CHARTA LEARNING · LESSON NOTES

Performance and returns of alternative investments

Alternative Investments

1 A story: one fund, two honest answers

- Start: £10m. Year 1: fund earns 10%
- Client adds £90m. Year 2: fund loses 5%
- Manager says +4.5%; investors say −3.7%
- Both are right: different questions

IN THE LESSON

- Harbor Capital is an investment firm, and it runs our funds in this lesson.
- The numbers are made up, round examples.
- Start with its growth fund.
- It begins the year with ten million pounds.
- In year one the fund earns ten per cent.
- Then a big client adds ninety million pounds.
- In year two the fund loses five per cent.
- So, how did the fund do?
- The manager says, plus four and a half per cent.
- The clients, who put in a hundred million and got back less, say minus three point seven.
- Both are telling the truth.
- They are answering two different questions, and this video shows you which question to ask.

2 Time-weighted return: strip out the timing

TWR = $(1+R_1) \times (1+R_2) \times \ldots \times (1+R_n) - 1$

Sub-period return (R)	Return in each sub-period
Number of sub-periods (n)	Split at every cash flow
Time-weighted return	Chained return, timing removed

IN THE LESSON

- The first measure is the time-weighted return.
- The idea is simple.
- Cut the story into sub-periods, cutting every time money goes in or out.
- In each sub-period, work out the plain return.
- Then chain the sub-periods together by multiplying growth factors.
- The letter R is the return in one sub-period, written as a decimal.
- A growth factor is one plus that return.
- The letter n is how many sub-periods there are.
- Multiply all the growth factors, then subtract one.
- Notice what this does.
- A pound earns its return whether the fund held ten million or a hundred million when it earned it.
- The size of the money, and so the investor's timing, drops out completely.

3 Example 1: the growth fund (set-up)

Harbor Growth Fund, £ millions

Start of year 1	fund holds	**£10m**
Year 1 return	10m grows to 11m	**+10%**
Start of year 2	client adds 90m: fund 101m	**£101m**
Year 2 return	101m falls to 95.95m	**−5%**

IN THE LESSON

- Example one, using the story.
- Here is the set-up, row by row.
- At the start of year one the fund holds ten million pounds.
- During year one it returns plus ten per cent, so it grows to eleven million.
- At the start of year two the client adds ninety million pounds, so the fund now holds a hundred and one million.
- During year two it returns minus five per cent.
- There are two sub-periods, because there was one cash flow in the middle.
- The question is, what is the time-weighted return, and what is the internal rate of return?
- We start with the time-weighted return, which is the easier one.

4 Example 1: the time-weighted return

TWR = growth factors multiplied, minus one

Year 1 growth factor	1 + 0.10	**1.10**
Year 2 growth factor	1 − 0.05	**0.95**
Multiply	1.10 × 0.95	**1.045**
Time-weighted return	1.045 − 1	**+4.5%**

IN THE LESSON

- Step one.
- Turn each sub-period return into a growth factor.
- Year one is one plus zero point one, which is one point one zero.
- Year two is one minus zero point zero five, which is zero point nine five.
- Step two.
- Multiply them.
- One point one zero times zero point nine five is one point zero four five.
- Step three.
- Subtract one.
- The time-weighted return is plus four point five per cent over the two years, which is about two point two per cent a year if you want it annual.
- Notice that the ninety million pounds never entered the sum.
- That is deliberate.
- The manager did not choose when the client added the money, so the manager is not marked on it.

5 Internal rate of return: follow the money

$0 = CF_0 + CF_1 \div (1+IRR) + CF_2 \div (1+IRR)^2 + \ldots$

Cash flow (CF)	Out negative, in positive
Time of the flow (t)	Years from the start
IRR	The rate that makes the sum zero

IN THE LESSON

- The second measure is the internal rate of return, or I R R.
- It is also called the money-weighted return.
- This time we look at the actual cash flows of the investor.
- Money paid in counts as negative, because it leaves your pocket.
- Money received, or the final value, counts as positive.
- The I R R is the single discount rate that makes the present value of all those flows equal to zero.
- In the formula, CF is a cash flow, and t is the time at which it happens, counted in years from the start.
- Because big balances sit invested for longer, they carry more weight, and so the I R R reflects when the money arrived.
- Usually you find it with a calculator.

