CHARTA LEARNING · LESSON NOTES

Judging alternatives and managers

Performance Measurement

1 A story: the ruler that stretches

- Clear, priced, investable, set in advance
- Plain shares pass these tests easily
- Alternatives break them, each in its own way
- No single benchmark will be enough

IN THE LESSON

- Imagine measuring a table with a ruler that stretches a little every time you pull it.
- You would get an answer, but you could not trust it.
- Benchmarks can behave like that.
- You met the tests of a good benchmark earlier in this topic: clear holdings, a price you can see, something an investor could actually buy, and a rule fixed in advance.
- A plain share portfolio passes these tests easily.
- Mara Quinn now has to benchmark the Brindle trustees' alternatives, and each one breaks a test in its own way.
- Hedge funds have no steady weights.
- Private assets have no daily prices.
- Commodities have many rival indices.
- So let us take them in turn.

2 Four ways to benchmark a hedge fund

Corvid hedge fund: what each benchmark gets wrong

Benchmark choice	Flaw
Broad index: fund is long, short, levered	poor fit
Cash plus 4%: easy to state	ignores market
Peer group median	biased
Factor model: right risks	unstable

IN THE LESSON

- Start with the hedge fund.
- The trustees own one called Corvid, which goes long some shares, short others, and uses borrowing.
- Its holdings and its market exposure shift all the time, so there is no stable set of weights to compare with an index.
- Mara Quinn runs through the candidates.
- A broad equity index fits poorly, because the fund is long, short and levered.
- An absolute target, such as cash plus four per cent, is easy to state, but it ignores what markets did.
- A peer group median can be biased, as we will see.
- And a factor model, which measures the fund against the risks it really takes, is better, but it needs the right factors and they can change.

3 Example 1: a target that flatters

Cash 1%, target cash + 4%, Corvid +8%, market +16%, beta 0.4

Absolute target	1% + 4%	**5.0%**
Corvid beats target by	8.0 − 5.0	**3.0 pts**
Return from market exposure	1 + 0.4 × (16 − 1)	**7.0%**
Corvid beats that by	8.0 − 7.0	**1.0 pt**
Target flattered it by	3.0 − 1.0	**2.0 pts**

IN THE LESSON

- Here is why an absolute target can mislead.
- Cash pays one per cent, so cash plus four gives a target of five per cent.
- Corvid returns eight per cent in a year when the market rises sixteen per cent.
- Against the target, Corvid wins by three points, and the trustees are pleased.
- But Mara Quinn estimates that Corvid carries a market beta of point four.
- Its market exposure alone should have earned one plus point four times the market's return above cash, sixteen minus one, which is seven per cent.
- Against that fairer figure, Corvid wins by only one point.
- So the target flattered the fund by two points.
- A rising market did most of the work, not the manager.

4 Peer groups can mislead

- Not investable: nobody can buy the median fund
- Survivorship: funds that close drop out
- Backfill: new members bring their good past
- Strategies inside one group differ widely

IN THE LESSON

- Many trustees compare a hedge fund with the median fund in a database of its peers.
- It sounds sensible, but there are four snags.
- First, you cannot invest in a peer index.
- Nobody can buy the median fund, so it fails the test of being investable.
- Second, survivorship bias.
- Funds that do badly close and drop out of the database, so the survivors look better than the whole group really was.
- Third, backfill bias.
- A fund that joins the database often brings its good early history with it, and its poor start is never recorded.
- Fourth, the strategies inside one group can be very different, so comparing them is not fair.
- Peer groups are useful context, but they are a weak yardstick on their own.

5 Appraisals move slowly and late

Chart: Value vs Time (quarters)

- Market value swings
- Appraised value is smooth

IN THE LESSON

- Now private assets, such as private equity and property.
- They are valued by appraisal, not by daily trades.
- Look at the picture.
- The horizontal axis is time, in quarters.
- The vertical axis is value.
- The red line is what a market price would do, swinging up and down.
- Now the green line, the appraised value.
- Each valuer starts from the last valuation, so the green line moves much less and turns later than the red one.
- Smoothed values make the asset look steadier than it is.
- Reported volatility is too low, and so is the correlation with shares, which makes private assets seem better diversifiers than they really are.
- This is the same smoothing we met in the earlier lesson on liquidity.

