CHARTA LEARNING · LESSON NOTES

Illiquid Assets

FRM Part II

1 A story: the plan that could not sell

- A crisis: share prices fall, pensions must still be paid
- Listed shares and bonds sell within a day, at low prices
- Private real estate stake: few buyers, deep discounts
- Its reported value barely moves: calm on paper
- Why illiquid, why so smooth, what it means for a portfolio

IN THE LESSON

- Imagine a crisis year for Linden.
- Share prices fall sharply, and pensions must still be paid, so the plan needs cash.
- Its listed shares and bonds can be sold within a day, though at low prices.
- Its stake in a private real estate fund cannot.
- There is no exchange, there are few buyers, and the only offers come at a deep discount to the last reported value.
- Stranger still, the fund's reported value has barely moved, so on paper the private holdings look calm while everything else is falling.
- This lesson explains each part of that story: why some markets are illiquid, why their reported returns look so smooth, and what that means for a portfolio.

2 What makes a market illiquid

- Illiquid: cannot sell quickly, in size, near fair value
- Quick sale: low price. Fair price: a long wait
- Infrequent trading
- Wide bid-ask spreads and high transaction costs
- Costly search: bilateral, over-the-counter trades

IN THE LESSON

- A market is illiquid when an investor cannot turn an asset into cash quickly, in size, and at a price close to fair value.
- All three parts matter.
- You can sell a house quickly if you accept a low price, and you can get a fair price if you wait for months.
- Illiquidity shows up in typical signs.
- Trading is infrequent.
- Bid ask spreads are wide, the spread being the gap between the price a dealer pays and the price a dealer charges, and other transaction costs are high.
- Finding a counterparty is costly, because trades are bilateral and over the counter, that is, negotiated privately between two parties rather than on an exchange.

3 More signs, and where values come from

- Assets are unique, not standardised
- A large order has price impact
- Few observed prices: values come from appraisals or models
- Liquidity varies over time
- In a crisis it vanishes in many markets at once

IN THE LESSON

- There are more signs.
- Illiquid assets are unique rather than standardised.
- One office building is not interchangeable with another, unlike one share of a listed company and the next.
- A large order has price impact, which means the act of selling pushes the price down.
- Because trades are rare, observed prices are scarce, so values come from appraisals, which are expert estimates, or from models.
- And liquidity is not fixed.
- It varies over time.
- In a crisis it tends to vanish in many markets at once, just when investors most need it.
- That last feature is what hurt Linden in our story.

4 Three dimensions of liquidity

A liquid market is tight, deep and resilient: all three

Dimension	What it measures
Tightness	The bid-ask spread
Depth	Size that trades without moving the price
Resiliency	How fast the price recovers after a trade
Exam point	Narrow spread, thin depth: not liquid

IN THE LESSON

- A liquid market is tight, deep and resilient, all three at once.
- The table names each dimension.
- Tightness is the bid ask spread: the gap between the bid and the ask.
- The mid price is halfway between them.
- Depth is the size that can trade without moving the price.
- Resiliency is how fast the price recovers after a trade has pushed it away.
- Now the exam point for this outcome.
- A market can be tight but shallow, with a narrow quoted spread that is good only for a small amount.
- So a narrow quoted spread does not prove a market is liquid if depth is thin.
- The next example puts money on that.

5 Example 1: a narrow spread, a thin book

Linden sells $50m of a bond. Bid 99.0, ask 99.4, per 100 of face value

Quoted spread and mid	99.4 − 99.0; mid 99.2	**0.4 points**
Depth at the bid	only $10m of the $50m	**$10m**
Average sale price	10 at 99, 15 at 98.5, 25 at 97.6	**98.15**
Actual cost against mid	(99.2 − 98.15)% × $50m	**$525,000**
Cost the spread suggests	half-spread 0.2% × $50m	**$100,000**

Actual cost is 5.25 times what the quoted spread suggests

IN THE LESSON

- Linden must sell fifty million dollars of face value of a corporate bond.
- Prices are quoted per one hundred of face value: the bid is ninety nine, and the ask is ninety nine point four.
- The quoted spread is only point four, and the mid price is ninety nine point two.
- But depth at the bid is only ten million dollars.
- To sell the rest, Linden must accept lower bids: ten million at ninety nine, fifteen million at ninety eight point five, and twenty five million at ninety seven point six.
- The average price is ninety eight point one five.
- Against the mid, that costs one point zero five per cent of fifty million, or five hundred and twenty five thousand dollars.
- The quoted half spread of point two per cent suggests only one hundred thousand dollars.
- The true cost is five point two five times larger, because depth was thin.