6 Example 1: the IRR, followed from the client's side

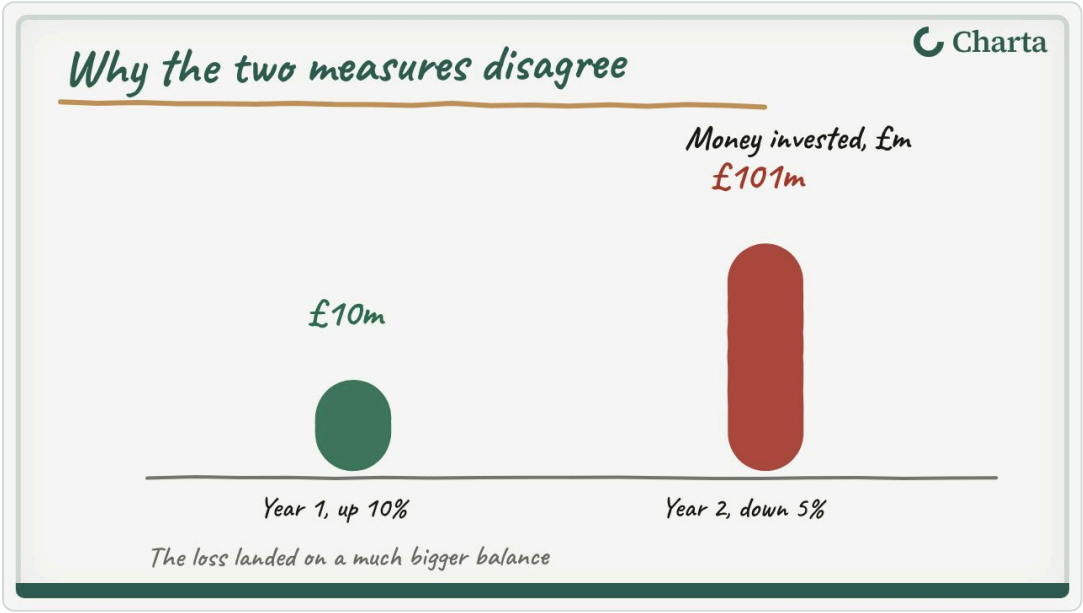
Cash flows from the client's side, £ millions

Time 0: pay in		**−10**
Time 1: pay in more		**−90**
Time 2: fund is worth		**+95.95**
IRR that makes the sum zero	solved by calculator	**−3.7%**

IN THE LESSON

- Now the client's side of example one.
- At time zero the client pays in ten million pounds, so the flow is minus ten.
- One year later, at the start of year two, the client pays in ninety million, minus ninety.
- Two years from the start, the fund is worth ninety-five point nine five million, plus ninety-five point nine five.
- Which discount rate makes these three flows add up to zero?
- The answer is minus three point seven per cent.
- Check it by plugging in.
- Ninety divided by zero point nine six three, and ninety-five point nine five divided by that squared, together bring the sum to about zero.
- The client paid in a hundred and got back ninety-five point nine five, so the I R R is negative.

