CHARTA LEARNING · LESSON NOTES

Liquidity Risk Management

FRM Part II

1 A story: the fund that could not sell in time

- Investors can redeem every day at NAV
- Many bonds take weeks to sell without cutting the price
- A bad month for credit: redemptions jump
- Early leavers get full value; stayers bear the cost
- Others rush to leave first

IN THE LESSON

- Picture the Harrow Bond Fund.
- Investors can redeem every day at the net asset value, or N A V, the value of the fund's holdings per unit.
- But many of its corporate bonds take weeks to sell without cutting the price.
- Then a bad month for credit markets arrives, and redemptions jump.
- The investors who leave early get full value, while the cost of selling bonds in a hurry falls on the investors who stay.
- Others rush to leave first, which forces more selling at worse prices.
- Nothing here needs a fraud: only a promise of daily cash, backed by assets that cannot become cash in a day.

2 Three problems a fund must control

- Liquidity risk: meeting redemptions without harming value
- Mismatch: terms faster than assets can be sold
- First-mover advantage: early leavers escape the cost
- Dilution: trading costs fall on the investors who stay
- Each tool targets one of these problems

IN THE LESSON

- Let us name the problem precisely.
- For a fund, liquidity risk is the risk that it cannot meet redemptions on time without harming the value left for its other investors.
- It shows up as three linked problems.
- The first is mismatch: the redemption terms are faster than the time the assets need to be sold.
- The second is first mover advantage.
- If leaving early lets you escape the cost of selling, everyone has a reason to be first out of the door, and a trickle of redemptions can become a run.
- The third is dilution: the trading costs of meeting redemptions fall on the investors who stay.
- Every recommended tool targets one of these problems.

3 Control by design: redemption terms

- Dealing frequency: how often investors can redeem
- Notice period: warning required before a redemption
- Minimum holding: how long money must stay in
- Match the terms to the time the assets need to be sold
- That time is measured without material price impact

IN THE LESSON

- The first and most powerful control is design, decided before a single investor joins.
- The fund chooses the redemption terms it offers.
- Dealing frequency is how often investors can redeem, for example daily, monthly or quarterly.
- A notice period is the warning an investor must give before a redemption is paid.
- A minimum holding period is how long money must stay in before it can leave.
- The rule is to match these terms to the time the assets need to be sold.
- That time is measured without material price impact, meaning without pushing the price down noticeably.

4 A liquidity risk framework

- Clear governance: who sets, monitors and acts
- Limit the share of hard-to-sell assets
- Limit investor concentration
- Stress-test asset liquidity
- Stress-test redemption behaviour

IN THE LESSON

- Design is not enough on its own, because markets and investors change.
- Recommended practice adds a liquidity risk framework built on clear governance: it says who sets the liquidity policy, who monitors it, and who acts when a limit is breached.
- The framework sets limits.
- One caps the share of the portfolio held in hard to sell assets.
- Another caps investor concentration, meaning how much of the fund belongs to one investor or a small group, because one large exit can force a large sale.
- Then come regular stress tests on both sides.
- Asset liquidity is stressed by asking how fast holdings could be sold if volumes dried up.
- Redemption behaviour is stressed by asking how much investors might take out in a bad period.

5 Example 1: checking Harrow's limits

NAV $500m. Limits: hard-to-sell at most 15% of NAV; largest investor at most 10%

Hard-to-sell assets	90 ÷ 500	**18%**
Excess over the limit	90 − 15% × 500 = 90 − 75	**$15m**
Largest investor	40 ÷ 500	**8%: inside**
Buffer: cash plus liquid bonds	15 + 85	**$100m**
Stress exit, before haircuts	100 ÷ 40	**2.5 times cover**

IN THE LESSON

- Let us check Harrow against its framework.
- The fund's N A V is five hundred million dollars.
- Its limits say hard to sell assets may be at most fifteen per cent of N A V, and the largest investor at most ten per cent.
- Harrow holds ninety million of hard to sell bonds, ninety over five hundred, eighteen per cent.
- Fifteen per cent of five hundred is seventy five million, so the fund is fifteen million over its limit and must reduce those holdings.
- The largest investor, a pension plan, holds forty million, eight per cent, which is inside its limit.
- Now a simple stress test.
- Harrow's buffer of cash and highly liquid bonds is fifteen plus eighty five, one hundred million.
- If the largest investor left at once, the buffer would cover it two and a half times, before any haircuts, which we return to later.

