational risk.
- And these risks interact.
- If rates rise on floating-rate mortgages, monthly payments jump, a payment shock, and what began as interest rate risk turns into credit risk.
- Exam questions usually ask you to classify a loss by its root cause, or to include undrawn limits in exposure.

5 Retail versus corporate credit risk

Same idea, two different ways of working

	Corporate	Retail
Assessed	Borrower by borrower	Pools by scorecard
Judgement	Analyst	Statistical model
Main worry	Single-name concentration	Systematic factors
Basel IRB default	Obligor level	Facility level allowed
Basel IRB capital	Maturity adjusted	No maturity adjustment

IN THE LESSON

- Now, how does retail credit risk part ways with corporate?
- A corporate loan is assessed borrower by borrower.
- Analysts read financial statements, look at ratings and use judgement.
- The exposures are large and lumpy, so one big name can matter a great deal, and single-name concentration is a real worry.
- Retail credit is the opposite.
- It uses statistical scorecards on large pools of small, similar exposures, and it estimates default probability, loss given default and exposure at default for segments, using historical data.
- The Basel rules reflect this.
- Under the internal ratings based approach, retail default can be defined at the facility level instead of for the borrower as a whole.
- Retail capital formulas have no maturity adjustment, and they use supervisory asset correlations that depend on the retail sub-class.

6 Why a pool of small loans diversifies

SD of pool default rate = √(p × (1 − p) ÷ N)
Bigger N: idiosyncratic risk fades, systematic risk stays

Default probability, p	Chance one loan defaults
Number of loans, N	Independent loans in the pool
Spread of the pool default rate	Shrinks as N grows

IN THE LESSON

- Why can a pool of small loans be treated statistically?
- Let p be the default probability of each loan, and N the number of loans, assumed independent.
- The standard deviation of the pool's default rate is the square root of the whole quantity: p times one minus p, divided by N.
- As the pool grows, the spread shrinks.
- Bad luck and good luck cancel, so risk specific to a borrower, called idiosyncratic risk, diversifies away.
- But the word independent matters.
- Borrowers share an economy.
- What remains in a big retail pool is systematic risk, driven by common factors such as unemployment and house prices.
- So retail losses are mainly systematic, and that is what Larkfield's capital must cover.

7 Example 2: pool size and spread

Each loan has p = 3%. Independent loans

One loan	√(0.03 × 0.97)	**17.06%**
Ten loans	√(0.03 × 0.97 ÷ 10)	**5.39%**
One hundred loans	√(0.03 × 0.97 ÷ 100)	**1.71%**
Ten thousand loans	√(0.03 × 0.97 ÷ 10,000)	**0.17%**
Spread falls by	√(10,000 ÷ 10)	**31.6 times**

IN THE LESSON

- Take a three per cent default probability.
- For one loan the standard deviation is the square root of point zero three times point nine seven, about seventeen per cent.
- A retail pool of one hundred loans cuts it to one point seven one per cent.
- A pool of ten thousand cuts it to point one seven per cent.
- Going from ten loans to ten thousand, a thousand times more loans, shrinks the spread by the square root of a thousand, about thirty two times.
- The usual error is to apply single-name concentration thinking to a granular pool.

8 The dark side of retail credit

- Predatory or subprime lending; income not verified
- Teaser rates that reset to unaffordable payments
- Over-indebtedness across several lenders
- Unfair or discriminatory practices
- Originate-to-distribute: sell every loan, ignore quality

IN THE LESSON

- Retail credit has a dark side.
- It is the harm that follows when lending growth outruns prudence.
- Predatory and subprime lending push expensive loans onto people who cannot afford them.
- Loans made without verifying income, sometimes called no-documentation loans, invite fraud and wishful thinking.
- Teaser-rate products start with a low payment, then reset to an unaffordable one.
- Customers borrow from several lenders at once, so none sees the total debt: that is over-indebtedness.
- And unfair or discriminatory practices harm particular groups.
- Originate-to-distribute made all of this worse.
- A lender that sells every loan to investors keeps little of the loss, so it has little reason to care about quality.

