Hedge funds, risk and the portfolio

Alternative Investments

1 Starting from the economy, not the stock

- Top-down: start from the economy, not one stock
- Bets across rates, currencies, commodities, shares
- Global macro = judgement
- Managed futures = rules

IN THE LESSON

- Think of two weather forecasters.
- One studies a single cloud, in great detail.
- The other looks at the whole map, the winds, the oceans and the seasons, and says: it is going to be a wet winter.
- Most hedge funds are the first kind.
- They pick individual securities.
- Opportunistic funds are the second kind.
- They are top-down.
- They begin with the economy and ask where interest rates, currencies, growth and prices are heading, and then they place bets across many markets.
- There are two main styles.
- Global macro funds use judgement.
- Managed futures funds use rules.

2 Global macro: a view on the world

- Discretionary views on growth, rates, currencies
- Expressed with futures, forwards and options
- Depends on skill and timing, so returns are lumpy
- Has tended to hold up when markets are stressed

IN THE LESSON

- A global macro manager forms views.
- Perhaps she thinks a central bank will raise rates faster than the market expects, or that a currency is too strong for its country's trade position.
- She then expresses the view with futures, forwards and options, because they are cheap and liquid, and she often uses leverage.
- Notice that the views are discretionary.
- A person decides, using judgement about growth, inflation and policy.
- That has two consequences.
- Results depend heavily on skill and timing, so returns are lumpy: a few big wins and long quiet stretches.
- And because the manager can position before a shock rather than after it, the strategy has tended to hold up in market stress.

3 Managed futures: following the trend

A trend rule

Look back at the recent trend →
Rising: go long. Falling: go short →
Trade liquid futures

IN THE LESSON

- Managed futures funds are run by commodity trading advisers, and most of them are systematic.
- That means a computer follows a fixed rule, with no gut feeling involved.
- The commonest rule is trend following, also called time-series momentum.
- Picture the process as three boxes.
- First, look back at the recent price trend of a market.
- Second, turn that into a signal: if the price has been rising, go long, and if it has been falling, go short.
- Third, trade the signal in liquid futures, across dozens of markets at once, shares, bonds, currencies and commodities.

4 Example 1: a trend rule in two markets

Rule: long after a rise, short after a fall. Starts long.

A: equities	100×0.95^4	**81.5**
A: trend follower	$100 \times 0.95 \times 1.05^3$	**110.0**
B: equities	$100 \times 0.90 \times 1.08$	**97.2**
B: trend follower	$100 \times 0.90 \times 0.92$	**82.8**

A: four months of −5%. B: −10%, then +8%.

IN THE LESSON

- Let us test a simple trend rule on one equity index future.
- The rule goes long after a month in which the price rose, and short after a month in which it fell.
- It starts the story long.
- In scenario A, the index falls five per cent a month for four months.
- Equities alone end at eighty one point five.
- The trend follower loses in month one, because it was still long, then flips short and gains five per cent in each of the next three months.
- That gives one hundred times zero point nine five times one point oh five cubed, which is one hundred and ten.
- In scenario B, the index falls ten per cent, then rises eight.
- The trend follower is long for the fall and short for the rebound, so it loses twice and ends at eighty two point eight.

5 Same rule, two very different outcomes

Value of 100 invested

A: equities	**81.5**
A: trend	**110.0**
B: equities	**97.2**
B: trend	**82.8**

Trend following needs a trend to last

IN THE LESSON

- Here are those four results as bars, each starting from one hundred.
- The first two bars are scenario A, a long, steady slide.
- Equities finish at eighty one point five, down almost a fifth.
- The trend follower finishes at one hundred and ten.
- That is the heart of the strategy: a trend that runs a long way pays, and that is why the returns are positively skewed and strongest in prolonged sell-offs.
- Now the last two bars, scenario B, a sharp drop that reverses at once.
- Equities end close to where they started, at ninety seven point two, but the trend follower falls to eighty two point eight.
- It needed the trend to last, and the shock was over before the signal could turn.