6 Dilution: who pays to sell

Dilution = selling cost ÷ NAV left after redemptions

Selling cost	Spread and impact paid to sell
NAV left after redemptions	Value owned by those who stay
Dilution	Loss borne by each stayer

IN THE LESSON

- Now the third problem, dilution, in numbers.
- When the fund sells bonds to pay a redeeming investor, it pays transaction costs: the bid ask spread, which is the gap between the price dealers pay and the price they charge, plus any price impact.
- If redeemers receive the full N A V and the fund has no anti dilution tool, those costs come out of the assets that are left.
- Dilution equals the selling cost, divided by the N A V left after redemptions.
- The selling cost is the spread and impact paid.
- The N A V left is the value owned by the investors who stay.
- And dilution is the loss each remaining holder bears, as a share of what they own.

7 Example 2: dilution with no protection

Harrow: NAV $500m; redemptions $50m; selling cost 1% of bonds sold

Redeemers paid at full NAV	10% × 500	**$50m**
Selling cost	1% × 50	**$0.5m**
NAV left after redemptions	500 − 50	**$450m**
Dilution	0.5 ÷ 450	**0.11%**
Stayers now own	450 − 0.5	**$449.5m**

IN THE LESSON

- Take a bad day at Harrow.
- Investors redeem fifty million dollars, ten per cent of the fund, and they are paid at the full N A V.
- To raise the cash, the manager sells fifty million of bonds, and the selling cost is one per cent of the bonds sold, half a million dollars.
- The N A V left after redemptions is five hundred minus fifty, four hundred and fifty million.
- Dilution is half a million over four hundred and fifty million, about point one one per cent.
- The investors who stayed now own four hundred and forty nine and a half million instead of four hundred and fifty.
- It repeats on every heavy day, and it rewards whoever leaves first.

8 Example 3: swing pricing

Threshold: net flows over 5% of NAV. Swing factor 1%. NAV per unit $100

Net outflow today	50 ÷ 500	**10%: swing**
Dealing price	100 × (1 − 1%)	**$99**
Redeemers receive	500,000 units × $99	**$49.5m**
Cost borne by redeemers	50 − 49.5	**$0.5m**
Quiet day: outflow 2%	below the 5% threshold	**no swing**

Stayers keep $450m: no dilution, and leaving first no longer pays

IN THE LESSON

- Anti dilution tools fix this, and the main one is swing pricing.
- When net flows pass a threshold, the dealing price moves by a swing factor, so the investors who trade bear the transaction costs instead of those who stay.
- Harrow sets the threshold at five per cent of N A V and a swing factor of one per cent, its estimated selling cost.
- The N A V per unit is one hundred dollars.
- Today's net outflow is fifty over five hundred, ten per cent, above the threshold, so the dealing price swings down to ninety nine dollars.
- Redeemers hand in five hundred thousand units and receive forty nine and a half million.
- So redeemers bear the half million cost.
- On a quiet day, with an outflow of two per cent, there is no swing.
- Stayers keep four hundred and fifty million, and leaving first no longer pays.

9 Tie each tool to its problem

Judge liquidity in stress, not in calm markets or by average spread

Tool	Problem it targets
Terms matched to asset liquidity	Mismatch
Hard-to-sell and concentration limits	Mismatch; one large exit
Liquidity stress tests	Mismatch, seen before it bites
Swing pricing, anti-dilution	Dilution; first-mover advantage

IN THE LESSON

- The grey line at the top states the usual mistake: liquidity must be judged in stress, not in calm markets, and not by the average spread alone, because spreads widen and volumes shrink exactly when redemptions arrive.
- Now tie each tool to its problem.
- Redemption terms matched to asset liquidity attack the mismatch.
- Limits on hard to sell assets and on investor concentration attack the mismatch too, and the risk of one large exit.
- Liquidity stress tests reveal a mismatch before it bites.
- Swing pricing and the other anti dilution tools attack dilution, and because leavers pay their own way, they also remove the first mover advantage.

10 Why bond funds find liquidity harder

- Bonds trade over the counter, through dealers
- No central order book showing buyers and sellers
- Most individual issues trade rarely
- One issuer may have dozens of separate bonds
- Equities: one share, an exchange, prices all day

IN THE LESSON

- Now to fixed income, which is Harrow's world.
- Bond funds face harder liquidity management than equity funds.
- Bonds trade over the counter, which means a fund must call dealers and negotiate, rather than send an order to an exchange.
- There is no central order book, the public list of buy and sell orders that an exchange shows, so a manager cannot see how much could be sold, or at what price.
- Most individual bond issues trade rarely.
- Part of the reason is that one issuer may have dozens of separate bonds.
- Equities are different: one share, an exchange, prices all day.