6 Where illiquidity comes from

- Frictionless: any quantity trades instantly at one price
- Search frictions: finding a counterparty takes time
- Asymmetric information: some traders know more
- Transaction costs and limited dealer capital
- Funding constraints: borrowers can be forced to sell

IN THE LESSON

- Why does illiquidity exist at all?
- Start with a frictionless market, the ideal of textbook models, in which any quantity trades instantly at one price.
- There, illiquidity cannot exist.
- It appears only when imperfections are added, and four matter.
- Search frictions: in over the counter markets, finding a counterparty takes time.
- Asymmetric information: some traders know more than others.
- Transaction costs and limited dealer capital: trading is costly, and dealers can carry only so much.
- And funding constraints: investors who borrow to hold assets can be forced to sell.
- Each imperfection produces its own symptom, and the exam asks you to match them.

7 Search frictions and dealer capital

- Search: finding a buyer takes time
- Dealers supply immediacy and gain bargaining power
- So trading is slow, and dealers charge more
- Limited dealer capital: small inventory, thin depth
- Dealers lose capital: depth shrinks further

IN THE LESSON

- Take search frictions first.
- When a trade must be arranged privately, the seller has to find a buyer, and that takes time.
- A dealer who can supply immediacy, meaning an instant trade, gains bargaining power over a seller in a hurry.
- So trading is slow and costly, and dealers can charge more.
- Next, transaction costs and limited dealer capital.
- A dealer buys a block, holds it in inventory and hopes to sell it later at a profit.
- Its capital limits how much it can carry, so the volume the market can absorb near a fair price is small.
- That is thin depth, the problem we met in Linden's bond sale.
- And when dealers lose capital, depth shrinks further.

8 Example 2: adverse selection widens the spread

Mid 100. Informed traders know the value is 105 or 95. 20% of traders are informed

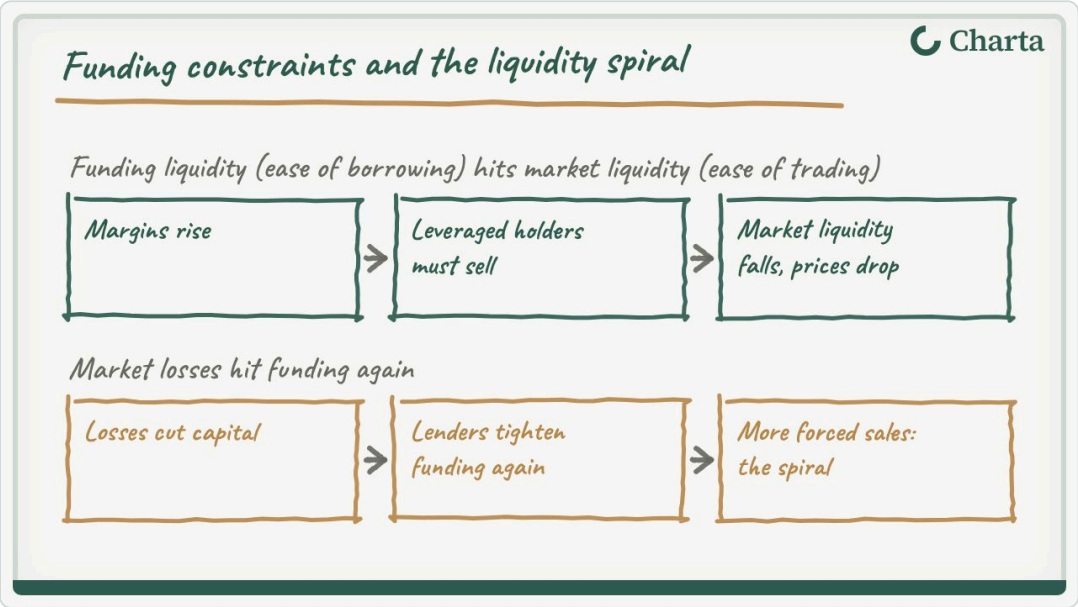
Loss to an informed buyer	105 − (100 + h)	**5 − h**
Gain from an uninformed buyer	pays 100 + h, worth 100	**h**
Break even	0.8 × h = 0.2 × (5 − h)	**h = 1**
Quoted bid and ask	100 − 1 and 100 + 1	**99 / 101**
If 40% are informed	0.6 × h = 0.4 × (5 − h)	**h = 2**

Spread rises from 2 to 4. In the extreme, buyers withdraw altogether

IN THE LESSON

- Now asymmetric information, worked through as a simple break even.
- A dealer quotes around a mid price of one hundred.
- Some traders are informed: they know the asset is really worth either one hundred and five or ninety five.
- Twenty per cent of traders are informed.
- The dealer sets the ask at one hundred plus a half spread, h.
- When an informed trader buys, the asset is worth one hundred and five, so the dealer loses five minus h.
- When an uninformed trader buys, the dealer gains h.
- To break even, point eight times h must equal point two times five minus h, so h is one.
- The dealer quotes ninety nine bid and one hundred and one ask.
- If forty per cent are informed, h is two, and the spread doubles from two to four.
- This is adverse selection.
- In the extreme, buyers withdraw altogether.