9 What is being done about it

Example: keep 5% of a $400m pool = $20m

No income checks	Ability-to-repay, verify income
Teaser rates	Qualify at fully indexed payment
Sell-everything lenders	Risk retention by securitisers
Hidden total debt	Bureau affordability checks
Biased scorecards	Fair-lending testing
Fragile books	Stress testing

IN THE LESSON

- The responses are rules and controls, and each targets a cause.
- Ability-to-repay rules require lenders to verify income and judge whether the borrower can afford the loan.
- Qualifying borrowers at the fully indexed payment, rather than the teaser payment, stops the reset shock being ignored.
- Risk retention makes securitisers keep a share of the loss.
- As an example, a securitiser selling a four hundred million dollar pool and required to keep five per cent holds twenty million dollars.
- Affordability checks use credit bureau data to see total debt.
- Fair-lending testing checks scorecards for disparate impact on protected groups.
- And stress testing asks whether the book survives a recession.
- The exam point is that the fix targets incentives and affordability, not just pricing.

10 Types of scoring model

Grouped by the question they answer

Application	New applicants: accept or reject
Behavioural	Existing accounts: limits, cross-sell
Collection	Delinquent accounts: cure or recovery
Fraud, attrition, response	Other tasks

IN THE LESSON

- Now to the engine of retail credit: scoring models.
- A credit score is a number that ranks customers by risk, and models are grouped by purpose.
- Application scores rate new applicants and drive accept or reject decisions.
- Behavioural scores use the payment history, balances and utilisation of existing accounts, to manage limits, cross-selling and early collections.
- Collection scores rank delinquent accounts by the likelihood of cure or recovery, so collectors call the right people first.
- Fraud, attrition and response scores serve other tasks: spotting fraud, predicting who will leave, and predicting who will respond to an offer.

11 Inputs, methods and a classic trap

- Inputs: demographics, employment, income, debt ratios
- Bureau history: delinquencies, utilisation, enquiries
- Methods: logistic regression, trees, neural networks
- Statistical models versus judgmental systems
- Trap: accepted-only data is biased; use reject inference

IN THE LESSON

- What goes into a score?
- Typical inputs are demographic and employment data, income and debt ratios, and bureau history: delinquencies, utilisation and the number of recent credit enquiries.
- For existing customers there is behavioural data as well.
- The methods are statistical.
- Logistic regression is the workhorse, with discriminant analysis, decision trees and neural networks as alternatives.
- They contrast with judgmental systems, where an experienced underwriter decides.
- There is one classic trap.
- If you build a model only on applicants who were accepted, you only see outcomes for those people.
- The model is biased, because it never saw how the rejected would have behaved.
- The remedy is reject inference, which estimates the likely outcomes of those rejected, so that the model describes the whole applicant population.

12 From logit to probability of default

logit = intercept + b_1x_1 + b_2x_2 + …
PD = 1 ÷ (1 + e^(−logit))

Logit	The log odds of default
Coefficients times variables	Weights from the data
Probability of default	Always between 0 and 1

IN THE LESSON

- Logistic regression deserves a closer look.
- It models the log odds of default, the natural logarithm of the odds of default against no default.
- The log odds equal an intercept plus a sum of coefficients times variables.
- We call this figure the logit.
- The conversion is this.
- Probability of default equals one divided by one plus e to the power of minus logit, where e is about two point seven one eight.
- Whatever the logit, the result lies between zero and one.
- A logit of zero gives fifty per cent.

13 Example 3: scoring an applicant

Intercept −4.2. Terms: 0.5×2, 1×0.8, 4×0.15, 6×(−0.1)

Sum of the four terms	1.0 + 0.8 + 0.6 − 0.6	**1.8**
Logit	−4.2 + 1.8	**−2.4**
e to the 2.4		**11.02**
Probability of default	1 ÷ (1 + 11.02)	**8.3%**

IN THE LESSON

- Suppose Larkfield's scorecard has an intercept of minus four point two.
- Utilisation, as a fraction of the limit, has a coefficient of two.
- Each delinquency in the last year adds point eight.
- Each credit enquiry in six months adds point one five.
- Each year in the current job subtracts point one.
- An applicant uses half her limit, has one delinquency, four enquiries and six years in her job.
- The four terms are one, point eight, point six and minus point six, which sum to one point eight.
- The logit is minus four point two plus one point eight, which is minus two point four.
- E to the power two point four is about eleven.
- So the probability of default is one divided by one plus eleven point oh two, about eight point three per cent.