11 Evaluated prices flatter the fund

- Evaluated or matrix prices: taken from comparable bonds
- Reported NAV can overstate what a sale would realise
- Smoothed prices understate volatility
- and understate correlation in risk estimates
- So risk and liquidity both look better than they are

IN THE LESSON

- Because most bonds do not trade on a given day, the fund cannot value them from today's trades.
- Instead many holdings use evaluated prices, also called matrix prices: a pricing service estimates a bond's price from comparable bonds, with a similar issuer, rating, maturity and coupon.
- That creates two problems.
- First, the reported N A V can overstate what a sale would realise, because a forced sale of a large block may happen well below the evaluated price.
- Second, evaluated prices are smoothed: they move less, and later, than true values.
- When they feed a risk model, they understate volatility, and they understate correlation between bonds too.
- So risk and liquidity both look better than they are.

12 Biased trade data and shrinking dealers

- Observed trades: mostly small, in calm, liquid conditions
- So average costs understate large sales in stress
- Dealer capacity shrinks when everyone sells at once
- Cost per unit rises with trade size
- Trap: assuming normal-times cost holds in stress

IN THE LESSON

- The trade data that do exist are biased as well.
- The trades we observe are mostly small, and they happen mostly in calm, liquid conditions, because that is when trading is easy.
- So an average cost computed from them understates the cost of selling a large position in stress, which is exactly the sale a fund must plan for.
- Dealers make it worse.
- A dealer buys a bond onto its own balance sheet and holds it until it finds a buyer, and that capacity shrinks when everyone sells at once.
- Cost per unit also rises with trade size, because a large block must be priced lower to tempt a dealer to take it.
- The usual error is assuming that normal times cost holds in stress.

13 How long would it take to sell?

Days = position ÷ (participation rate × expected daily volume)

Position	Amount of the bond held
Participation rate	Share of daily volume taken
Expected daily volume	What the market trades a day
Days to sell	Time to sell it all

IN THE LESSON

- A practical measure ties these ideas together: the number of days needed to sell a position.
- Days equals the position, divided by the participation rate times the expected daily volume.
- The position is the amount of the bond the fund holds.
- The participation rate is the share of the market's daily volume the fund is willing to take, because selling more than that would push the price down.
- The expected daily volume is the amount of that bond the whole market trades on a normal day.
- The result is the time to sell the whole position without material price impact.
- If it is longer than the dealing terms, there is a mismatch.

14 Example 4: days to sell, calm and stressed

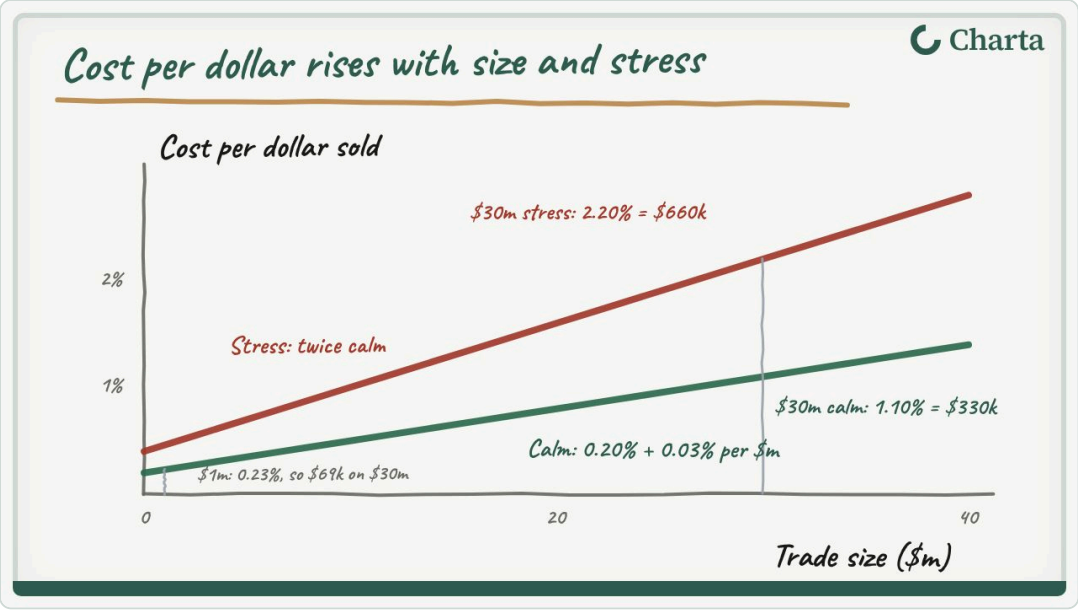
Harrow holds $30m of one bond. Daily volume $4m; participation 25%