9 Funding constraints and the liquidity spiral



IN THE LESSON

- The fourth imperfection links two kinds of liquidity.
- Market liquidity is how easily an asset trades.
- Funding liquidity is how easily a trader can borrow against it.
- The first row reads: funding liquidity hits market liquidity.
- Margins rise, meaning lenders demand more of the trader's own money for each dollar of assets.
- Leveraged holders, those investing with borrowed money, must sell.
- Market liquidity falls and prices drop.
- The second row reads: market losses hit funding again.
- Losses cut the holders' capital.
- Lenders tighten funding again.
- And more forced sales follow: the spiral.
- Each turn makes the next one worse, which is why liquidity can vanish in many markets at once.

10 Example 3: a margin call turns into a fire sale

A fund holds $100m of bonds with $10m of its own money: a 10% margin

Prices fall 2%	100 × 2%	**loss $2m**
Own money left	10 − 2	**$8m**
Most it can hold at 10%	8 ÷ 10%	**$80m**
Most it can hold at 16%	8 ÷ 16%	**$50m**
Forced sale at 16% margin	98 − 50	**$48m**

Funding constraint, higher margin: sells $48m, not $18m

IN THE LESSON

- Here is one turn of the spiral.
- A leveraged hedge fund holds one hundred million dollars of the same bond that Linden sold, with ten million dollars of its own money, so the margin, its own money as a share of the position, is ten per cent.
- Prices fall two per cent, a loss of two million.
- Its own money falls from ten million to eight.
- Its holding is now worth ninety eight million.
- If the margin stayed at ten per cent, eight million would support eighty million of bonds, so the fund would sell eighteen million.
- But lenders raise the margin to sixteen per cent.
- Now eight million supports only fifty million, and the forced sale is ninety eight minus fifty, forty eight million dollars.
- That selling pushes prices down again, and the next turn begins.
- The specific imperfection here is the funding constraint: the higher margin.

11 Match the symptom to the imperfection

Exam point: name the imperfection behind each symptom

Symptom	Imperfection behind it
Wide spreads	Adverse selection: informed traders
Slow, costly to arrange a trade	Search frictions in OTC markets
Large trades move the price	Costs and limited dealer capital
Fire sales that feed on themselves	Funding constraints: margins
Buyers withdraw altogether	Extreme asymmetric information

IN THE LESSON

- Here is the exam point for this outcome: name the specific imperfection behind each symptom.
- Read the table row by row.
- Wide spreads come from adverse selection, the dealer's fear of informed traders.
- Trades that are slow and costly to arrange come from search frictions in over the counter markets.
- Large trades that move the price come from transaction costs and limited dealer capital.
- Fire sales that feed on themselves come from funding constraints, above all rising margins.
- And buyers who withdraw altogether reflect asymmetric information in its extreme form.
- If a question describes a symptom, find its row, and you have the cause.

12 Why reported returns are biased

- Reported returns are biased, usually flatteringly
- Smoothing: appraisals anchor on the last value
- Stale prices: the last trade may be months old
- A sudden fall shows up bit by bit
- Reported series: a moving average of the truth

IN THE LESSON

- Now to the returns that illiquid assets report.
- They are biased, usually in a flattering direction.
- The first bias is smoothing.
- Linden's property fund is valued by appraisal, often only once a year, and appraisers anchor on the last value and adjust it slowly.
- Some holdings are carried at stale prices, meaning the last trade, which may be months old.
- So a sudden fall in true value shows up in reported value only bit by bit, over several periods.
- The reported series behaves like a moving average of the truth: its ups and downs are trimmed, and its moves run on in the same direction.