14 Mortgage variables: LTV and income ratios

LTV = Loan ÷ lower of price and appraisal
Front-end = Housing payment ÷ Gross income
Back-end = All debt payments ÷ Gross income

Loan to value	Loan against property value
Combined LTV	Adds any second lien
Front-end ratio	Housing costs only
Back-end ratio	Every debt payment

IN THE LESSON

- Now the mortgage.
- Which variables matter?
- First, how much of the property's value the loan covers.
- Loan to value, L T V, is the loan divided by the lower of the purchase price and the appraised value.
- Combined L T V adds any second lien, a second loan secured on the same home.
- Second, affordability, measured by two ratios of gross income.
- The front-end ratio counts housing costs only: the housing payment divided by gross income.
- The back-end ratio counts every debt payment, including cards and car loans, divided by gross income.
- The exam trap is mixing them up: front is housing only, back is everything.

15 Example 4: one mortgage applicant

Price $300k, appraisal $280k, loan $224k, second lien $14k

LTV	224 ÷ 280 (lower value)	**80%**
Combined LTV	(224 + 14) ÷ 280	**85%**
Front-end ratio	1,500 ÷ 6,000	**25%**
Back-end ratio	(1,500+300+150) ÷ 6,000	**32.5%**

IN THE LESSON

- Larkfield considers a home priced at three hundred thousand dollars but appraised at two hundred and eighty thousand, with a first loan of two hundred and twenty four thousand.
- The lower value is the appraisal, so L T V is eighty per cent, rather than the seventy five per cent the price would give.
- A second lien of fourteen thousand brings combined L T V to two hundred and thirty eight over two hundred and eighty, which is eighty five per cent.
- The borrower earns six thousand dollars a month.
- Her housing payment is fifteen hundred, so the front-end ratio is twenty five per cent.
- She also pays three hundred on a car and one hundred and fifty on cards, so all debt payments are nineteen hundred and fifty, and the back-end ratio is thirty two point five per cent.

16 Other variables and the cutoff score

- Also: bureau score, income stability, occupancy, purpose
- Also: loan term and rate type (fixed or floating)
- Cutoff score = the minimum score accepted
- Lower cutoff: more volume, higher default and loss rates
- Loss rate = default rate × LGD (4% × 25% = 1%)

IN THE LESSON

- Other variables matter too: the credit bureau score and history, how stable and well documented the income is, occupancy, meaning owner-occupier or investor, the loan purpose, the term, and whether the rate is fixed or floating.
- A scoring model turns all of these into one score, and the cutoff score is the minimum accepted.
- Lowering the cutoff raises volume, because more applicants pass, but it also raises the default rate and the loss rate of the accepted pool.
- The loss rate is the default rate times the loss given default.
- Suppose the accepted pool defaults at four per cent and the loss given default is twenty five per cent: the loss rate is one per cent.

17 Where to put the cutoff

Profit of band = R × (1 − b) − L × b
Cutoff where profit = 0: b = R ÷ (R + L)

Revenue per good account, R	Earned on each good account
Loss per bad account, L	Lost on each bad account
Bad rate of the band, b	Share of the band that defaults
Profit of the band	Zero at the best cutoff

IN THE LESSON

- Where should the cutoff go?
- Divide applicants into score bands and look at the marginal band, the last band you would accept.
- Let R be the revenue from a good account, L the loss on a bad account, and b the bad rate of that band.
- So profit is R times one minus b, minus L times b.
- The profit-maximising cutoff is where the marginal band earns exactly zero.
- Solving, b equals R divided by R plus L.
- Accept bands while their bad rate is below that break-even bad rate.
- One warning: compare marginal bands, not the portfolio average.
- The average can stay positive while the last band loses money.