Sold per day, calm	25% × 4m	**$1m**
Days to sell, calm	30 ÷ 1	**30 days**
Volume in stress	4m halves	**$2m**
Sold per day, stress	25% × 2m	**$0.5m**
Days to sell, stress	30 ÷ 0.5	**60 days**

IN THE LESSON

- Harrow holds thirty million dollars of one corporate bond.
- The market trades about four million of it on a normal day, and Harrow's policy is to take no more than twenty five per cent of the volume.
- In calm markets it can sell a quarter of four million, one million dollars a day.
- So the position takes thirty divided by one, thirty days to sell.
- Now stress it.
- In a sell off, volume in this bond halves to two million a day.
- At the same participation rate, Harrow can sell only half a million a day, and the position takes sixty days.
- Harrow promises cash every day, so this is a mismatch, and the bond is hard to sell.

15 Cost per dollar rises with size and stress



IN THE LESSON

- This chart shows the cost of selling one bond, as a percentage of the amount sold.
- Along the bottom is the trade size, from zero to forty million dollars.
- Up the side is the cost per dollar sold.
- Suppose a model for Harrow's bond says the calm cost is point two per cent, plus point zero three per cent for each million sold: that is the green line.
- In stress, the cost is twice the calm cost: the red line.
- A small trade of one million costs point two three per cent, the kind of figure seen in the trade data.
- Selling the whole thirty million costs one point one per cent in calm markets, three hundred and thirty thousand dollars, and two point two per cent in stress, six hundred and sixty thousand.
- The small trade figure, applied to the whole thirty million, would suggest just sixty nine thousand dollars.

16 Ways to model asset liquidity

Each measure needs data, and each fails somewhere

Approach	What it measures
Quoted or effective spread	Sale cost: ½ × spread × amount
Roll's estimator	Effective spread from price bounce
Amihud ratio	Price move per dollar traded
Zero-trading days	Share of days with no trades
T-cost regression	Cost from size, rating, maturity, age

IN THE LESSON

- Now to modelling asset liquidity, where there are several approaches, and the premise is that each needs data and each fails somewhere.
- The quoted or effective bid ask spread gives the cost of a sale directly, as half the spread times the amount sold.
- Roll's estimator infers the effective spread from the way prices bounce between bid and ask, when quotes are not available.
- The Amihud ratio measures how far the price moves per dollar traded, so a higher value means less liquid.
- Zero trading day measures count the share of days with no trades at all.
- Finally, a transaction cost regression, or t cost model, predicts the cost from trade size, rating, maturity and age.
- We take them in turn.

17 The spread, and Roll's estimator

Cost of a sale = ½ × spread × amount sold
Roll: Spread = 2 × √(−cov)

Spread	Quoted or effective spread
Cost of a sale	What the seller gives up
Amount sold	Amount of the bond sold
Serial covariance	Covariance of price changes

IN THE LESSON

- First, the spread.
- The quoted spread is the gap between a dealer's ask and bid.
- The effective spread uses actual trades: twice the distance between the trade price and the midpoint.
- A seller receives roughly the midpoint minus half the spread, so the cost of a sale is half the spread times the amount sold.
- But what if no quotes are available?
- Roll's insight is bid ask bounce.
- Trades land alternately at the bid and at the ask, so a price rise tends to be followed by a fall, and successive price changes are negatively correlated.
- Write C O V for their first order serial covariance, the covariance of each price change with the one before.
- Roll's estimator infers the spread as two times the square root of minus C O V.

18 Example 5: Roll's estimator and the cost of a sale

Daily price changes measured in %; their serial covariance is −0.09

Minus the covariance	−(−0.09)	**0.09**
Square root	√0.09	**0.30%**
Spread: times 2	2 × 0.30%	**0.60%**
Cost of selling $20m	½ × 0.60% × 20m	**$60,000**
Wrong: drop the 2	½ × 0.30% × 20m	**$30,000**

Positive covariance or a thin sample: Roll fails

IN THE LESSON

- Apply it to a Harrow bond.
- Its daily price changes are measured in per cent, and their serial covariance is minus point zero nine.
- Minus the covariance is plus point zero nine.
- Its square root is point three per cent.
- Times two, the estimated effective spread is point six per cent.
- Now use it.
- Selling twenty million dollars of this bond costs half of point six per cent times twenty million, which is sixty thousand dollars.
- The usual mistake is to drop the factor two, giving a spread of point three per cent and a cost of only thirty thousand.
- And Roll can fail.
- If the measured covariance is positive, there is no square root to take, and a thin sample of few trades makes the estimate unreliable.