13 Example 4: smoothing a crash

Linden's property fund: true returns −20%, 0%, 0%. The appraiser moves half-way

Year 1 reported	0.5 × (−20) + 0.5 × 0	**−10.0%**
Year 2 reported	0.5 × 0 + 0.5 × (−10)	**−5.0%**
Year 3 reported	0.5 × 0 + 0.5 × (−5)	**−2.5%**
Worst year, true and reported	−20% and −10%	**half the drop**
Three reported years added	−10 − 5 − 2.5	**−17.5%**

IN THE LESSON

- Here is smoothing at work in Linden's property fund.
- Suppose the true return is minus twenty per cent in year one, then zero in years two and three.
- The appraiser closes only half the gap each year, so the reported return is half this year's true return plus half last year's reported return, which was zero.
- Year one reports half of minus twenty plus half of zero, minus ten per cent.
- Year two reports half of zero plus half of minus ten, minus five.
- Year three reports minus two point five.
- In the worst year the true return is minus twenty, but the reported one is minus ten: half the drop.
- Adding the three reported years gives minus seventeen point five, creeping toward the true twenty.

14 What smoothing does to risk measures

- Reported volatility is understated
- Positive autocorrelation: −10, −5, −2.5
- Correlation and beta with liquid markets understated
- Sharpe ratio overstated: volatility too small
- Smoothing hides risk; the average is unchanged

IN THE LESSON

- What does smoothing do to the numbers an investor relies on?
- Four things.
- Reported volatility is understated, because the swings are trimmed.
- Returns show positive autocorrelation, meaning a return tends to be followed by one of the same sign, like Linden's minus ten, minus five, minus two point five.
- Correlation and beta with liquid markets are understated, because when shares crash the appraisal reacts only partly in the same period.
- And the Sharpe ratio, the excess return over the risk free rate divided by volatility, is overstated, because its denominator is too small.
- Put simply, smoothing hides risk.
- On its own, it leaves the average return unchanged.

15 Three biases that inflate the average

Database averages overstate what investors earned

Bias	How it arises
Survivorship	Failed funds drop out of the database
Selection	Managers choose whether to report
Backfill (instant history)	A joiner's earlier, good record is added

IN THE LESSON

- Database averages overstate what investors earned.
- Fund returns are collected in databases that managers report to voluntarily, and three biases push the average return itself above the truth.
- Survivorship bias: funds that fail drop out of the database, so the average of the survivors overstates what an investor in the full group earned.
- Selection bias: managers choose whether to report at all, and those with good results are keener to do so.
- Backfill bias, also called instant history bias: when a fund joins a database, its earlier record is added, and that record is usually good, because funds tend to join after a good run.

16 Example 5: the funds that disappeared

A database starts with 10 property funds. After five years, 3 have failed and dropped out

Survivors' average return	7 funds still reporting	**9.0%**
Failed funds' average return	3 funds, dropped out	**−8.0%**
Full group average	(7 × 9 + 3 × (−8)) ÷ 10	**3.9%**
Survivorship bias	9.0 − 3.9	**5.1 points**

Database shows 9.0%; an investor in all ten earned 3.9%

IN THE LESSON

- Linden is choosing a property manager, and looks at a database of property funds.
- Ten funds started reporting five years ago.
- Three have since failed and dropped out.
- The seven survivors earned an average of nine per cent a year, and that is the figure the database shows.
- The three failed funds earned an average of minus eight per cent before they vanished.
- An investor who had spread money across all ten earned seven times nine, plus three times minus eight, divided by ten, which is three point nine per cent.
- The survivorship bias is nine minus three point nine, five point one percentage points.
- The database average flatters the group by more than the whole of the true return.

17 Consequences, and the exam point

- High return, low volatility, low correlation
- An apparent free lunch in diversification
- So too large an allocation to illiquid assets
- Exam: smoothing hides risk
- Survivorship, selection, backfill inflate the average

IN THE LESSON

- Put the biases together and the picture is dangerously attractive.
- Returns look high, volatility looks low, and correlation with shares looks small.
- That creates an apparent free lunch in diversification: an asset that seems to add return while cutting portfolio risk.
- An optimiser fed these numbers puts too large an allocation into illiquid assets.
- Linden's board, shown a smooth nine per cent from a database of survivors, could easily be tempted.
- Remember the exam point for this outcome.
- Smoothing hides risk, while survivorship, selection and backfill inflate the average return.
- The next outcome shows how to undo the first of these.