18 Example 5: choosing the cutoff

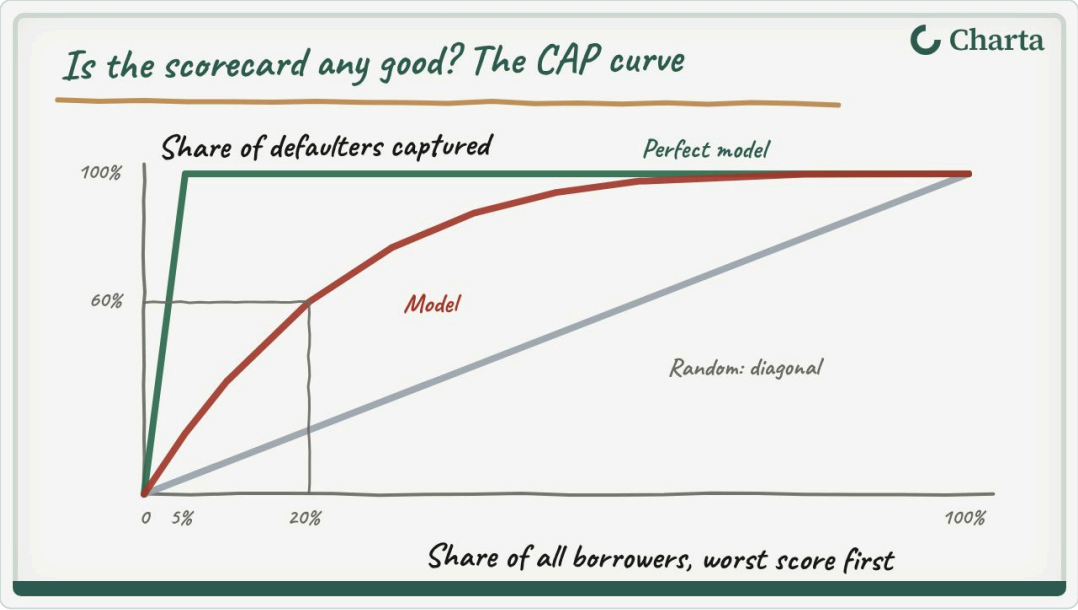
R $200, L $2,000. Break-even bad rate 200 ÷ 2,200 = 9.1%

Band A, bad rate 2%	200×0.98 − 2,000×0.02	**+$156**
Band B, bad rate 5%	200×0.95 − 2,000×0.05	**+$90**
Band C, bad rate 8%	200×0.92 − 2,000×0.08	**+$24**
Band D, bad rate 12%	200×0.88 − 2,000×0.12	**−$64: reject**
Average of the four bands	(156+90+24−64) ÷ 4	**+$51.50**

IN THE LESSON

- Say a good account earns R equals two hundred dollars, and a bad account loses L equals two thousand dollars.
- Four equal score bands, A to D, have bad rates of two, five, eight and twelve per cent.
- Band A earns two hundred times point nine eight, minus two thousand times point zero two: one hundred and fifty six dollars.
- Band B earns two hundred times point nine five, minus two thousand times point zero five: ninety dollars.
- Band C earns two hundred times point nine two, minus two thousand times point zero eight: one hundred and eighty four minus one hundred and sixty, a profit of twenty four dollars.
- Band D earns one hundred and seventy six and loses two hundred and forty, so minus sixty four.
- The break-even bad rate is two hundred over twenty two hundred, nine point one per cent, and band D sits above it.
- The average over all four bands is positive, fifty one dollars fifty, yet band D loses money.
- Accept A to C.

19 Is the scorecard any good? The CAP curve



IN THE LESSON

- To judge a scorecard we first ask whether it ranks well.
- The cumulative accuracy profile, or C A P, does this.
- Rank borrowers from worst score to best.
- Along the bottom is the cumulative share of all borrowers.
- Up the side is the cumulative share of defaulters captured.
- A random model, with no skill, lies on the diagonal: the worst twenty per cent of borrowers hold twenty per cent of the defaulters.
- A perfect model captures every defaulter first, so the green line shoots up and then runs flat, once we pass the five per cent of borrowers who actually default.
- A real model, the red curve, lies in between.
- Here the worst twenty per cent of borrowers hold sixty per cent of the defaulters.
- The more it bows to the top left, the better.