6 The public market equivalent

PME = (payouts grown at index + remaining value) ÷ calls grown at index

Payouts grown at index	Payout × end index ÷ its index
Remaining value	Fund's reported value at the end
Calls grown at index	Call × end index ÷ its index
Public market equivalent	Above 1: fund beat the index

IN THE LESSON

- Private equity brings a second problem.
- The manager decides when to call money and when to give it back, so the investor does not control the timing.
- Returns are therefore measured on a money weighted basis, and compared with a public market equivalent.
- The idea is simple.
- Suppose every payment into the fund had instead been put into a public index on the same date.
- Grow each call and each payout at the index, up to the end date.
- Then divide the fund's grown payouts plus what it still holds by the grown calls.
- A result above one means the private fund beat an identical public investment.
- Below one means it lagged.

7 Example 2: Halden private equity

Call 100; paid back 50 in year 3; value 100 in year 5; index 100, 120, 150

Multiple of capital	(50 + 100) ÷ 100	**1.50×**
Call grown at index	100 × 150 ÷ 100	**150**
Payout grown at index	50 × 150 ÷ 120	**62.5**
Payout plus remaining value	62.5 + 100	**162.5**
Public market equivalent	162.5 ÷ 150	**1.08**

IN THE LESSON

- The Brindle Pension Fund owns a private equity fund called Halden.
- It called one hundred at the start, paid back fifty in year three, and still has a reported value of one hundred in year five.
- A public index stood at one hundred, then one hundred and twenty, then one hundred and fifty.
- The simple multiple is fifty plus one hundred, over one hundred, which is one point five.
- That looks strong.
- But the index itself rose fifty per cent.
- Grow the call to one hundred and fifty.
- Grow the payout to sixty two point five.
- Add the remaining one hundred to get one hundred and sixty two point five.
- Divide by one hundred and fifty.
- The public market equivalent is only one point zero eight, a modest margin.

8 A blend, not one benchmark

Hedge fund

Custom benchmark → Factor exposures → Peer comparison

Private assets

Money-weighted return → Public market equivalent → Unsmoothed risk

IN THE LESSON

- So what does Mara Quinn recommend?
- Not one benchmark, but a blend, and the blend differs by holding.
- For the hedge fund, use three comparisons together.
- A custom benchmark built from the fund's stated strategy.
- Its factor exposures, to see which risks it really takes.
- And a peer comparison, read with the biases in mind.
- For the private assets, measure money weighted returns, compare them with a public market equivalent, and unsmooth the risk figures so volatility is not understated.
- The answer is to explain why no single benchmark suffices, and to propose the blend.

9 A story: two hikers

- Same average return can hide very different journeys
- Drawdown: the worst fall from a peak to a later trough
- Sortino: return above a target, per unit of shortfall risk
- Appraisal and capture ratios judge the manager's edge

IN THE LESSON

- Two hikers both finish a long walk at the same height.
- One climbed steadily all day.
- The other went up a mountain, down into a deep valley, and climbed out again late in the afternoon.
- Their average progress is the same, but you would not describe the days in the same way.
- Investments are like that.
- Two managers can have the same average return and the same volatility, yet one fell hard and took years to recover, and the other held up well in bad markets.
- Mara Quinn needs measures of the path, not just the end point.
- Four are on the board: drawdown and its duration, the Sortino ratio, the appraisal ratio and the capture ratios.

10 Peak, trough, recovery

Chart: Portfolio value vs Month

- peak 118
- trough 92
- regained at 121

IN THE LESSON

- Here is the value of Fenner's portfolio over eight months.
- The horizontal axis is the month, and the vertical axis is portfolio value, starting at one hundred.
- The line rises to a peak of one hundred and eighteen in month two.
- It then falls to a trough of ninety two in month five.
- The dashed line marks the old peak.
- Notice that the value only climbs back above it in month eight, at one hundred and twenty one.
- The depth of the fall from the peak is the maximum drawdown.
- The time from the peak, down through the trough, and back to the old peak is the drawdown duration.
- An investor who needed cash in month five lived through the trough, not through the average.