18 The smoothing model

r(t) = (1 − c) × T(t) + c × r(t − 1)
c = first-order autocorrelation of r (true returns independent)

Reported, r(t)	This period's reported return
Smoothing coefficient, c	Between 0 and 1; 0 is no smoothing
True return, T(t)	This period's true return
Last reported return	r(t − 1), last period's

IN THE LESSON

- Unsmoothing tries to recover the true return from the smoothed reported one.
- It starts from a model of how smoothing happens.
- The reported return, r of t, is a mix.
- It equals one minus c, times the true return, T of t, plus c times last period's reported return, r of t minus one.
- The letter c is the smoothing coefficient, between zero and one.
- If c is zero there is no smoothing, and reported equals true.
- The larger c is, the more weight falls on the past.
- In Linden's example, c was one half.
- How do we find c?
- If true returns are independent, it equals the first order autocorrelation of reported returns: the correlation of each return with the one before.

19 Unsmoothing, and the volatility factor

T(t) = [r(t) − c × r(t − 1)] ÷ (1 − c)
True vol = reported vol × √((1 + c) ÷ (1 − c))

True return, T(t)	Unsmoothed, or true, return
Smoothing coefficient, c	Divide by 1 − c
Reported volatility	Understated by smoothing
Volatility factor	Factor: at least 1, rises with c

IN THE LESSON

- Now solve the model for the true return.
- Move c times last period's reported return to the other side, and divide by one minus c.
- The unsmoothed return equals r of t minus c times r of t minus one, all divided by one minus c.
- Divide by one minus c, never by one plus c.
- Unsmoothing leaves the average return roughly unchanged, but it raises volatility.
- If true returns are independent from one period to the next, true volatility equals reported volatility times a factor.
- Reported volatility is understated by smoothing, and the factor is the square root of one plus c, over one minus c.
- That factor is always at least one, and it rises fast as c rises.

20 Example 6: unsmoothing Linden's returns

Linden reports −10%, −5%, −2.5%; it reported 0% the year before. c = 0.5

Year 1 true	(−10 − 0.5 × 0) ÷ 0.5	**−20%**
Year 2 true	(−5 − 0.5 × (−10)) ÷ 0.5	**0%**
Year 3 true	(−2.5 − 0.5 × (−5)) ÷ 0.5	**0%**
Wrong: divide by 1 + c	−10 ÷ 1.5	**−6.7%**

Unsmoothing recovers the true path: −20%, 0%, 0%

IN THE LESSON

- Let us unsmooth Linden's reported returns of minus ten, minus five and minus two point five per cent, with c of one half and a reported return of zero the year before.
- Year one: minus ten minus half of zero, divided by one half, gives minus twenty per cent.
- Year two: minus five minus half of minus ten is zero, so the true return is zero.
- Year three: minus two point five minus half of minus five is also zero.
- Unsmoothing has recovered the true path exactly: minus twenty, then zero, then zero.
- Now the common mistake.
- Dividing by one plus c, year one would give minus ten over one point five, about minus six point seven per cent, which makes the crash look even smaller instead of larger.

21 Example 7: true volatility and the Sharpe ratio

Linden's property fund reports excess return 5% a year and volatility 8%. c = 0.5

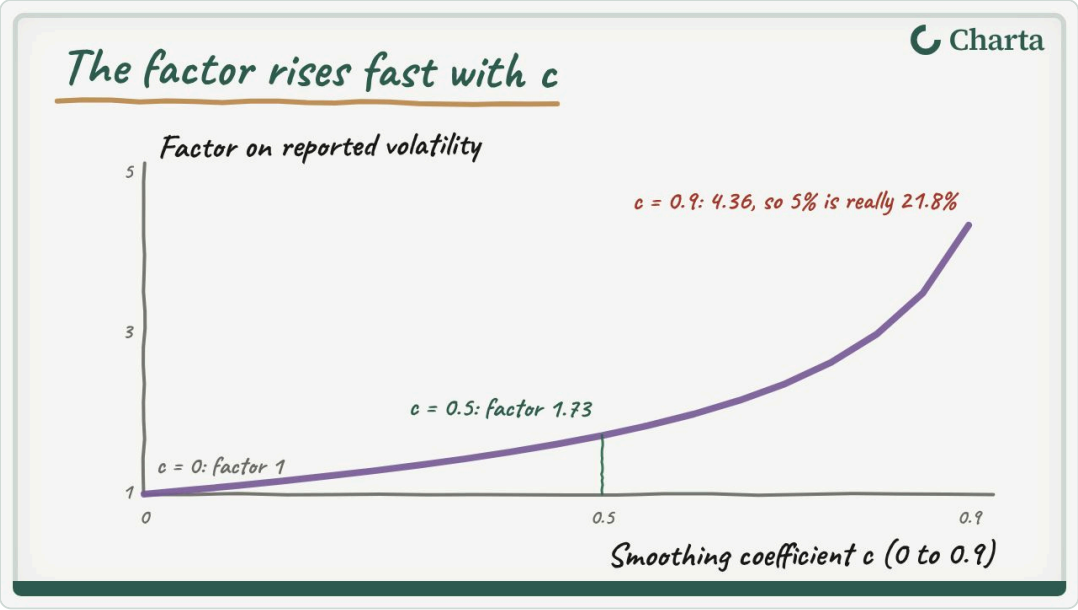
Volatility factor	√((1 + 0.5) ÷ (1 − 0.5)) = √3	**1.732**
True volatility	8% × 1.732	**13.86%**
Sharpe ratio, reported	5 ÷ 8	**0.625**
Sharpe ratio, unsmoothed	5 ÷ 13.86	**0.36**
Wrong: no square root	8% × 3	**24%**