19 Example 6: Amihud ratio and zero-trading days

Five days of one bond: trades on days 1, 3 and 4; none on days 2 and 5

Day 1: |return| ÷ volume	0.4% ÷ $2m	**0.20**
Day 3	0.6% ÷ $1m	**0.60**
Day 4	0.2% ÷ $4m	**0.05**
Amihud: average of trade days	(0.20 + 0.60 + 0.05) ÷ 3	**0.283**
Zero-trading days	2 ÷ 5	**40%**

Wrong: dividing by all 5 days gives 0.17

IN THE LESSON

- Next, the Amihud ratio: the average, over days with trades, of the absolute return divided by trading volume.
- Take five days of a Harrow bond.
- It traded on days one, three and four, and not on days two and five.
- On day one the price moved point four per cent on two million dollars of volume, a ratio of point two, in per cent per million.
- On day three it moved point six per cent on one million, point six.
- On day four, point two per cent on four million, point zero five.
- The Amihud ratio averages the three trade days: point eight five over three, about point two eight three.
- A higher value means less liquid.
- Zero trading days are two out of five, forty per cent.
- Averaging over all five days would wrongly give point one seven.

20 Bonds that rarely trade: a cost model

- Rare trades: no spread, Roll or Amihud data to use
- Fit a t-cost regression on bonds that do trade
- Cost explained by size, rating, maturity and age
- Apply it to the bond in question
- It also gives corporate-bond transaction costs

IN THE LESSON

- What about a bond that has not traded for months?
- There is no spread to observe, too few trades for Roll's covariance, and no volumes for the Amihud ratio.
- The answer is to borrow information from similar bonds.
- Fit a t cost regression, a statistical model of transaction cost, on the bonds that do trade.
- It explains each trade's cost by the characteristics that drive it: the trade size, the bond's credit rating, its maturity, and its age, meaning the time since it was issued.
- Once fitted, the model is applied to the bond in question.
- The same model gives transaction cost estimates for corporate bonds generally, including the cost of a planned sale.

21 Example 7: applying a t-cost model

Cost (bp) = 10 + 2 × size ($m) + 15 if high yield + 1.5 × maturity + 2 × age

Constant		**10 bp**
Size: $5m	2 × 5	**10 bp**
Rating: high yield	15 × 1	**15 bp**
Maturity 8 years, age 3 years	1.5 × 8 + 2 × 3	**18 bp**
Total, and cost of the trade	53 bp × $5m	**$26,500**

IN THE LESSON

- Here is a made up t cost model fitted on traded bonds.
- Cost, in basis points, where a basis point is one hundredth of a per cent, equals ten, plus two for each million dollars of trade size, plus fifteen if the bond is high yield, meaning rated below investment grade, plus one and a half per year of maturity, plus two per year of age.
- Harrow wants to sell five million of a high yield bond that has not traded for months, with eight years to maturity, issued three years ago.
- The constant gives ten basis points.
- Size adds two times five, ten.
- The rating adds fifteen.
- Maturity and age add one and a half times eight plus two times three, eighteen.
- The total is fifty three basis points, so the trade costs twenty six thousand five hundred dollars.

22 Emergency tools for unforeseen redemptions

Each needs a basis in the fund documents

Tool	What it does
Suspension	Stops all dealing
Gate	Caps redemptions on a dealing date
Side pocket	Separates illiquid or hard-to-value assets
Redemption in kind	Pays investors in securities
Liquidity fees, swing pricing	Redeemers bear the cost of exit

IN THE LESSON

- Now suppose redemption requests arrive that nobody foresaw.
- Regulators can allow funds to use emergency tools, each with its own conditions, and the premise is that every one of them needs a basis in the fund documents, so investors know in advance they may be used.
- A suspension stops all dealing.
- A gate caps redemptions on a dealing date, usually as a percentage of N A V.
- A side pocket separates illiquid or hard to value assets from the rest of the fund.
- Redemption in kind pays investors in securities instead of cash, so the fund does not have to sell.
- And liquidity fees, or swing pricing, make redeemers bear the cost of their exit.
- Next, when each is permitted.