11 Example 3: Fenner's drawdown

Fenner's value, months 0 to 8: 100, 106, 118, 112, 101, 92, 98, 109, 121

Peak, month 2	highest before the fall	**118**
Trough, month 5	lowest after it	**92**
Maximum drawdown	(92 − 118) ÷ 118	**−22.0%**
Drawdown duration	month 2 to month 8	**6 months**
Gain needed from trough	118 ÷ 92 − 1	**28.3%**

IN THE LESSON

- Now the numbers.
- Fenner's values are one hundred, one hundred and six, one hundred and eighteen, one hundred and twelve, one hundred and one, ninety two, ninety eight, one hundred and nine and one hundred and twenty one.
- The highest point before the big fall is one hundred and eighteen.
- The lowest point after it is ninety two.
- The fall is ninety two minus one hundred and eighteen, over one hundred and eighteen, which is minus twenty two per cent.
- The old peak is regained in month eight, so the duration is six months, from month two.
- And notice the climb back.
- From ninety two to one hundred and eighteen needs a gain of twenty eight per cent, which is more than the fall itself.

12 The Sortino ratio

Sortino = (Rp − MAR) ÷ DD
DD = √(Σ min(Ri − MAR, 0)² ÷ n)

Portfolio return	Average return of the portfolio
Minimum acceptable return	The least the client can accept
Downside deviation	Counts only returns below the MAR
Sortino ratio	Reward per unit of shortfall

IN THE LESSON

- The Sortino ratio improves on the Sharpe ratio.
- Sharpe divides by standard deviation, which treats a big gain as risk, exactly as it treats a big loss.
- But no investor complains about surprises on the upside.
- The Sortino ratio takes the portfolio's return, subtracts the minimum acceptable return, the least the client can live with, and divides by the downside deviation.
- To find the downside deviation, take each return below the minimum acceptable return, find the shortfall, square it, add the squares, divide by the number of all the observations, and take the square root.
- Returns above the target count as zero.
- A higher Sortino ratio means more reward per unit of risk that actually hurts.

13 Example 4: Fenner's Sortino ratio

Fenner's returns: 12%, −4%, 9%, 15%, −2%. Minimum acceptable return 3%

Average return	(12 − 4 + 9 + 15 − 2) ÷ 5	**6.0%**
Shortfalls below 3%	−4 → −7; −2 → −5	**−7, −5**
Squared, summed, ÷ 5	(49 + 25) ÷ 5	**14.8**
Downside deviation	√14.8	**3.85%**
Sortino ratio	(6.0 − 3.0) ÷ 3.85	**0.78**

IN THE LESSON

- Here are Fenner's returns over five years: twelve, minus four, nine, fifteen and minus two per cent.
- The Brindle trustees set a minimum acceptable return of three per cent.
- The average return is six per cent.
- Only two years fall short of the target.
- Minus four is seven below it, and minus two is five below it.
- Square the shortfalls to get forty nine and twenty five.
- They add to seventy four.
- Divide by all five years to get fourteen point eight.
- The square root, the downside deviation, is three point eight five per cent.
- So the Sortino ratio is six minus three, over three point eight five, which is point seven eight.
- The big gains in twelve and fifteen did not count against the manager.

14 Appraisal ratio and capture ratios

Appraisal ratio = α ÷ σ(residual)
Capture ratio = upside capture ÷ downside capture

Alpha	Return beyond the benchmark
Residual risk	Unexplained, non-systematic risk
Appraisal ratio	Alpha per unit of residual risk
Upside capture	Fund ÷ index in rising periods
Downside capture	Fund ÷ index in falling periods

IN THE LESSON

- Two more measures judge the manager's edge.
- The appraisal ratio divides alpha, the return the benchmark or factor model cannot explain, by the residual risk, the standard deviation of the returns left over after the model.
- It rewards alpha earned with little unexplained risk.
- Then the capture ratios.
- Upside capture is the manager's average return in the periods when the benchmark rose, divided by the benchmark's average return in those periods.
- Downside capture does the same for the periods when the benchmark fell.
- The capture ratio is upside capture divided by downside capture.
- Above one, the manager gets more of the gains than of the losses, which is a favourable lopsidedness.