20 The accuracy ratio

AR = Area(model, diagonal) ÷ Area(perfect, diagonal)
AR = 2 × AUC − 1

Accuracy ratio	0 for no power, 1 for perfect
Area under the ROC curve	Area under the ROC curve

IN THE LESSON

- The accuracy ratio, A R, turns the picture into one number.
- Take the area between the model's curve and the diagonal, and divide it by the same area for the perfect model.
- A random model gives zero, a perfect model gives one.
- There is a link to another measure.
- The R O C curve plots, for every cutoff, the share of defaulters caught against the share of good borrowers wrongly flagged.
- Let A U C be the area under that curve.
- Then A R equals two times A U C, minus one.
- An A U C of one half, a coin toss, gives zero.
- An A U C of one gives one.

21 Example 6: measuring the scorecard

1,000 borrowers, 50 defaulters, AUC 0.82

Defaulters in worst 20%	30 of 50	**60%**
Accuracy ratio	2 × 0.82 − 1	**0.64**
Random model	AUC 0.50	**AR = 0**
Perfect model	AUC 1.00	**AR = 1**
All PDs raised by a fifth	ranking unchanged	**AR still 0.64**

IN THE LESSON

- Larkfield tests its mortgage scorecard on one thousand borrowers, of whom fifty defaulted.
- Ranked from the worst score, the worst two hundred borrowers, twenty per cent, contain thirty of the fifty defaulters.
- So the C A P captures sixty per cent of defaulters at the twenty per cent point, against twenty per cent for a random model.
- The analyst also computes the area under the R O C curve and finds point eight two.
- A R is two times point eight two, minus one, which is point six four.
- Now a trap.
- If stress raises every borrower's default probability by a fifth but the order does not change, A R stays at point six four, because it is rank-based, yet the scorecard now under-predicts defaults.

22 Monitoring drift: the PSI

PSI = Σ (actual% − expected%) × ln (actual% ÷ expected%)
Below 0.10: usually read as stable

Population stability index	One number for drift
Actual share now	Share of borrowers in band now
Expected share at development	Share in band at development

IN THE LESSON

- A scorecard is built on a snapshot of customers.
- If the customers change, the scorecard can quietly go wrong.
- The population stability index, P S I, watches for that.
- Split borrowers into score bands.
- For each band, take the actual share, the share of borrowers in the band now, and the expected share, the share at development.
- The index sums, over the bands, actual minus expected, times the natural log of actual divided by expected.
- A P S I below point one zero is usually read as stable.
- That is an industry rule of thumb, not a regulatory standard.
- Larger values call for investigation.

23 Example 7: a stability check

Shares by band, development to now: 15→20, 30→30, 35→30, 20→20

Band 1	0.05 × ln(0.20 ÷ 0.15)	**0.0144**
Band 3	−0.05 × ln(0.30 ÷ 0.35)	**0.0077**
Bands 2 and 4	no change	**0**
PSI	0.0144 + 0.0077	**0.0221**
Reading	below 0.10	**Stable**

IN THE LESSON

- Larkfield compares the score distribution of this year's applicants with the development sample.
- There are four bands, worst to best.
- At development the shares were fifteen, thirty, thirty five and twenty per cent.
- Now they are twenty, thirty, thirty and twenty.
- Band one contributes point zero five times the natural log of twenty over fifteen, which is point zero one four four.
- Band three contributes minus point zero five times the log of thirty over thirty five, which is also positive, point zero zero seven seven.
- Bands two and four did not move, so they add nothing.
- The P S I is point zero two two one, well below point one zero, so the population looks stable.