23 When each tool is permitted

- Suspension: exceptional circumstances only
- For example a market closure, or no way to value
- Serves investors as a whole; regulator informed
- Gate: heavy redemptions; pro rata or carry forward
- Side pocket: assets that cannot be valued or sold

IN THE LESSON

- A suspension is the most drastic tool, so it is reserved for exceptional circumstances.
- Examples are a market closure, when the fund's assets cannot be traded, or the impossibility of valuing the portfolio, when no fair N A V can be struck.
- It must serve the investors as a whole, not one group at another's expense, and the regulator is normally informed.
- A gate is used when redemptions on a dealing date are heavy relative to the fund.
- Requests above the cap are scaled back pro rata, meaning in proportion, or the excess is carried forward to the next dealing date.
- A side pocket is used when particular assets cannot be valued or sold, so that the rest of the fund can continue to deal.

24 When in kind and fees are permitted

- In kind: a cash sale would harm remaining holders
- Fees, swing pricing: redemptions cause dilution
- Both need support in the fund documents

IN THE LESSON

- Two tools remain.
- Redemption in kind is used when a cash sale would harm the remaining holders, for example when one large redeemer would force sales of hard to sell bonds.
- Liquidity fees or swing pricing are used when redemptions cause dilution, and they fall on those who leave.
- Both need support in the fund documents, so that investors know in advance which tools may be used and on what conditions.

25 Example 8: a gate at Harrow

NAV $500m. Gate: 10% of NAV per dealing date. Requests today: $80m

Gate cap	10% × 500	**$50m**
Share of each request paid	50 ÷ 80	**62.5%**
Investor asking $8m is paid	62.5% × 8	**$5m**
Carried forward	8 − 5	**$3m**
Under a suspension instead	all dealing stops	**$0 dealt**

IN THE LESSON

- Now a gate at Harrow.
- The fund's documents allow a gate of ten per cent of N A V on any dealing date.
- N A V is five hundred million dollars, and requests today total eighty million.
- The gate caps redemptions at ten per cent of five hundred, fifty million.
- So each request is scaled back pro rata: fifty over eighty, sixty two and a half per cent of each request is paid.
- An investor asking for eight million receives sixty two and a half per cent of eight, five million, and the remaining three million is carried forward to the next dealing date.
- Under a suspension instead, all dealing stops and nothing is dealt today.
- The usual mistake is confusing the two: a gate limits redemptions and pays part of each request, while a suspension pays none.

26 Example 9: a side pocket and redemption in kind

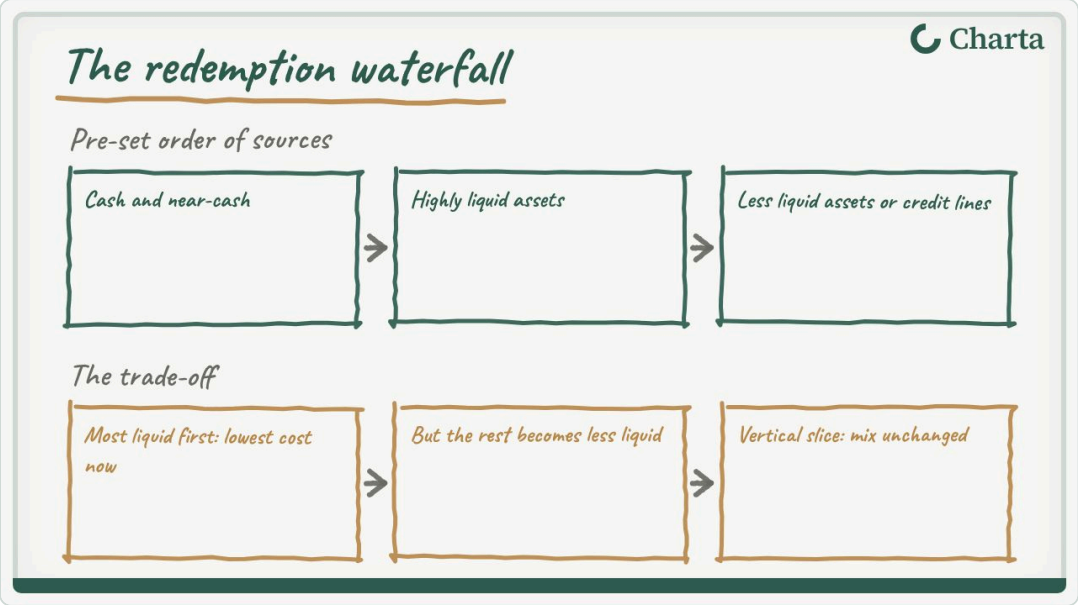
Defaulted bonds, carried at last value $20m, cannot be valued. Investor owns 1%