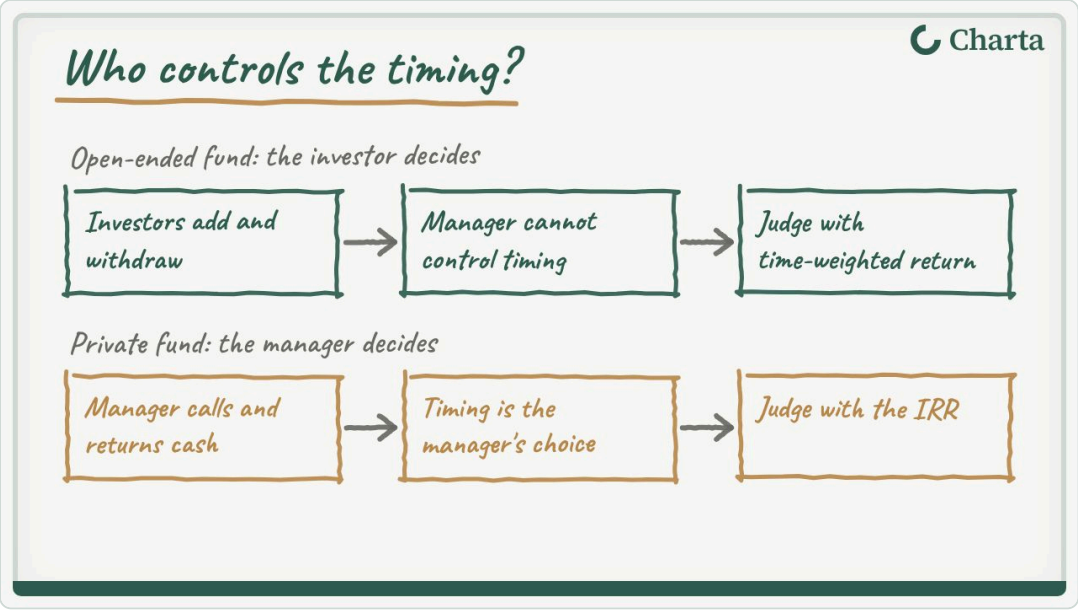
7 Why the two measures disagree



IN THE LESSON

- Look at the picture.
- The bars show how much money was invested in each of the two years.
- In year one, only ten million pounds was in the fund, and it earned ten per cent, a gain of one million.
- In year two, a hundred and one million pounds was in the fund, and it lost five per cent, a loss of more than five million.
- So the good year was small and the bad year was huge.
- The time-weighted return gives each year the same vote, so the good year cancels most of the bad one, and the answer is positive.
- The I R R lets the big balance dominate, so the answer is negative.
- The client put in a hundred million in all and finished with ninety-five point nine five, so the client really did lose money.

8 Who controls the timing?



IN THE LESSON

- So which measure should you use?
- Ask one question: who decides when the money goes in and out?
- Look at the top row.
- In an open-ended fund, the investors choose when to add and when to withdraw.
- The manager cannot control that.
- If you strip out the flows with a time-weighted return, you judge the manager only on the skill they showed.
- Now the bottom row.
- In a private fund, such as private equity, infrastructure or real estate, the general partner, the manager, decides when to call the investors' money and when to pay it back.
- The timing of the cash is part of the manager's job.
- So the money-weighted measure, the I R R, is the right one, because the timing is part of what you judge.

9 Example 2: a private equity fund

Cedar Private Equity: commitment £100m

Capital called	t0: 40, t1: 60	**− £100m**
Cash paid back	t2: 30, t3: 50, t4: 90	**+ £170m**
Money multiple	170 ÷ 100	**1.70×**
IRR of these flows	NPV = 0, by calculator	**21.6%**

IN THE LESSON

- Example two.
- Cedar Private Equity is a private fund.
- An investor commits a hundred million pounds.
- The manager calls forty million at the start, which we call time zero, and sixty million one year later.
- Then the fund sells its holdings and pays back thirty million at the end of year two, fifty million at the end of year three, and ninety million at the end of year four.
- The investor paid in a hundred million and received a hundred and seventy million, so the money multiple is one point seven times.
- That is the simple version.
- The I R R also asks how long the money was out.
- Solve the cash flows, and the I R R is twenty-one point six per cent.
- A time-weighted return cannot do this, because a private fund has no price each time cash moves.

10 Example 3: when the manager controls the timing

Same asset, same £50m gain, different call date

Cash called at year 0	−100 then +150 at year 3	**14.5%**
Cash called at year 2	−100 then +150 at year 3	**50.0%**
Money multiple, both cases	150 ÷ 100	**1.5×**
Gain created, both cases	unchanged	**+ £50m**

IN THE LESSON

- Example three shows why the timing matters.
- Suppose a fund buys an asset worth a hundred million pounds and sells it three years later for a hundred and fifty.
- If the investors hand over the hundred million on day one, the I R R is fourteen point five per cent.
- Now suppose the manager borrows the money to buy the asset and calls the investors' cash only at the end of year two.
- The same asset, the same gain of fifty million, the same multiple of one point five.
- But the investors' money was out for just one year, so the I R R jumps to fifty per cent.
- We ignore the loan interest to keep it simple.
- Nothing about the investment improved.
- Only the timing changed, and it was the manager's choice.

11 Which measure fits which fund?

Match the measure to who controls the cash

Situation	Measure
Open fund, investor times the flows	TWR
Hedge fund open to redemptions	TWR
Private equity, manager calls cash	IRR
Private IRR beside index TWR	Mismatch

IN THE LESSON

- Let us pull part one together into a table.
- For an open-ended fund where investors choose when to enter and leave, judge the manager with the time-weighted return.
- A hedge fund open to redemptions is the same.
- For a private fund where the general partner calls and returns the cash, use the I R R.
- That covers private equity and most real estate and infrastructure funds.
- One last warning, and it is a favourite of questions.
- Never put a private fund's I R R next to a public index's time-weighted return and call it a fair comparison.
- One counts the timing of the cash and the other deliberately ignores it.
- They answer different questions, so the comparison is not like for like.