15 Example 5: Fenner's edge

Alpha 2.1%, residual risk 6.0%. Index up quarters 4.0%, down −3.6%

Appraisal ratio	2.1% ÷ 6.0%	**0.35**
Upside capture	3.8 ÷ 4.0	**95%**
Downside capture	−2.7 ÷ −3.6	**75%**
Capture ratio	95% ÷ 75%	**1.27**
Verdict	keeps gains, cushions falls	**favourable**

IN THE LESSON

- Back to Fenner.
- Its alpha against the benchmark is two point one per cent a year, and the residual risk is six per cent.
- The appraisal ratio is two point one over six, which is point three five.
- Now capture.
- In the quarters when the benchmark rose, it averaged four point zero per cent and Fenner averaged three point eight, so upside capture is ninety five per cent.
- In the quarters when the benchmark fell, it averaged minus three point six and Fenner averaged minus two point seven, so downside capture is seventy five per cent.
- The capture ratio is ninety five divided by seventy five, which is one point two seven.
- Fenner keeps most of the gains and loses less in the falls.

16 Capture, drawn as bars

Share of index move (%)

Index up	**100%**
Fenner up	**95%**
Index down	**100%**
Fenner down	**75%**

Capture ratio: 95 ÷ 75 = 1.27

IN THE LESSON

- Here are the capture ratios as bars.
- The vertical axis shows each move as a share of the index's move, in per cent.
- The first bar is the index in its rising periods, set at one hundred, the yardstick.
- The second bar is Fenner in those same periods, at ninety five.
- It captured almost all of the rise.
- Now the falling periods.
- The third bar is the index again at one hundred.
- The fourth bar is Fenner, at seventy five, so it took only three quarters of the fall.
- Compare the two Fenner bars.
- Ninety five on the way up and seventy five on the way down is the favourable lopsidedness that gives a capture ratio of one point two seven.
- Plenty of managers show the reverse.

17 Mistakes to avoid so far

- Drawdown duration runs until the old peak is regained
- Sortino: divide by downside deviation, over all observations
- Appraisal ratio: divide by residual risk, not total risk
- Capture ratio: up divided by down, never a difference

IN THE LESSON

- Before we move on, here are four slips that cost marks.
- First, stopping the drawdown duration clock at the trough.
- The clock keeps running until the old peak is regained, so count peak to recovery, not peak to trough.
- Second, dividing the Sortino numerator by the ordinary standard deviation.
- The Sortino ratio divides by the downside deviation, which counts only returns below the target, and which divides by every observation, not just the bad ones.
- Third, using total risk in the appraisal ratio.
- It uses the residual risk, after the benchmark's effect is removed.
- Fourth, treating the capture ratio as a difference.
- It is upside divided by downside, a ratio.

18 A story: the speedometer

- Exact arithmetic can sit on flawed inputs
- Every ratio comes from one sample of past returns
- Ask: what does this measure assume?
- Then ask: does this portfolio break the assumption?

IN THE LESSON

- A speedometer can be perfectly accurate in its sums and still show the wrong speed, if the wheels are the wrong size.
- Performance ratios are the same.
- The arithmetic can be flawless, and the answer can still mislead, because of what went into it.
- The returns may be smoothed, the benchmark may be wrong, or the record may be too short to mean much.
- Mara Quinn therefore asks two questions of every measure.
- What does it assume?
- And does this particular portfolio break that assumption?
- A good written answer names the measure, names the assumption, and explains why it fails in this case.
- Let us go through the weak points one by one.

19 What each measure leans on

Each measure rests on something that can fail

Measure	Leans on
Sharpe ratio, M^2: total volatility	normality
Treynor, alpha, appraisal ratio	the benchmark
Sortino ratio	the chosen MAR
Capture ratios, drawdown	the window

IN THE LESSON

- This table sets each family of measures against the thing it leans on.
- The Sharpe ratio and M squared both divide by total volatility and work best when returns are roughly normal.
- Selling options, using leverage or holding smoothed assets can make them look better than the risk really is.
- The Treynor ratio, Jensen's alpha and the appraisal ratio all depend on the benchmark or factor model.
- A poor market proxy gives the wrong beta, and so the wrong alpha.
- The Sortino ratio depends on the minimum acceptable return you choose, and on having enough observations below it.
- Capture ratios and drawdowns depend on how you cut the periods and on the window you pick.