Mean unchanged; measured risk up 73%; the Sharpe ratio falls

IN THE LESSON

- Over a long record, Linden's property fund reports an average excess return of five per cent a year, with volatility of eight per cent, and its smoothing coefficient is one half.
- The factor is the square root of one point five over point five, the square root of three, which is one point seven three two.
- True volatility is eight times one point seven three two, thirteen point eight six per cent.
- The average excess return stays at five.
- So the reported Sharpe ratio, five over eight, point six two five, falls to five over thirteen point eight six, about point three six.
- Measured risk has risen by about seventy three per cent.
- The other common mistake is to forget the square root, which gives eight times three, or twenty four per cent.

22 The factor rises fast with c



IN THE LESSON

- This chart shows how far reported volatility understates true volatility.
- Along the bottom is the smoothing coefficient c, from zero to point nine.
- Up the side is the factor that multiplies reported volatility, starting from one.
- At c of zero the factor is one: no smoothing, so no correction.
- At Linden's c of one half, the factor is one point seven three.
- Then the curve steepens.
- At c of point nine the factor is about four point three six, so an asset reporting five per cent volatility would really have about twenty one point eight per cent.
- Small changes in an already high c make a big difference, which is why c must be estimated with care.

23 What unsmoothing changes

- Average return: roughly unchanged
- Volatility, correlation and beta all rise
- Only (1 − c) of today's market move shows today
- Beta 0.3 reported, c = 0.5: true ≈ 0.3 ÷ 0.5 = 0.6
- Sharpe ratio falls; diversification looks smaller

IN THE LESSON

- Step back and list what unsmoothing does.
- The average return is left roughly unchanged, because over a long record the reported and true series add up to nearly the same total.
- Volatility, correlation and beta with liquid markets all rise.
- Here is a quick way to see it for beta.
- Only one minus c of today's market move shows up in today's reported return.
- So if Linden's fund reports a beta to shares of point three, with c of one half the true beta is roughly point three divided by point five, which is point six.
- Finally, the Sharpe ratio falls, and diversification looks smaller.
- The exam point: unsmoothing raises risk measures and lowers apparent diversification, with the mean unchanged.

24 What an illiquidity premium is

- Extra expected return over an identical liquid asset
- Compensates for trading costs: spread, price impact
- And the risk of being unable to sell when cash is needed
- Within a category: a clean comparison
- Across categories: leverage, factors, credit risk, biases

IN THE LESSON

- Now to the reward for bearing illiquidity.
- An illiquidity premium is the extra expected return an asset must offer over an otherwise identical liquid asset.
- Investors demand it for two reasons.
- First, trading costs: every purchase and sale loses the spread and the price impact.
- Second, the risk of being unable to sell when cash is needed, exactly what Linden faced in its crisis year.
- Comparing premiums is clean within a single category of asset.
- Across categories, such as private equity, real estate and hedge funds, it is much harder, because the assets also differ in leverage, in factor exposure, in credit risk and in the biases of their reported returns.

25 Example 8: premiums within a category

Same issuer or similar risk: only the ease of trading differs

On-the-run 10-year Treasury	newest, most traded	**yield 4.00%**
Off-the-run, similar maturity	older, rarely traded	**yield 4.05%**
Premium within Treasuries	4.05 − 4.00	**0.05 points**
Liquid stock, expected return	frequently traded	**8.5%**
Illiquid stock, expected return	similar risk, rarely traded	**10.0%**

Illiquid stock premium: 10.0 − 8.5 = 1.5 points

IN THE LESSON

- Within a category, the comparison is clean, because the only real difference is how easily the asset trades.
- Take United States Treasuries.
- The on the run issue is the newest and most traded bond of a given maturity.
- Suppose the on the run ten year Treasury yields four per cent.
- An older, off the run issue of similar maturity, rarely traded, yields four point zero five per cent.
- Same issuer, same credit, so the gap of point zero five percentage points is an illiquidity premium.
- Shares work the same way.
- Suppose a liquid stock is expected to return eight point five per cent, and an illiquid stock of similar risk, rarely traded, ten per cent.
- Its premium is ten minus eight point five, one point five percentage points.