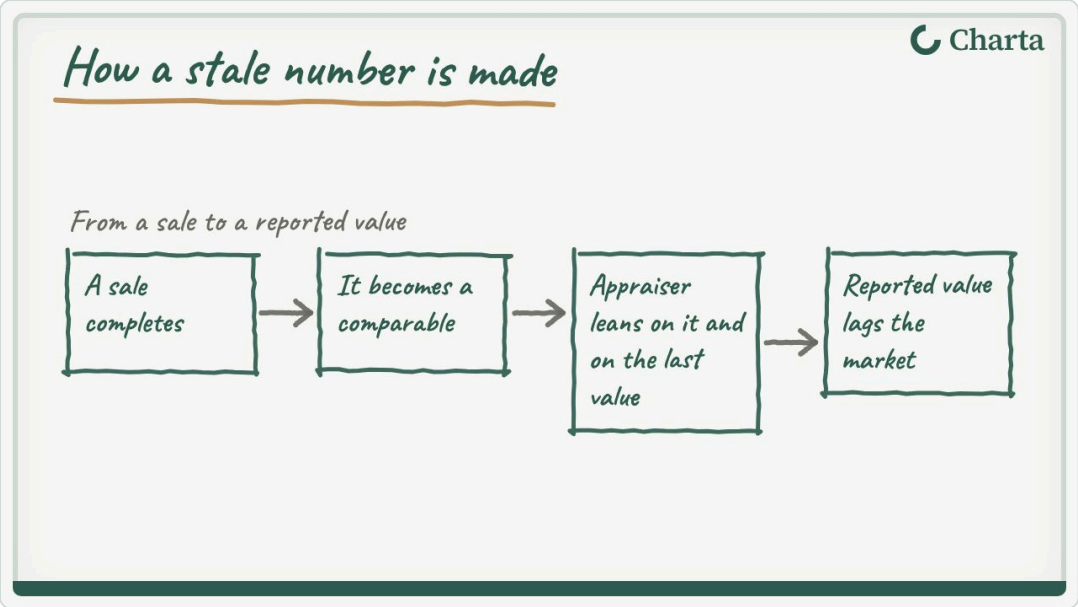
12 A story: the stale house price

- Unlisted assets rarely trade, so they are appraised
- Appraisal leans on old comparable sales
- Values move late, and move less
- This is valuation lag and smoothing

IN THE LESSON

- Imagine you want to know what your house is worth.
- Houses do not trade every day.
- So an estate agent looks at a house that sold six months ago down the road, and adjusts a little.
- Now suppose prices have fallen since.
- Nobody on your street has sold, so the old sale is all the agent has.
- The estimate stays high.
- Only slowly, as new sales trickle in, does it catch up.
- Many alternative assets work this way: unlisted property, private equity and infrastructure do not trade every day, so they are valued by appraisal, an estimate by a valuer.
- The reported returns lag the truth, and they are smoother than the truth.
- That is called valuation lag and smoothing.

13 How a stale number is made



IN THE LESSON

- Here is the chain, step by step.
- First, a sale completes at that day's price.
- By the time the sale is recorded and used by others, it is already months old.
- Second, it becomes a comparable, a reference point for valuing similar assets.
- Third, the appraiser leans on the comparable, and also on their own last valuation.
- They are usually cautious about changing a number sharply, since a big jump needs a lot of justification.
- Fourth, the value reported to investors sits behind the market.
- Think of a sluggish thermometer that is slow to rise and slow to fall.
- The reading is not wrong, but it is always a little late, and it never swings as far as the real temperature.

14 A simple model of smoothing

Reported = (1 − φ) × True + φ × Last reported

Reported return	What investors see
True return	What really happened
Smoothing weight (φ)	Weight on the old number
Last reported return	Last quarter's reported number

IN THE LESSON

- We can describe the smoothing with one formula.
- Each quarter, the reported return is a blend of two things: the true return in that quarter, and the return reported in the previous quarter.
- The letter phi, written as the Greek letter on the board, is the weight put on the old number.
- If phi is zero, there is no smoothing, and you simply report the truth.
- If phi is one half, half of each reported return is a repeat of the last one.
- The bigger phi gets, the more the series is smoothed.
- This is a standard model in the textbooks.
- The point is not the exact weight, but the idea that today's reported number contains yesterday's number.