6 Macro and managed futures side by side

Both are top-down, but they decide very differently

How they compare	Answer
Global macro decisions	Judgement
Managed futures decisions	A rule
Managed futures skew	Positive
Weak spot of trend following	Reversals

IN THE LESSON

- Let us put the two opportunistic strategies side by side.
- Global macro decisions come from a person's judgement.
- Managed futures decisions come from a rule.
- That is why macro returns are lumpy and depend on one mind being right at the right time, while managed futures returns tend to be positively skewed: many small losses while the market chops about, then a large gain when a long trend arrives.
- The weak spot of trend following is a reversal, as we just saw.
- In choppy markets with no trend, the rule keeps buying high and selling low, and it bleeds.

7 Specialist funds: volatility as an asset

- Volatility traded directly: options, futures, variance swaps
- Long volatility: pay premium, gain in a sell-off
- Short volatility: collect premium, risk a severe loss
- A tail hedge versus a seller of insurance

IN THE LESSON

- Now the specialists.
- These funds need unusual skills or unusual access, and they aim for returns that are unrelated to the usual stock and bond markets.
- Start with volatility traders.
- Volatility is how violently prices move, and it can be traded in its own right, through options, futures on a volatility index, and variance swaps.
- A long volatility book is like paying for fire insurance.
- Each calm month you pay a little premium, which is lost as the options fade.
- But in a sharp sell-off, volatility leaps and the book gains a lot.
- That makes it a tail hedge.
- A short volatility book is the other side of the contract.
- It collects the premium steadily, like an insurer, and it suffers a severe loss when volatility jumps.

8 Long and short volatility, calm and storm

Each side of the contract earns in a different weather

Position	Result
Long volatility, calm month	Loses premium
Long volatility, sell-off	Large gain
Short volatility, calm month	Collects premium
Short volatility, sell-off	Large loss

IN THE LESSON

- Here is the same idea as a little grid.
- Take a long volatility book first.
- In a calm month it loses a little, because the options it bought decay and it pays the premium.
- In a sell-off it makes a large gain.
- Now the short volatility book.
- In a calm month it collects its premium, small and steady, and the record looks wonderfully smooth.
- In a sell-off it takes a large loss, and the loss can wipe out years of premium.
- The trap, as always, is to trust the smooth record.
- A short volatility fund with five quiet years behind it has not proved that it is safe.
- It may simply not have met the storm yet.

9 Insurance risk: life settlements and catastrophes

- Life settlement: buy a policy, pay premiums, collect
- Return depends mainly on how long the insured lives
- Catastrophe reinsurance: premiums in, disasters paid out
- Driven by events, not by market prices

IN THE LESSON

- The other specialist strategies buy insurance risk.
- In a life settlement, the fund buys a life insurance policy from its holder, for more than the policy would pay if the holder cashed it in, but for less than the death benefit.
- The fund then keeps paying the premiums until the insured person dies, and collects the death benefit.
- So the fund is really betting on a lifespan.
- The longer the insured lives, the more premiums the fund pays and the later the big payment arrives.
- What matters is how long someone lives, which no stock market sets.
- Catastrophe reinsurance is similar: premiums arrive every year, and a hurricane or earthquake triggers a payout.

10 Example 2: a life settlement

Price $200k, premium $40k a year, benefit $1,000k

Outlay, lives 8 years	200 + 8 × 40	**$520k**
Return, lives 8 years	what the price assumed	**11.2%**
Lives 5 years	outlay $400k	**24.4%**
Lives 12 years	outlay $680k	**4.7%**

Yearly return on the policy, in thousands of dollars

IN THE LESSON

- Let us put numbers on a life settlement.
- A fund pays two hundred thousand dollars for a policy with a one million dollar death benefit.
- It must pay forty thousand dollars of premium at the start of each year until the insured dies, and the benefit arrives at the end of that final year.
- The medical estimate is eight more years.
- Over eight years the fund pays the price plus eight premiums, which is five hundred and twenty thousand dollars, and it receives one million.
- The internal rate of return, the yearly rate that makes the payments and the receipt balance, is eleven point two per cent, which is what the price assumed.
- If the insured lives only five years, the outlay is four hundred thousand and the return jumps to twenty four point four per cent.
- If they live twelve years, the outlay is six hundred and eighty thousand and the return falls to four point seven per cent.