26 Example 9: premiums across categories

Linden's private equity reported 13% a year; listed equity returned 9%

Raw return gap	13 − 9	**4.0 points**
Less extra leverage	more debt in the companies	**−1.5**
Less factor tilt	smaller, riskier companies	**−1.0**
Less biases in the record	survivorship, backfill	**−1.0**
Illiquidity premium left	4.0 − 1.5 − 1.0 − 1.0	**0.5 points**

IN THE LESSON

- Across categories, the raw numbers mislead.
- Linden's private equity has reported thirteen per cent a year, against nine per cent for listed equity.
- The raw gap is four percentage points, but it is not all an illiquidity premium.
- Suppose Linden's analysts estimate the other sources.
- The private companies carry more debt, so extra leverage explains one and a half points.
- They tilt toward smaller, riskier companies, a factor tilt worth one point.
- And biases in the reported record, such as survivorship and backfill, add another point.
- What is left is four minus one and a half, minus one, minus one: half a percentage point.
- That is the pure illiquidity premium, far smaller than the raw gap.

27 Who holds illiquid assets, and why it matters

- Premium per unit of illiquidity: smaller for illiquid assets
- Round trip 6%, held 10 years: 6 ÷ 10 = 0.6% a year
- A 1-year holder bears the full 6% that year
- Patient owners hold them and accept less
- Exam: strip out other risks and biases first

IN THE LESSON

- A surprising result: the premium per unit of illiquidity tends to be smaller for the most illiquid assets.
- The reason is who holds them.
- Suppose an asset costs six per cent to buy and later sell, a round trip.
- A long horizon investor like Linden, holding for ten years, spreads that cost over the decade, six divided by ten, about point six per cent a year.
- A trader holding for one year bears the full six per cent in that year.
- So illiquid assets end up with patient owners, who trade rarely and accept less for each unit of illiquidity.
- The exam point: strip out other risks and biases before calling a return gap an illiquidity premium.

28 What illiquid holdings do to a portfolio

- Cannot be rebalanced cheaply or quickly
- Weights drift with relative performance
- Liquid assets fall: the illiquid share rises unchosen
- Rebalancing and cash needs fall on liquid assets
- Liquid sleeve: more volatile and sale-prone

IN THE LESSON

- Finally, what happens when illiquid assets go into a portfolio?
- They change how it is run.
- They cannot be rebalanced cheaply or quickly, so their weights drift with relative performance.
- When liquid assets fall, the illiquid share rises without the investor choosing it.
- And because the illiquid assets cannot be traded, any rebalancing and any cash need falls on the liquid assets.
- The liquid sleeve, the part of the portfolio that can be traded, ends up more volatile and more likely to be sold than its weight suggests.
- Let us see this in Linden's numbers.

29 Example 10: the share that drifts up

Linden: $2,000m. Liquid $1,500m; private $500m, its 25% target

Liquid assets fall 20%	1,500 × 0.8	**$1,200m**
Private, reported down 4%	500 × 0.96	**$480m**
Private share now	480 ÷ (1,200 + 480)	**28.6%**
Pay $80m pensions from liquid	1,200 − 80	**$1,120m**
Private share after paying	480 ÷ (1,120 + 480)	**30.0%**

IN THE LESSON

- Linden starts with two thousand million dollars: one thousand five hundred million in liquid assets and five hundred million in private assets, its target of twenty five per cent.
- In the crisis year, liquid assets fall twenty per cent, to one thousand two hundred million.
- The private assets report a fall of only four per cent, to four hundred and eighty million, a figure flattered by smoothing.
- The private share is now four hundred and eighty over one thousand six hundred and eighty, twenty eight point six per cent.
- Then Linden pays eighty million dollars of pensions, and only the liquid assets can pay them, leaving one thousand one hundred and twenty million.
- The private share rises to four hundred and eighty over one thousand six hundred, thirty per cent, though Linden never chose it.