15 Example 4: a property shock, set-up

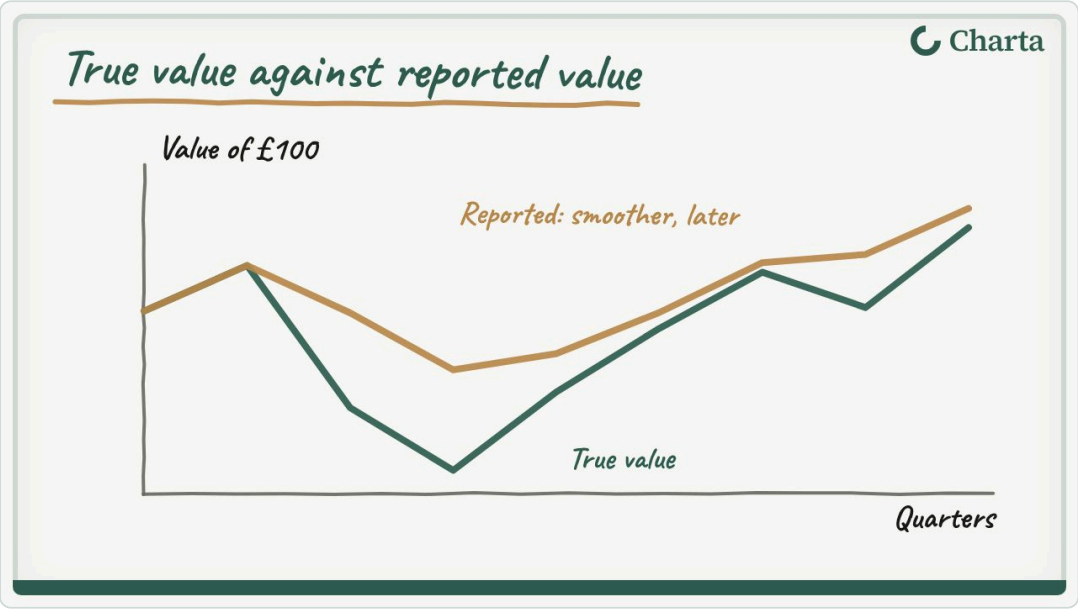
Harbor Property Fund, φ = 0.5, returns by quarter

Quarter 1: true +4%, reported	starts equal	**+4%**
Quarter 2: true −12%, reported	0.5×(−12) + 0.5×4	**−4%**
Quarter 3: true −6%, reported	0.5×(−6) + 0.5×(−4)	**−5%**
Quarter 4: true +8%, reported	0.5×8 + 0.5×(−5)	**+1.5%**

IN THE LESSON

- Example four.
- Harbor Property Fund holds unlisted buildings and reports every quarter.
- Suppose phi is one half.
- The true returns in the first four quarters are plus four, minus twelve, minus six and plus eight per cent.
- Say the first reported return equals the first true return, plus four.
- Now apply the formula.
- In quarter two, half of minus twelve is minus six, and half of the old plus four is plus two, so the report is minus four.
- In quarter three, half of minus six is minus three, and half of minus four is minus two, so minus five.
- In quarter four, half of plus eight is plus four, and half of minus five is minus two and a half, so plus one and a half.
- The market fell twelve.
- The fund said minus four.

16 True value against reported value



IN THE LESSON

- Here is the same fund drawn as a picture.
- The bottom axis is time in quarters.
- The side axis is the value of a hundred pounds invested in the fund.
- Follow the dark line first, the true value.
- It climbs in the first quarter, then plunges over the next two, down to about eighty-six pounds, and then recovers.
- Now the lighter line, the reported value.
- It dips only to about ninety-five pounds, and it keeps showing a higher value through the fall.
- The reported line is smoother, and it is late.
- Both lines end in about the same place, over the whole period, so the average return is almost the same.
- What the smoothing hides is the size and the timing of the swings, and that is the risk.

17 Example 4: what smoothing does to the statistics

Eight quarters, φ = 0.5, annual = quarterly × 2

Volatility a year, true	7.26% × 2	**14.5%**
Volatility a year, reported	3.73% × 2	**7.5%**
Correlation with the market	reported against true	**0.87**
Sharpe a year, true, reported	cash 0.5% a quarter	**0.17, 0.35**

IN THE LESSON

- Now run eight quarters through the same model and compare.
- The true quarterly returns are plus four, minus twelve, minus six, plus eight, plus six, plus five, minus three and plus seven.
- Their average is about one point one per cent a quarter, and the reported series has almost the same average.
- But look at the risk.
- The true standard deviation is seven point three per cent a quarter, or fourteen and a half per cent a year.
- The reported one is only three point seven per cent a quarter, seven and a half a year, half the true level.
- If the market moved exactly with the true returns, the reported correlation with the market is only zero point eight seven, not one.
- The Sharpe ratio, return per unit of risk, is about doubled, from zero point one seven to zero point three five, with a cash return of half a per cent a quarter.