11 Paid by how long someone lives

Chart: Yearly return vs Years the insured lives

- 8 years: 11.2%
- Dies early: big return
- Lives long: small return

IN THE LESSON

- Here is that example as a picture.
- Along the bottom is the number of years the insured lives, from four to fourteen.
- Up the side is the yearly return to the fund.
- Follow the curve from the left.
- If the insured dies early, the return is enormous, above thirty per cent, because the large payment arrives quickly after little premium has been paid.
- As the years go by, the curve falls steeply and keeps falling, because every extra year costs another premium and pushes the payment further away.
- The labelled point is the medical estimate, eight years, which gives eleven point two per cent.
- Past twelve years the return is under five per cent.
- Nothing on the horizontal axis is a market price, so a stock market crash does not move it.
- That is where the low correlation comes from.

12 The trap: low correlation is not low risk

- Low correlation tells you the bad day has a different cause
- It does not tell you how bad, or how likely
- Risks are idiosyncratic and sit in the tails
- Funding costs and the insurer's credit still matter

IN THE LESSON

- Here is the most important lesson about specialist strategies.
- Their returns are largely unrelated to the stock market, and that is genuinely valuable for diversification.
- But unrelated is not the same as safe.
- The life settlement fund can earn four point seven per cent or twenty four point four per cent on the same policy, depending on one variable that nobody can control.
- These risks are idiosyncratic, they sit in the tails, and they are hard to model.
- A low correlation only tells you that the bad outcome has a different cause.
- It tells you nothing about how big, or how likely, that outcome is.
- Investors are paid for bearing exactly those tails.

13 Multi-manager funds: a fund that holds funds

- Fund of funds: allocates to external managers
- Gives access, due diligence, diversification
- Multi-strategy: several teams inside one firm
- One fee on net performance, capital moves quickly

IN THE LESSON

- Now the multi-manager funds, which combine strategies to diversify.
- The fund of funds is like a food critic who picks restaurants for you.
- It allocates your money across external hedge fund managers.
- It offers access to managers you could not reach alone, careful due diligence, and diversification across firms.
- The price is a second layer of fees, and something called netting risk, which we will measure shortly.
- The other approach is a multi-strategy fund.
- One firm runs several strategies inside it, with its own teams, and moves capital between them quickly.
- It charges one fee, on the net performance.

14 Example 3: the fee layers, step by step

Two managers, $100m each. Fees: 20% each, fund of funds 10%.

Gross result	+20 − 10	**+$10m**
Manager A's fee	20% × 20	**$4m**
After manager fees	10 − 4	**$6m**
Fund of funds fee	10% × 6	**$0.6m**
Investor keeps	6 − 0.6	**$5.4m**

Multi-strategy: one fee, 10 − 20% × 10 = $8m

IN THE LESSON

- Here is how netting risk works.
- An investor puts two hundred million dollars into a fund of funds, split equally between Manager A and Manager B.
- Each manager charges an incentive fee of twenty per cent of its own gains, and the fund of funds charges ten per cent of its gains.
- Manager A gains twenty per cent and Manager B loses ten per cent.
- In dollars, A makes twenty million and B loses ten million, a gross gain of ten million, which is five per cent.
- Now the fees.
- A charges twenty per cent of twenty million, which is four million.
- B earned nothing, so it charges nothing, but it also refunds nothing.
- After manager fees six million is left.
- The fund of funds takes ten per cent of that, six hundred thousand, and the investor keeps five point four million.

15 What reaches the investor

Return on the $200m

Gross	**5.0%**
Multi-strategy	**4.0%**
Fund of funds	**2.7%**

Same trading, three different results

IN THE LESSON

- Let us compare the three results as bars, each as a return on the two hundred million.
- The first bar is the gross return, five per cent.
- That is what the underlying trading earned, and it is the number that often appears in brochures.
- The second bar is the multi-strategy fund.
- It charges twenty per cent of the net gain of ten million, which is two million, leaving eight million, or four per cent.
- The third bar is the fund of funds.
- After both layers of fees the investor keeps five point four million, which is two point seven per cent.
- The fund of funds investor paid forty six per cent of the net gain in fees, not twenty.
- Netting turned a nominal twenty per cent fee into forty per cent of the net gain, before the fund of funds added its own layer.