20 Example 6: a flattering Sharpe ratio

Larch property fund: return 9%, cash 3%, reported volatility 6%, smoothing weight 0.5

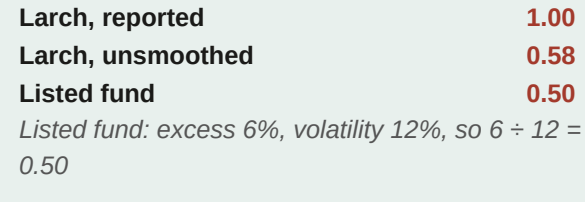
Excess return	9% − 3%	**6%**
Reported Sharpe ratio	6 ÷ 6	**1.00**
Smoothing factor	√((1 + 0.5) ÷ (1 − 0.5))	**1.73**
Unsmoothed volatility	6 × 1.73	**10.4%**
Unsmoothed Sharpe ratio	6 ÷ 10.4	**0.58**

IN THE LESSON

- Take the Brindle property fund, Larch.
- It returns nine per cent while cash pays three, and its reported volatility is six per cent, based on appraisals.
- The excess return is six.
- The reported Sharpe ratio is six over six, which is one point zero zero.
- Now suppose the appraisals have a smoothing weight of point five.
- Recall the correction from the earlier lesson: multiply by the square root of one plus the weight, over one minus the weight.
- That factor is one point seven three.
- The unsmoothed volatility is about ten point four per cent.
- The Sharpe ratio drops to six over ten point four, which is point five eight.
- Same fund, same return.
- Only the input changed, and the ratio fell sharply.

21 Reported against realistic

Sharpe ratio



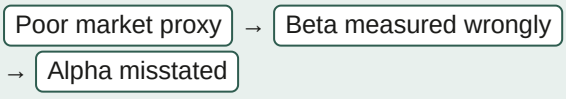
Listed fund: excess 6%, volatility 12%, so 6 ÷ 12 = 0.50

IN THE LESSON

- Here is the same story as a picture.
- The vertical axis is the Sharpe ratio.
- The first bar is Larch on its reported figures, at one point zero zero.
- It looks far ahead of everything.
- The second bar is Larch after unsmoothing, at point five eight.
- Nothing about the fund changed, only the measure of its risk.
- The third bar is a listed equity fund, priced every day, which earns six points over cash with volatility of twelve per cent, so its Sharpe ratio is point five zero.
- On reported numbers Larch seems to beat the listed fund by a mile.
- On comparable numbers the gap is small.
- The lesson for trustees is to compare like with like before drawing any conclusion.

22 A wrong benchmark, a wrong alpha

How one weak input spreads



IN THE LESSON

- The next weak point is the benchmark itself.
- Treynor, Jensen's alpha and the appraisal ratio all rest on beta, and beta is only as good as the market proxy it is measured against.
- Follow the boxes from left to right.
- If the market proxy does not really represent the fund's market, the beta is measured wrongly.
- A wrong beta means the return that the fund should have earned for its risk is wrong.
- And that means alpha is wrong, so skill can be overstated or understated.
- The fault is hidden, because every figure looks exact.
- Before trusting any alpha, the adviser asks whether the benchmark or factor model fits the strategy.

23 Example 7: alpha from the wrong beta

Manager +10%, cash 2%, market 8%. Proxy beta 0.9, true beta 1.2

Required, beta 0.9	2 + 0.9 × (8 − 2)	**7.4%**
Alpha, beta 0.9	10 − 7.4	**2.6%**
Required, beta 1.2	2 + 1.2 × (8 − 2)	**9.2%**
Alpha, beta 1.2	10 − 9.2	**0.8%**
Alpha overstated by	2.6 − 0.8	**1.8 pts**

IN THE LESSON

- Let us put numbers on it.
- A manager returns ten per cent, cash pays two, and the market returns eight.
- Using a poor proxy, the beta comes out at point nine.
- The return the manager should earn is two plus point nine times six, which is seven point four.
- Alpha looks like two point six per cent, and the manager appears skilled.
- But the fund's true beta is one point two.
- Then the required return is two plus one point two times six, which is nine point two, and alpha is only point eight.
- The poor proxy overstated alpha by one point eight points.
- A trustee reading the first number would reward a manager who was mostly just taking more market risk.