18 Unsmoothing: undoing the lag

True = (Reported − φ × Last reported) ÷ (1 − φ)

True return	The unsmoothed estimate
Reported return	Reported this quarter
Last reported return	Reported last quarter
Smoothing weight (φ)	Estimated smoothing weight

IN THE LESSON

- Analysts can undo the smoothing.
- If reported equals one minus phi times true plus phi times last reported, then a little algebra turns the formula round.
- True equals reported, minus phi times last reported, all divided by one minus phi.
- This is called unsmoothing, or de-smoothing.
- You take each reported return, subtract the part that is just a repeat of the last one, and scale the rest back up.
- In practice, nobody knows phi for certain, so the analyst estimates it from the data, and the unsmoothed series is only an estimate.
- But it generally gives a much higher volatility, and a higher correlation with public markets, which is the honest picture.

19 Example 5: unsmoothing the property fund

Unsmoothed = (reported − 0.5 × last reported) ÷ 0.5

Quarter 2	(−4 − 0.5×4) ÷ 0.5	**−12%**
Quarter 3	(−5 − 0.5×(−4)) ÷ 0.5	**−6%**
Quarter 4	(1.5 − 0.5×(−5)) ÷ 0.5	**+8%**

IN THE LESSON

- Example five.
- Take the reported returns from example four and unsmooth them with phi equal to one half.
- Quarter two.
- Reported minus four, minus a half of last quarter's plus four, is minus six.
- Divide by one minus a half, which is a half, and we get minus twelve.
- That is exactly the true return.
- Quarter three.
- Reported minus five, minus a half of minus four, is minus three.
- Divide by a half and we get minus six.
- Quarter four.
- Reported plus one and a half, minus a half of minus five, is plus four.
- Divided by a half, plus eight.
- All three true returns come back.
- Real data would not be this tidy, since phi must be estimated, but you can see that unsmoothing lifts the apparent risk.

20 Example 6: why the optimiser over-buys

Extra return 6%, risk aversion 10: w = 6% ÷ (10 × variance)

Weight using reported vol 10%	0.06 ÷ (10 × 0.01)	**60%**
Weight using unsmoothed vol 20%	0.06 ÷ (10 × 0.04)	**15%**
Over-weighted by	60 ÷ 15	**4 times**

IN THE LESSON

- Example six is why this matters to a real portfolio.
- A mean-variance optimiser chooses how much to put in a risky asset using its return and its variance, which is the standard deviation squared.
- Here is a simple rule.
- Take the extra return over cash, which is six per cent here, and divide it by risk aversion times variance.
- Take a risk aversion of ten.
- If you feed in the reported volatility of ten per cent, the variance is zero point zero one, and the rule says put in sixty per cent.
- If you feed in the unsmoothed volatility of twenty per cent, the variance is zero point zero four, and the rule says fifteen per cent.
- Halving the volatility makes the variance four times smaller, so the weight is four times too large.
- Raw reported statistics tend to overweight the asset.

21 Mistakes to avoid, so far

- 1. Private IRR beside index TWR: not like for like
- 2. Ask who controls the timing of the cash
- 3. Low volatility and correlation: partly lag
- 4. Unsmoothing raises volatility and correlation
- 5. Smoothing is mostly a feature of appraisal

IN THE LESSON

- Time to stop and check the traps so far.
- One.
- Comparing a private fund's I R R with a public index's time-weighted return.
- They are not like for like.
- Two.
- Using the time-weighted return for a private fund where the manager controls the cash, or the I R R for a fund whose investors control the cash.
- Ask who decides the timing.
- Three.
- Believing the low volatility and low correlation of appraised assets are real diversification.
- Part of it is just lag.
- Four.
- Forgetting that unsmoothing pushes volatility up, not down.
- Five.
- Thinking smoothing means someone cheated.
- Mostly it is the honest result of valuing things by appraisal, though a manager could abuse it.

22 A story: selling storm insurance

- Many small gains, rare large losses
- Negative skew: long tail on the loss side
- Fat tails: extremes likelier than a bell curve
- Standard deviation counts up and down alike

IN THE LESSON

- Imagine a firm that sells storm insurance on a coast.
- Most months it collects premiums and pays nothing.
- Month after month, small steady gains.
- Then, once in a long while, a great storm arrives and wipes out years of premiums at a stroke.
- Many alternative strategies look like this: selling options, lending in credit markets, and short-volatility trades.
- The picture is many small gains and rare large losses.
- Statisticians call this negative skew, a long tail on the loss side, and fat tails, meaning extreme outcomes are more likely than the bell curve suggests.
- The usual measure of risk, standard deviation, treats a gain and a loss of the same size as the same thing.
- That is the problem.