16 Fund of funds versus multi-strategy

Different structures put the costs and risks in different places

Compare	Answer
Fund of funds: fee layers	Two
Multi-strategy: fee layers	One
Fund of funds: netting risk	Investor
Multi-strategy: netting risk	Manager

IN THE LESSON

- So how do the two structures compare?
- A fund of funds has two layers of fees, the managers and the fund itself, and the investor carries the netting risk: a fee is paid on each winner, with no credit for the losers.
- A multi-strategy fund has one layer of fees on net performance, so the netting problem lands on the manager, who still has to reward the team that made money while another team lost.
- But there are costs on that side too.
- Operational risk and key-person risk are concentrated in one firm, and leverage is often higher.
- Neither structure is simply better.
- What you cannot do is compare their gross returns and call it a day.

17 Four mistakes to avoid so far

- Trend followers need a trend to form
- Low correlation is not low risk
- Gross returns hide the second fee layer and netting
- Judge the effect on the whole portfolio

IN THE LESSON

- Before we change gears, here are four mistakes that exam questions love.
- One: expecting trend followers to protect you against any shock.
- They need a trend to form, so a sudden drop that reverses before the signal turns can hurt them.
- Two: treating low correlation as low risk.
- Specialist strategies carry idiosyncratic and tail risks that are hard to model.
- Three: comparing gross returns across structures.
- A fund of funds has a second layer of fees, and netting means you pay incentive fees on winners even when losers cancel them out.
- And four, which we will meet in the next section: judging a strategy by its own figures, rather than by what it does to the whole portfolio.

18 Is it skill, or hidden market risk?

- Explain returns with market factors
- Equities, credit, currencies, rates, commodities
- What is left over is alpha
- Reveals hidden beta behind apparent skill

IN THE LESSON

- A hedge fund has made nine per cent a year with low volatility.
- Is that skill?
- Before paying skill fees, a sensible allocator asks something else: how much of this could ordinary market factors have earned?
- A factor model is the taste test.
- It explains the fund's returns using exposures to market factors, such as equity returns, credit spreads, currencies, interest rates, commodities and volatility.
- Whatever the factors cannot explain is called alpha.
- If most of a fund's return is explained by a credit factor, then it is really selling credit risk, whatever it calls itself, and an investor could probably buy that exposure far more cheaply.

19 The factor model

Fund return = alpha + sum of (beta × factor return) + error

Fund return	Month by month
Alpha	What the factors do not explain
Beta	Exposure to one factor
Factor return	Equities, credit, rates...
Error	Noise left over

IN THE LESSON

- Here is the factor model written down.
- A fund's return in a month equals alpha, plus the sum of each beta times its factor return, plus an error.
- Take the pieces one at a time.
- A factor return is how a chosen market factor did that month, for instance equities, or credit spreads.
- A beta is the fund's exposure to that factor.
- A beta of point three to equities means the fund tends to move about thirty cents for every dollar the equity market moves.
- Alpha is what is left once the factor exposures are paid for.
- It is the part that looks like skill, and it is only as credible as the betas that were taken out.
- The error is month to month noise.

20 Reading the regression output

Illustrative fund. Alpha is 0.3% a month in both regimes.

Beta of the fund	Estimate
Equity, normal months	0.2
Equity, crisis months	0.6
Credit, normal months	0.1
Credit, crisis months	0.5

IN THE LESSON

- Let us read some output.
- An analyst regresses a fund's monthly returns on an equity factor and a credit factor, and separately estimates the betas for normal months and for crisis months.
- Alpha comes out at point three per cent a month in both.
- Now look at the betas.
- In normal months the equity beta is point two and the credit beta is point one.
- The fund barely responds to the market, and that is the picture its marketing would show.
- In crisis months the equity beta is point six and the credit beta is point five.
- The equity exposure has tripled, and the credit exposure has risen fivefold.
- This is exactly why one set of average betas can mislead.

21 A conditional model: add a stress dummy

Fund return = alpha + sum of (beta × factor) plus D × sum of (extra beta × factor)

Stress dummy D	1 in a stress month, else 0
Beta	Normal-month exposure
Extra beta	Added exposure in a crisis
Crisis beta	Crisis beta = beta + extra
Alpha	Same in both regimes

IN THE LESSON

- To let the betas change in a crisis, we use a conditional model.
- We add a dummy variable, called D, which is just a switch.
- It equals one in a stress month, and zero in every other month.
- Then we add a second set of terms that are switched on only when D is one.
- Each of them is an extra beta times the factor return.
- In a normal month, D is zero, those terms vanish, and only the normal betas remain.
- In a stress month, D is one, and the crisis beta is the normal beta plus the extra beta.
- In our example, the equity beta of point two plus an extra point four gives point six.
- So in each regime, the expected return is alpha plus the sum of that regime's betas times the factor returns.