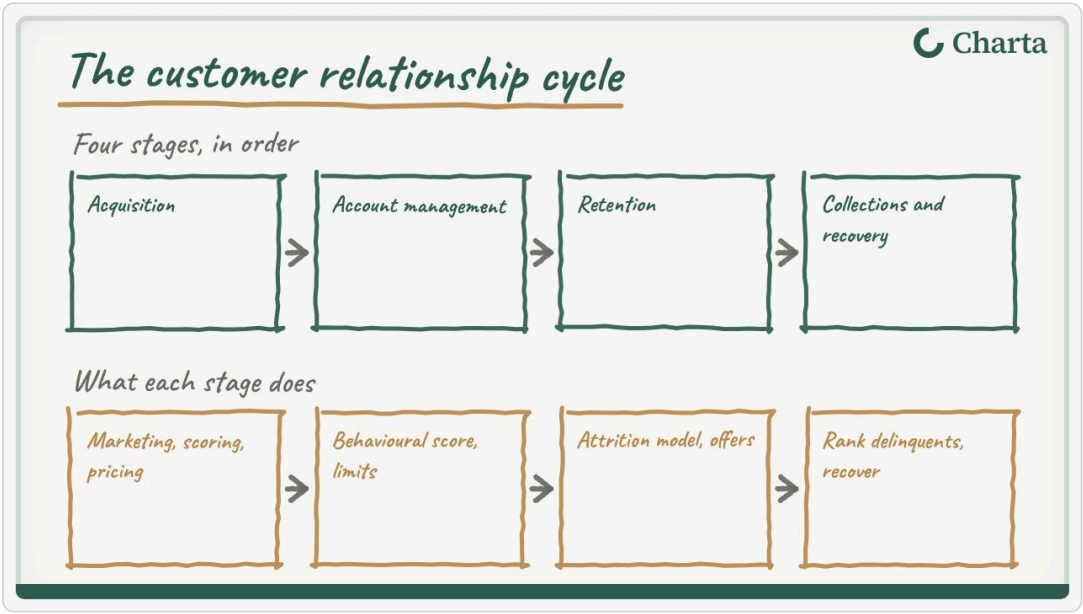
24 Monitoring a scorecard over time

- Judge first on rank ordering: CAP and AR
- Monitor population drift with the PSI
- Backtest AR and calibration by vintage
- Vintage = accounts opened in the same period
- AR is rank-based: a calibration shift leaves it unchanged

IN THE LESSON

- Putting it together, a scorecard is first judged on rank ordering, using the C A P and the accuracy ratio.
- Then it is monitored over its life.
- Monitoring tracks the population stability index, to see whether applicants still look like the development sample.
- It also backtests two things by vintage, meaning groups of accounts opened in the same period.
- One is the accuracy ratio, to check ranking power.
- The other is calibration, comparing predicted default rates with those actually observed.
- Because A R depends only on ranking, a calibration shift does not change it.

25 The customer relationship cycle



IN THE LESSON

- Now zoom out to the customer relationship cycle.
- Follow the boxes.
- First is acquisition: marketing, application scoring and pricing.
- Then account management: behavioural scoring, limit changes and cross-selling.
- Then retention: attrition models predict who may leave, and offers try to keep the good ones.
- Finally collections and recovery, where delinquent accounts are ranked and chased.
- Each stage uses different data and makes different decisions.
- The scoring models from earlier slot into every box.
- A customer is not one decision, but a long series of them.

26 Safe is not the same as profitable

- Creditworthiness and profitability are different
- Transactors pay in full: safe, earn mainly interchange
- Revolvers carry balances: pay interest and fees
- Some risky customers are worth more than their EL
- A risk score is not a profit score: set cutoffs on profit

IN THE LESSON

- Here is a surprise.
- The safest customer is not always the best customer.
- Creditworthiness and profitability are different things.
- Many of the safest customers are transactors: they pay their card balance in full every month.
- They almost never default, but they pay no interest, so Larkfield earns mainly interchange income, the fee merchants pay on card purchases.
- Revolvers carry balances from month to month, so they pay interest and fees.
- They are riskier, yet some of them are worth much more than their expected loss.
- A risk score alone is therefore not a profit score.
- Exam questions test exactly this: a lender that rejects everyone above a certain risk may be throwing away profitable customers, and cutoffs can be set on profit.