30 Example 11: rebalance now, or wait?

Private share 30% of $1,600m. Target 25%, band 20% to 30%. Sales at a 15% discount

Private above target	480 − 25% × 1,600	**$80m**
Discount lost if sold now	15% × $80m	**$12m**
Inside the band?	30% is the upper edge	**yes, just**
Policy	wide band, trade rarely	**hold**

IN THE LESSON

- Should Linden sell private assets to get back to twenty five per cent?
- Its policy sets a tolerance band of twenty to thirty per cent, and private stakes can be sold only on the secondary market, to another investor, at an assumed discount of fifteen per cent.
- At twenty five per cent of one thousand six hundred million, the target holding is four hundred million, so the private assets are eighty million above target.
- Selling eighty million at a fifteen per cent discount would lose twelve million dollars.
- Is thirty per cent inside the band?
- Yes, just, at the upper edge.
- So the policy is to hold.
- With high trading costs, sensible policy is wide tolerance bands and infrequent trading.

31 How large should the allocation be?

- Smaller than a frictionless model gives
- Keep a liquidity buffer for obligations
- Cannot exit in a bad market
- Larger for long horizons with few liquidity needs
- Premium weighed against lost rebalancing flexibility

IN THE LESSON

- How large should the illiquid allocation be?
- Smaller than a frictionless model would give, for two reasons.
- First, the investor must keep a liquidity buffer: enough liquid assets to meet its obligations, such as Linden's pensions, even after a crash.
- Second, it cannot exit in a bad market, so it cannot cut risk when it most wants to.
- The allocation can be larger for a long horizon investor with few liquidity needs.
- And it is smaller once risk is measured with unsmoothed volatility and correlation.
- The exam point: the illiquidity premium must be weighed against lost rebalancing flexibility.

32 Example 12: sizing with unsmoothed risk

One-asset rule: weight rises with excess return ÷ variance. Reported risk gives 60%

Variance, reported	0.08 × 0.08	**0.0064**
Variance, unsmoothed	0.0064 × 3	**0.0192**
Weight, unsmoothed	60% × 0.0064 ÷ 0.0192	**20%**
Then: a liquidity buffer	pensions must be paid	**lower still**

Unsmoothed risk cuts the weight to a third: from 60% to 20%

IN THE LESSON

- Here is a rule from portfolio theory: the weight in a risky asset rises with its excess return divided by its variance, the square of its volatility.
- A simple one asset rule, using reported risk and a stated risk aversion, gives sixty per cent in this fund.
- This is a textbook illustration, not Linden's actual weight.
- The reported variance is point zero eight times point zero eight, point zero zero six four.
- Unsmoothing multiplies volatility by the square root of three, so it multiplies variance by three, to point zero one nine two.
- With the same excess return, the weight falls in proportion: sixty per cent times point zero zero six four over point zero one nine two, which is twenty per cent.
- Then Linden must keep a liquidity buffer for the pensions it pays, so the weight is lower still.

33 Mistakes to avoid

- A narrow quoted spread is not proof: check depth
- Reported volatility and correlation understate risk
- Unsmooth with (1 − c); take the root for volatility
- A raw gap across categories is not the premium
- The illiquid share drifts up when liquid assets fall

IN THE LESSON

- Before the questions, here are the mistakes to avoid.
- First, a narrow quoted spread is not proof of liquidity: always check depth.
- Second, the reported volatility of an appraised asset understates its risk, and its reported correlation understates how it moves with shares.
- Third, when unsmoothing, divide by one minus c, and take the square root for the volatility factor.
- Fourth, do not call a raw return gap across categories an illiquidity premium until leverage, factor tilts and biases are stripped out.
- Fifth, expect the illiquid share to drift up when liquid assets fall, and plan the liquid sleeve for it.

34 Try three quick questions

1. Reported volatility 4%, c = 0.8. True volatility?
   Answer: √(1.8 ÷ 0.2) = √9 = 3, so 4% × 3 = 12%.
2. Reported 2% now, 4% last period, c = 0.6. Unsmoothed return?
   Answer: (2 − 0.6 × 4) ÷ 0.4 = −0.4 ÷ 0.4 = −1%.
3. Dealers widen spreads for fear of informed traders. Imperfection?
   Answer: Asymmetric information: adverse selection.

IN THE LESSON

- Now try three quick questions.
- Pause the video, think, then listen to the answers.
- Question one.
- An appraised fund reports volatility of four per cent, and its smoothing coefficient is point eight.
- What is its true volatility?
- The factor is the square root of one point eight over point two, the square root of nine, which is three.
- True volatility is twelve per cent.
- Question two.
- A fund reports two per cent this period and four per cent last period, with c of point six.
- What is the unsmoothed return?
- Two minus point six times four is minus point four, divided by point four, which is minus one per cent.
- Question three.
- Dealers widen spreads because they fear better informed traders.
- Which imperfection is this?
- Asymmetric information, which causes adverse selection.