24 Small samples and other weak spots

- Short records give noisy ratios
- Sortino needs enough below-target observations
- Drawdown rests on one or two episodes and the window
- Option sellers look calm until the tail arrives
- Smoothed illiquid values flatter most measures

IN THE LESSON

- There are more limits to remember.
- Every appraisal measure is a backward looking estimate from a single sample, so a short record gives a noisy number, and the same fund could easily have scored higher or lower.
- The Sortino ratio needs enough returns below the target to estimate the downside deviation, and with only a few shortfalls the figure is unreliable.
- Maximum drawdown and its duration are path dependent.
- They are driven by one or two episodes, and they change if the observation window moves.
- Strategies that sell options look calm for years and then lose heavily, so a calm record hides the tail.
- And smoothed values for illiquid assets flatter most of the measures at once.

25 A story: the lucky coin flippers

- 1,024 people flip a coin ten times each
- Ten heads in a row: chance 1 in 1,024
- So about one flipper looks like a genius
- A strong record alone is not skill

IN THE LESSON

- Suppose one thousand and twenty four people each flip a fair coin ten times.
- The chance that any one of them gets ten heads in a row is one in one thousand and twenty four.
- So we should expect about one flipper to do it, purely by luck.
- Nobody would call that person a coin-flipping genius.
- Fund managers are in the same position.
- With thousands of managers, some will produce a great five year record by chance.
- Skill means value added that is repeatable, and that is not explained by luck or by exposures a cheap index could supply.
- Mara Quinn has to decide whether Fenner is a genuine stock picker.
- She asks three things: is it measured correctly, is it significant, and did it come from the claimed edge.

26 How long skill takes to show

IR = active return ÷ tracking error
t ≈ IR × √years
Years for t = 2: (2 ÷ IR)²

Information ratio	Active return ÷ tracking error
Years of record	Years of independent returns
t-statistic	About 2 or more is convincing
Years needed	Years the record must last

IN THE LESSON

- Now the statistics.
- Start with the information ratio, which is the average active return divided by tracking error, the standard deviation of the active return.
- To test whether the active return is more than luck, use the t statistic.
- It is roughly the information ratio times the square root of the number of years, provided the yearly returns are independent over time.
- A t statistic of about two is the usual hurdle.
- Turn the formula around and you can ask how many years are needed.
- Divide two by the information ratio, and square the result.
- The square root means evidence builds slowly.
- Four times the data only doubles the t statistic.

27 Example 8: is Fenner skilled?

Fenner: active return 3.0% a year, tracking error 5.0%, record 4 years

Information ratio	3.0% ÷ 5.0%	**0.60**
t statistic	0.60 × √4	**1.20**
Years needed for t = 2	(2 ÷ 0.60)²	**11.1**
More years than the record	11.1 − 4	**7.1**
If the ratio were 1.0	(2 ÷ 1.0)²	**4 years**

IN THE LESSON

- Fenner calls itself a stock picker.
- Over four years its active return averaged three per cent a year, with tracking error of five per cent.
- The information ratio is three over five, which is point six.
- The t statistic is point six times the square root of four, which is one point two.
- That is well below two.
- To reach a t statistic of two at this rate would take two divided by point six, squared, which is eleven point one years.
- That is about seven years more than the record.
- Compare a rare manager with an information ratio of one, who would need only four years.
- Fenner may well be skilled, but the numbers cannot yet tell it apart from a lucky manager.

28 Years of record needed

Chart: t-statistic vs Years of record

- IR 1.0: 4 years
- IR 0.5: 16 years
- t = 2

IN THE LESSON

- Here is the same idea as a picture.
- The horizontal axis is the years of record, up to sixteen.
- The vertical axis is the t statistic.
- The dashed horizontal line marks a t statistic of two, the usual hurdle.
- The upper curve is a manager with an information ratio of one.
- It reaches the dashed line at four years.
- The lower curve is a manager with an information ratio of point five, still a good long run result.
- It reaches the line only at sixteen years.
- Both curves flatten, because evidence grows with the square root of time.
- A record of only a few years sits well to the left of where the lower curve becomes convincing, which is why a short, strong record proves little.