27 Account profit and lifetime value

Profit = Interest margin + Fees − Operating cost − EL
EL = PD × LGD × EAD

Interest margin	Interest earned less funding cost
Fees	Including interchange income
Operating cost	Cost of serving the account
Expected loss	PD × LGD × EAD
Customer lifetime value	Adds retention and discounting

IN THE LESSON

- Let us make that precise.
- Expected loss, once more, is P D times L G D times E A D.
- Account profit is what remains of income after costs and expected loss: the interest margin, plus fees, minus operating cost, minus expected loss.
- Customer lifetime value goes further.
- It adds retention, the chance the customer stays, and discounting, because future profit is worth less than profit today.
- The key message is that risk enters profit only through expected loss, which is one item among four.

28 Example 8: transactor versus revolver

Annual figures in dollars. LGD is 80% for both

Transactor: EL	0.5% × 80% × 2,000	**$8**
Transactor: profit	0 + 150 − 70 − 8	**$72**
Revolver: EL	6% × 80% × 3,000	**$144**
Revolver: profit	400 + 90 − 110 − 144	**$236**
Revolver earns	236 ÷ 72	**3.3 times**

IN THE LESSON

- The transactor has two thousand dollars of exposure, pays no interest margin, and generates one hundred and fifty dollars of fees, mostly interchange.
- Operating cost is seventy.
- Default probability is half of one per cent and loss given default is eighty per cent, so expected loss is eight dollars.
- Profit is zero plus one fifty, minus seventy, minus eight: seventy two dollars.
- The revolver has three thousand of exposure, a four hundred dollar interest margin, ninety dollars of fees and a hundred and ten of cost.
- Her default probability is six per cent, so expected loss is one hundred and forty four dollars.
- Profit is four hundred plus ninety, minus one ten, minus one forty four: two hundred and thirty six dollars.
- The riskier customer earns more than three times as much.

29 Risk-based pricing: the break-even rate

r = (1 + cost − PD × (1 − LGD)) ÷ (1 − PD) − 1
Break-even: (1 + r)(1 − PD) + PD(1 − LGD) = 1 + cost

Loan rate, r	Rate that just breaks even
Funding plus operating cost	Per dollar lent
Probability of default	Chance the borrower defaults
Loss given default	Share of the loan not recovered

IN THE LESSON

- Last, pricing.
- Risk-based pricing sets each borrower's rate to cover four things: funding cost, operating cost, expected loss, and a return on the capital the loan uses.
- So higher-risk borrowers pay more.
- Here is the break-even rate for a one-year loan, before the return on capital.
- Let r be the loan rate, cost the funding plus operating cost per dollar lent, and P D and L G D as before.
- If the borrower repays, with probability one minus P D, Larkfield receives one plus r.
- If the borrower defaults, it recovers one minus L G D per dollar lent.
- Expected repayment must equal one plus cost.
- Solving gives r equals one plus cost, minus P D times one minus L G D, all divided by one minus P D, minus one.

30 Example 9: three borrowers, three prices

Cost 3% (funding + operating), LGD 50%

Low risk, PD 1%	(1.03 − 0.005) ÷ 0.99 − 1	**3.54%**
Medium risk, PD 4%	(1.03 − 0.02) ÷ 0.96 − 1	**5.21%**
High risk, PD 8%	(1.03 − 0.04) ÷ 0.92 − 1	**7.61%**
Check, medium	0.96 × 1.0521 + 0.02	**1.03**
Single 5.21% for low risk	5.21 − 3.54	**1.67 too high**

IN THE LESSON

- Take a cost of three per cent and a loss given default of fifty per cent.
- For a low-risk borrower with a one per cent default probability, the numerator is one point oh three minus point zero zero five, which is one point zero two five.
- Divide by point nine nine and subtract one: three point five four per cent.
- At four per cent the rate is five point two one per cent, and at eight per cent it is seven point six one.
- Check the medium case: point nine six times one point oh five two one, plus point oh two, is one point oh three, the amount lent plus cost.
- Now a single price of five point two one per cent charges the low-risk borrower one point six seven percentage points too much.