Side pocket	defaulted bonds	**$20m**
Main fund keeps dealing	500 − 20	**$480m**
Investor: main fund	1% × 480	**$4.8m**
Investor: side pocket	1% × 20; paid when sold	**$0.2m**
Or in kind: securities	a 1% slice of holdings	**no sale**

IN THE LESSON

- Last, a side pocket and redemption in kind.
- Suppose twenty million dollars of Harrow's bonds default and cannot be valued reliably, so Harrow carries them at their last value of twenty million.
- Harrow moves them into a side pocket, a separate part of the fund that does not deal.
- The main fund, five hundred minus twenty, four hundred and eighty million, keeps dealing at a fair price.
- An investor who owns one per cent of Harrow holds one per cent of the main fund, four point eight million, which can be redeemed as normal.
- They also hold one per cent of the side pocket, two hundred thousand dollars, paid out only when those bonds are sold or recovered.
- Alternatively, a large investor may be redeemed in kind, receiving a one per cent slice of the holdings as securities, so no sale is needed.

27 The redemption waterfall



IN THE LESSON

- Finally, redemption risk, which a fund handles by deciding in advance how redemptions will be funded and how large they could be.
- The first decision is the redemption waterfall: a pre set order of sources of cash.
- Follow the boxes.
- First, cash and near cash, such as short term deposits.
- Second, highly liquid assets, such as government bonds.
- Third, less liquid assets, or a credit line from a bank.
- The second row is the trade off.
- Paying from the most liquid assets first keeps the cost low now.
- But the rest of the portfolio becomes less liquid for the investors who remain.
- Selling a vertical slice, a proportional piece of every holding, keeps the mix unchanged, at a higher cost today.

28 Example 10: waterfall or vertical slice?

$m: cash 15 (free); liquid 85 (cost 0.1%); less liquid 400 (cost 0.8%). Redeem 75

Waterfall: sources used	15 cash + 60 liquid	**$75m**
Waterfall: cost	60 × 0.1%	**$60,000**
Waterfall: less liquid share	400 ÷ 425	**94.1%**
Slice: sell 15% of each	2.25 + 12.75 + 60	**$75m**
Slice: cost; less liquid share	12.75×0.1% + 60×0.8%	**$492,750; 80%**

IN THE LESSON

- Harrow, in millions of dollars, holds fifteen of cash, which costs nothing to use, eighty five of liquid bonds that cost point one per cent to sell, and four hundred of less liquid bonds that cost point eight per cent.
- Less liquid here is simply everything beyond cash and highly liquid bonds, a broader group than hard to sell.
- Redemptions of seventy five million arrive.
- Under the waterfall, Harrow uses all fifteen of cash and sixty of liquid bonds.
- The cost is point one per cent of sixty, sixty thousand dollars.
- But afterwards the less liquid bonds are four hundred out of four hundred and twenty five, ninety four point one per cent of the fund.
- Under a vertical slice, Harrow sells fifteen per cent of everything: two point two five of cash, twelve point seven five of liquid bonds and sixty of less liquid ones, seventy five in all.
- The cost is four hundred and ninety two thousand seven hundred and fifty dollars, but the less liquid share stays at eighty per cent.

29 Redemption-at-risk

RaR = mean + z × standard deviation
Over T periods: mean × T, SD × √T

Mean net redemption	Average net redemption
z	Normal quantile, e.g. 2.33 rounded
Standard deviation	How much net flows vary
T	Periods in the horizon

IN THE LESSON

- The second decision is how large redemptions could be.
- Redemption at risk, or R A R, is the net redemption over a horizon that is not expected to be exceeded at a chosen confidence level.
- Net means redemptions minus new subscriptions.
- If net flows are normally distributed, R A R equals the mean plus z times the standard deviation.
- The mean is the average net redemption per period.
- Z is the normal quantile for the confidence level, about two point three three, rounded, for ninety nine per cent.
- The standard deviation measures how much net redemptions vary.
- Over T periods, with independent flows, the mean scales with T and the standard deviation with the square root of T.