23 Steady climb, then a cliff



IN THE LESSON

- Here is a short-volatility strategy of Harbor Capital, drawn as a picture.
- The bottom axis is months, thirty-five of them.
- The side axis is the value of a hundred pounds invested.
- Follow the line from the left.
- For thirty-four months it climbs steadily, gaining either six tenths or one per cent a month, up to about a hundred and thirty-one pounds.
- Nothing about it looks dangerous.
- Then comes the last month.
- The line drops by nineteen per cent in a single step, down to about a hundred and six pounds, wiping out most of three years of gains.
- A strategy like this could show a high hit rate, and a stunning risk-adjusted ratio, for years, before the cliff arrives.

24 Example 7: the usual numbers look fine

35 months: 34 gains of 0.6% or 1.0%, then −19%

Months with a gain	34 of 35	**97%**
Mean monthly return		**0.23%**
Standard deviation	all 35 months	**3.35%**
Worst month, in std devs	(−19 − 0.23) ÷ 3.35	**5.7 below**

IN THE LESSON

- Example seven.
- Take the same thirty-five months.
- In thirty-four of them the strategy gained, so the hit rate is ninety-seven per cent.
- The mean monthly return is positive, at about zero point two three per cent.
- The standard deviation of all thirty-five months is three point four per cent.
- That sounds modest.
- But look at the last row.
- The worst month, minus nineteen, is more than five and a half standard deviations below the mean.
- A normal curve says that should almost never happen.
- And before the bad month, the standard deviation of the first thirty-four was only about zero point two per cent, so the reported Sharpe ratio was enormous.
- The usual numbers were looking away from the one thing that mattered.

25 Downside deviation and the Sortino ratio

Downside deviation = √[Σ min(0, R − MAR)² ÷ (n − 1)]
Sortino = (Mean − MAR) ÷ Downside deviation

Target (MAR)	Return you must at least earn
Downside deviation	Square root of average shortfall²
Mean return	Average return
Sortino ratio	Return per unit of loss risk

IN THE LESSON

- Now for a measure that looks only at the loss side.
- First pick a target, the minimum acceptable return, which we write as M A R.
- It might be zero, or the cash rate.
- Downside deviation counts only the returns below the target.
- For each return, take the shortfall, how far below the target it fell, or zero if it did not fall short.
- Square the shortfalls, add them up, divide by n minus one, where n is all the observations, and take the square root.
- Gains above the target do not count.
- The Sortino ratio then replaces the standard deviation in the Sharpe ratio with this downside deviation: the mean return minus the target, divided by the downside deviation.

26 Example 8: Sharpe and Sortino disagree

Two funds, 8 months, mean 4% each, target 0%

Standard deviation: Up, Down	spread of all months	**6.4 and 5.7**
Sharpe ratio: Up, Down	mean ÷ std dev	**0.62 and 0.71**
Downside deviation: Up, Down	months below 0 only	**1.51 and 3.78**
Sortino ratio: Up, Down	mean ÷ downside dev	**2.65 and 1.06**

IN THE LESSON

- Example eight shows the difference.
- Take two funds over eight months, with a target of zero and cash at zero.
- Fund Up earns plus ten and minus two, alternating, four times each.
- Fund Down earns plus six in seven months, then minus ten once.
- Both average four per cent a month.
- Fund Up has a standard deviation of six point four, and Fund Down only five point seven.
- So the Sharpe ratios are zero point six two for Up and zero point seven one for Down, and Sharpe prefers Down.
- Now look below the target.
- Up has downside deviation of one point five, and Down three point eight.
- The Sortino ratios are two point six five for Up and one point zero six for Down.
- Sortino prefers Up, as most investors would, since Up's volatility is all on the good side.

27 Example 9: maximum drawdown

Harbor Credit Fund, value at month-ends, £ millions

Highest value (the peak)	after 108, 115	**£115m**
Lowest value after the peak	104, then 92	**£92m**
Maximum drawdown	(92 − 115) ÷ 115	**−20%**
Gain needed to recover	115 ÷ 92 − 1	**+25%**

IN THE LESSON

- Another loss-side measure is maximum drawdown.
- Walk along the value of the fund.
- Whenever it sets a new high, that is a peak.
- A drawdown is the fall from the latest peak to the lowest point afterwards, and the maximum drawdown is the biggest such fall in the period.
- Example nine.
- Harbor Credit Fund starts at a hundred million pounds and at seven month-ends is worth one hundred, one hundred and eight, one hundred and fifteen, one hundred and four, ninety-two, ninety-eight and one hundred and twelve.
- The peak is a hundred and fifteen and the lowest point after it is ninety-two.
- The fall is twenty-three, divided by one hundred and fifteen, which is minus twenty per cent.
- Note that to get back to the peak, the fund needs a gain of twenty-five per cent, not twenty.