31 What pricing for risk gains

- Clear profitability by segment; better capital allocation
- Applicants rejected at one price can pass at a higher one
- Less cross-subsidy from low-risk to high-risk borrowers
- Protection against adverse selection by risk-pricing rivals

IN THE LESSON

- What does risk-based pricing gain?
- First, a clearer view of profitability by segment, and so better capital allocation: Larkfield can see which segments earn their keep.
- Second, wider access.
- Applicants rejected at a single price can be accepted at a higher one that covers their risk.
- Third, less cross-subsidy.
- Under one price, low-risk borrowers pay for the losses of high-risk ones, as we just saw.
- Fourth, protection against adverse selection.
- If a competitor prices on risk and Larkfield does not, the competitor takes the good borrowers with cheaper offers, and Larkfield is left with the bad ones at an average price that is too low.

32 Limits of pricing, and the trap

- Very high prices can raise default
- They discourage the best of a risky group
- Fair-lending rules limit the risk factors used
- Break-even rate excludes the return on capital
- Trap: a single price makes good borrowers subsidise bad

IN THE LESSON

- There are limits.
- Very high prices can themselves raise default, because a heavy payment is harder to meet.
- And a high price discourages the best of a risky group, who have other options, leaving the worst.
- Fair-lending rules constrain the risk factors a lender may use, so pricing cannot use everything that predicts default.
- Remember also that the formula gives the break-even rate: a full price adds a return on the capital the loan uses.
- The exam trap is the single price.
- One price for everyone means good borrowers subsidise bad ones, which is both unfair and unstable.
- Pricing by risk is the cure, within these limits.

33 Recap: the main ideas

- EAD = drawn + CCF × undrawn; EL = PD × LGD × EAD
- Retail pools diversify borrower risk; systematic risk stays
- Dark side: fix incentives and affordability, not just price
- Logit to PD; AR = 2 × AUC − 1; PSI below 0.10 is stable
- Cutoff and price on profit: marginal band, break-even rate

IN THE LESSON

- Let us pull it together.
- Retail exposure must include undrawn limits, and expected loss is probability times loss given default times exposure.
- Retail risk comes in many small, similar loans, so borrower-specific risk diversifies and systematic factors remain.
- The dark side of retail credit is fixed by targeting incentives and affordability, not just price.
- Scores turn a logit into a probability of default, mortgages are judged on loan to value and income ratios, and the cutoff sits where the marginal band earns zero.
- A scorecard is measured by the accuracy ratio, and monitored with the population stability index.
- Safe is not the same as profitable.
- And a risk-based price covers cost, expected loss and a return.

34 Try three quick questions

1. Limit $5,000, drawn $2,000, CCF 50%, PD 2%, LGD 75%. EL?
   Answer: EAD = 2,000 + 1,500 = 3,500. EL = 2% × 75% × 3,500 = $52.50.
2. Cost 2%, PD 5%, LGD 60%. Break-even rate?
   Answer: (1.02 − 0.05 × 0.4) ÷ 0.95 − 1 = 5.26%.
3. AUC of 0.90. What is the accuracy ratio?
   Answer: AR = 2 × 0.90 − 1 = 0.80.

IN THE LESSON

- Now try three quick questions.
- Question one.
- A card has a limit of five thousand dollars, of which two thousand is drawn.
- The credit conversion factor is fifty per cent, the default probability two per cent and the loss given default seventy five per cent.
- What is expected loss?
- Exposure is two thousand plus half of three thousand, thirty five hundred dollars.
- Expected loss is point zero two times point seven five times thirty five hundred: fifty two dollars and fifty cents.
- Question two.
- A one-year loan has a cost of two per cent, a default probability of five per cent and a loss given default of sixty per cent.
- What is the break-even rate?
- One plus point zero two, minus point zero five times point four, is one, and one over point nine five, minus one, is five point two six per cent.
- Question three.
- A scorecard has an area under the R O C curve of point nine zero.
- What is its accuracy ratio?
- Two times point nine, minus one, is point eight.