29 Where Fenner's gain came from

Active return (% a year)

Sector tilt	**2.2**
Stock selection	**0.8**

Total active return 3.0% a year

IN THE LESSON

- Statistics are only half of the check.
- Attribution tells you where the gain came from.
- Fenner says its edge is stock selection.
- Its active return is three per cent a year.
- Mara Quinn's attribution splits this.
- The first bar shows two point two points from a sector tilt, a persistent overweight in cheap value shares.
- The second bar shows only point eight points from picking individual stocks.
- So most of the gain came from an exposure that a low cost value index could supply for a fraction of the fee, not from the skill the manager is selling.
- If a stock picker's value does not come from selection, the claimed edge is not showing up in the numbers.

30 A checklist for judging skill

Four questions before calling a result skill

Question	Evidence
Measured against the right benchmark?	alpha, IR
Statistically significant?	t near 2
From the claimed edge?	attribution
Repeatable and believable?	persistence

IN THE LESSON

- Put the pieces together as a checklist.
- First, was the result measured against the right benchmark, using alpha and the information ratio?
- A wrong benchmark makes everything else meaningless.
- Second, is it statistically significant?
- Look for a t statistic near two or more, and remember how many years that can take.
- Third, did the gain come from the claimed edge?
- Use attribution, and look for selection in a stock picker, not allocation or factor tilts.
- Fourth, is it repeatable and believable?
- Check that results persist across periods, and test the story against the manager's philosophy, people and process.
- The trap is to infer skill from a short, strong record.

31 Our recommendation for Brindle

- Hold Fenner; do not add money yet
- t of 1.2 after 4 years: far from proof
- 2.2 of the 3.0 points came from a sector tilt
- Review yearly: look for positive selection

IN THE LESSON

- Now Mara Quinn must decide, and a decision needs reasons.
- Her advice to the trustees is to hold Fenner but not to add money yet.
- The reasons are three.
- The t statistic is one point two after four years, which is far from proof.
- Most of the active return, two point two points of three, came from a sector tilt and not from stock picks.
- And four years is a short record.
- What could go wrong with this advice?
- Fenner might be skilled, and the trustees might miss out on gains.
- So she suggests they keep reviewing each year, look for a positive selection effect, and ask Fenner to explain the tilt.
- If selection builds up, the case for adding strengthens.

32 Recap: the lesson in five lines

- Alternatives: blend of custom, factor and peer benchmarks
- Drawdown, Sortino, appraisal, capture: the path
- Each ratio rests on assumptions that can fail
- Skill: t near 2, gains from the claimed edge
- Short strong records prove little

IN THE LESSON

- Let us pull the lesson together.
- First, alternatives break the tests of a good benchmark, so use a blend: custom, factor and peer comparisons, with money weighted returns and a public market equivalent for private assets.
- Second, drawdown measures the worst fall and the duration runs until the peak is regained, and the Sortino ratio divides by downside deviation.
- Third, the appraisal ratio is alpha over residual risk, and the capture ratio is upside over downside.
- Fourth, every measure rests on assumptions about smoothing, the benchmark, the target and the window.
- Fifth, skill needs a significant t statistic, the right source of gains, and persistence over time.

33 Try three quick questions

1. Peak 150, trough 120. Maximum drawdown?
   Answer: (120 − 150) ÷ 150 = −20%.
2. Capture 108% up and 90% down. Capture ratio?
   Answer: 108 ÷ 90 = 1.2, a favourable lopsidedness.
3. Information ratio 0.4. Years for t = 2?
   Answer: $(2 \div 0.4)^2 = 5^2 = 25$ years.

IN THE LESSON

- Now try three quick questions yourself.
- Pause after each one.
- Question one.
- A portfolio peaks at one hundred and fifty and falls to a trough of one hundred and twenty.
- What is the maximum drawdown?
- One hundred and twenty minus one hundred and fifty is minus thirty, over one hundred and fifty, which is minus twenty per cent.
- Question two.
- A manager captures one hundred and eight per cent of the index in rising periods and ninety per cent in falling periods.
- What is the capture ratio?
- One hundred and eight divided by ninety is one point two.
- Question three.
- A manager has an information ratio of point four.
- How many years are needed to reach a t statistic of two?
- Two over point four is five, and five squared is twenty five years.
- That is a lot of patience.