30 Example 11: Harrow's redemption-at-risk

Daily net redemptions: mean 0.2% of NAV, SD 1.5%. 99%: z = 2.33. NAV $500m

One day	0.2% + 2.33 × 1.5%	**3.695%**
Five days: mean; SD	5 × 0.2%; 1.5% × √5	**1.0%; 3.354%**
Five-day RaR	1% + 2.33 × 3.354%	**8.815% = $44.1m**
Wrong: scaled linearly	5 × 3.695%	**18.475%**
Wrong: mean ignored	2.33 × 3.354%	**7.815%**

IN THE LESSON

- Harrow's daily net redemptions average point two per cent of N A V, with a standard deviation of one point five per cent.
- At ninety nine per cent confidence, z is two point three three.
- For one day, R A R is point two plus two point three three times one point five, three point six nine five per cent.
- For five days, the mean is five times point two, one per cent, and the standard deviation is one point five times the square root of five, three point three five four per cent.
- Five day R A R is one plus two point three three times three point three five four, eight point eight one five per cent, or forty four point one million dollars of the five hundred million fund.
- Scaling the one day figure linearly would give eighteen point four seven five per cent.
- Ignoring the mean gives seven point eight one five.

31 Example 12: tuning liquidity to cover RaR

Buffer: cash $15m plus liquid bonds $85m after a 5% stress haircut

Buffer after haircut	15 + 85 × (1 − 5%)	**$95.75m**
Five-day RaR	8.815% × 500	**$44.1m: covered**
Twenty-day SD	1.5% × √20	**6.708%**
Twenty-day RaR	4% + 2.33 × 6.708%	**19.63% = $98.15m**
Shortfall	98.15 − 95.75	**$2.4m: tune**

Tune: hold more cash, lengthen notice, or cut less liquid assets

IN THE LESSON

- Tuning liquidity means holding enough cash and liquid assets, after stress haircuts, to cover R A R.
- A stress haircut is the cut in value expected when an asset is sold in stress.
- Harrow's buffer is fifteen million of cash plus eighty five million of liquid bonds after a five per cent haircut: ninety five point seven five million.
- The five day R A R of forty four point one million is covered.
- Now a twenty day horizon.
- The mean is twenty times point two, four per cent, and the standard deviation is one point five times the square root of twenty, six point seven zero eight per cent.
- Twenty day R A R is nineteen point six three per cent, ninety eight point one five million.
- The shortfall is two point four million, so Harrow must tune: hold more cash, lengthen its notice period, or cut less liquid assets.

32 Recap: Harrow's liquidity toolkit

- Control: matched terms, limits, stress tests, swing pricing
- Bonds: OTC, rare trades, smoothed prices, biased data
- Measures: spread, Roll, Amihud, zero days, t-cost model
- Emergencies: suspend, gate, side pocket, in kind, fees
- Redemptions: waterfall, RaR, buffer after haircuts

IN THE LESSON

- Let us pull it together.
- To control liquidity risk, a fund matches its redemption terms to its assets, sets limits on hard to sell assets and investor concentration, stress tests both sides, and uses swing pricing against dilution.
- Bond funds find this harder, because bonds trade over the counter and rarely, evaluated prices are smoothed, and trade data are biased toward small, calm trades.
- Liquidity is measured with the spread, Roll's estimator, the Amihud ratio, zero trading days and a t cost model.
- In an emergency, regulators can allow a suspension, a gate, a side pocket, redemption in kind or liquidity fees.
- And redemptions are planned with a waterfall, redemption at risk, and a buffer counted after stress haircuts.

33 Try three quick questions

1. Price changes in %; serial covariance −0.04. Roll's spread?
   Answer: √0.04 = 0.2%; × 2 = 0.4%.
2. Net redemptions: mean 0.1%, SD 2% a day. 95%: z = 1.645. Four-day RaR?
   Answer: 0.4% + 1.645 × 4% = 6.98%.
3. Selling cost $0.6m; NAV left after redemptions $240m. Dilution?
   Answer: 0.6 ÷ 240 = 0.25%.

IN THE LESSON

- Now try three quick questions.
- Pause the video, think, then listen to the answers.
- Question one.
- Daily price changes are measured in per cent, and their serial covariance is minus point zero four.
- What is Roll's estimate of the spread?
- The square root of point zero four is point two per cent, and times two the spread is point four per cent.
- Question two.
- Daily net redemptions have a mean of point one per cent and a standard deviation of two per cent.
- At ninety five per cent confidence, z is one point six four five.
- What is the four day redemption at risk?
- Over four days the mean is point four per cent and the standard deviation four per cent, so the answer is six point nine eight per cent.
- Question three.
- Selling costs are six hundred thousand dollars, and the N A V left after redemptions is two hundred and forty million.
- What is the dilution?
- Point two five per cent.