28 Value at risk: a cut-off, not a worst case

Normal VaR = −(Mean − z × s)

Mean return	Average return per period
Cut-off (z)	1.645 at 95%, 2.33 at 99%
Standard deviation (s)	Spread of returns
Value at risk	Loss not expected to be exceeded

IN THE LESSON

- A third measure is value at risk, or V a R.
- It answers this question.
- At a chosen confidence level, such as ninety-five per cent, what is the loss that we expect not to exceed in a month?
- The simplest version assumes returns are normal.
- Take the mean return and subtract z standard deviations, where z is the cut-off from the normal curve: one point six four five at ninety-five per cent, and two point three three at ninety-nine.
- The loss is the negative of that.
- A historical version just sorts past returns instead.
- Two warnings: the normal version understates fat tails, and neither says how bad the loss is beyond the cut-off.

29 Example 10: value at risk meets the cliff

35 months: mean 0.23%, std dev 3.35%, z = 1.645

z × std dev	1.645 × 3.35	**5.51**
Mean − z × std dev	0.23 − 5.51	**−5.28%**
95% monthly VaR, as a loss		**5.3%**
Actual worst month beyond VaR	19 − 5.3	**13.7 points**

IN THE LESSON

- Example ten.
- Use the same thirty-five months.
- The mean is zero point two three per cent a month and the standard deviation is three point three five.
- Step one.
- Take the cut-off for ninety-five per cent confidence, one point six four five, times three point three five, which is five point five one.
- Step two.
- Subtract it from the mean, which gives minus five point two eight.
- Step three.
- State it as a loss: the monthly ninety-five per cent value at risk is about five point three per cent.
- In plain words, the normal model says that in nineteen months out of twenty, the loss will be smaller than five point three.
- Yet the actual worst month was nineteen per cent, nearly fourteen points beyond.
- The measure did not fail by arithmetic.
- It failed because the tail was fatter than the model assumed.

30 What each measure tells you, and what it hides

Each measure answers one narrow question

Measure	It answers
Standard deviation	Up and down
Downside deviation, Sortino	Below target
Maximum drawdown	Peak to trough
Value at risk	A cut-off

IN THE LESSON

- Let us line the measures up.
- Standard deviation says how much returns wobble, upwards and downwards alike.
- It misleads when losses are rare and large.
- Downside deviation counts only the shortfalls below a target.
- The Sortino ratio puts the return from that.
- Maximum drawdown shows the worst fall from a peak, the pain a real investor felt.
- And value at risk gives a loss cut-off at a confidence level, but says nothing beyond the cut-off.
- No single measure is enough.
- The surest guard is to understand the strategy's payoff.
- If it sells options or insurance, expect negative skew and fat tails, whatever the history of calm months says.
- Use several measures, and read the strategy as well as the numbers.

31 Recap: three questions about alternative returns

Alternative returns: choose the measure, correct the lag, look at the loss side

Which return?	TWR open fund, IRR private fund
Valuation lag	Unsmooth: more risk, less Sharpe
Look below target	Downside deviation, Sortino
Maximum drawdown	Worst peak to trough
Value at risk	A cut-off, not the tail

IN THE LESSON

- Let us pull the whole lesson together.
- Which return measure?
- Ask who controls the timing of the cash.
- The time-weighted return suits an open fund, where the investor times the flows.
- The I R R suits a private fund, where the manager does.
- Next, how reliable is the reported risk?
- For appraised assets, valuation lag smooths the series, so volatility and correlation are understated and the Sharpe ratio flattered.
- Unsmooth to correct it, and expect an optimiser fed raw numbers to overweight the asset.
- Last, for returns that are not normal, look below the target with downside deviation and the Sortino ratio, check the maximum drawdown, use value at risk with care, and understand the payoff.

32 Try three quick questions

1. General partner decides when capital is called. Which return measure?
   Answer: The IRR: the manager controls the timing.
2. Reported volatility is 7%. What does unsmoothing do to it?
   Answer: It raises it; the Sharpe ratio falls.
3. Mean 4%, target 0%, downside deviation 2%. Sortino ratio?
   Answer: (4 − 0) ÷ 2 = 2.0.

IN THE LESSON

- Now try three quick questions yourself.
- Pause after each one.
- Question one.
- The general partner decides when capital is called.
- Which return measure suits it?
- The answer is the I R R, because the manager controls the timing.
- Question two.
- Reported volatility is seven per cent.
- What does unsmoothing do to it?
- The answer is that it raises the volatility, and the Sharpe ratio falls.
- Question three.
- Mean four per cent, target zero per cent, downside deviation two per cent.
- What is the Sortino ratio?
- The answer is four minus zero, divided by two, which is two.
- Well done.