22 Example 4: a calm month and a crisis month

Calm: equities +1%, credit +0.3%. Crisis: −8% and −4%.

Calm month, normal betas	0.3 + 0.2×1 + 0.1×0.3	**0.53%**
Crisis, crisis betas	0.3 + 0.6×(−8) + 0.5×(−4)	**−6.5%**
Crisis, full-period betas	0.3 + 0.3×(−8) + 0.2×(−4)	**−2.9%**
Understated by	6.5 − 2.9	**3.6 points**

Full-period betas: equity 0.3, credit 0.2

IN THE LESSON

- Now let us use those betas.
- In a calm month, equities return one per cent and the credit factor returns point three per cent.
- With normal betas, the fund returns point three plus point two times one, plus point one times point three, which is point five three per cent.
- Now a crisis month: equities fall eight per cent, and the credit factor falls four.
- The conditional model uses crisis betas: point three plus point six times minus eight, plus point five times minus four.
- That is minus six point five per cent.
- What would the single full-period regression have said?
- Its averaged betas are point three and point two.
- It predicts point three plus point three times minus eight, plus point two times minus four, which is minus two point nine per cent.
- The average model understates the crisis loss by three point six points, in the one month that matters.

23 Averages hide the crisis

Expected loss in a crisis month

Full-period betas	**2.9%**
Crisis betas	**6.5%**

Crisis month: equities −8%, credit −4%

IN THE LESSON

- Here are the two crisis predictions as bars, in percentage losses for the month.
- The first bar is the full-period model, which expects a loss of two point nine per cent.
- It averages the many calm months with the few stressed ones, so the crisis gets diluted.
- The second bar is the crisis model, which expects a loss of six point five per cent, more than twice as much.
- The fund behaves much more like an equity and credit position than its average betas suggest, and it does so just when an investor was counting on diversification.
- With full-period betas, part of the crisis loss leaks into the error term; the conditional model keeps the regimes apart.
- The lesson is simple: do not rely on full-period average exposures.

24 What factor models tell you

- Separates alpha from beta bought cheaply elsewhere
- Exposes hidden beta behind a smooth record
- A stress dummy lets betas change in a crisis
- Trap: full-period averages hide crisis exposure

IN THE LESSON

- Let us gather up what a factor model gives an investor.
- First, it separates beta from alpha.
- A fund that earns nine per cent mostly because it holds equity and credit risk is not offering skill, it is offering a bundle of market exposures that could be bought more cheaply.
- Second, it shows the exposures behind a smooth record.
- A strategy that looks like a bond fund may carry a hidden credit beta.
- Third, a conditional model shows how exposures change in stress, which is when diversification matters most.
- A strategy with a modest equity beta in normal times can carry a much higher equity and credit beta in a crisis.
- And the trap: relying on full-period average exposures, which hide exactly that.

25 Adding a hedge fund to what you already own

- You hold the fund inside a portfolio, not alone
- Return, volatility and tail risk all change
- Judge the whole portfolio, not the standalone figures
- Sharpe ratio is the first test, the worst year the second

IN THE LESSON

- Now the practical question.
- You rarely hold a hedge fund alone.
- You add one to a portfolio of shares and bonds, so the right question is not how good the fund looks by itself, but what it does to the whole portfolio.
- Adding a strategy changes three things at once: the portfolio's return, its volatility, and its tail risk, meaning the worst outcomes.
- A strategy can lower volatility and raise the Sharpe ratio even while trimming return, and still disappoint in the worst year.
- We will measure all three.

26 Return and risk of a two-part portfolio

Return = w1×R1 + w2×R2
Variance = $w1^2SD1^2 + w2^2SD2^2 + 2\ w1\ w2\ Corr\ SD1\ SD2$

Weights w	Share of money, adds to 1
Returns R	Expected return of each
Volatilities SD	Volatility of each
Correlation	How much they move together
Variance	Volatility is its square root

IN THE LESSON

- Two formulas do the work.
- The portfolio's expected return is just the weighted average of the returns of its parts.
- The weights are the shares of your money in each, and they must add up to one.
- Risk is trickier, because risk does not average.
- Portfolio variance has three pieces: the first fund's weight times its volatility, squared; the second fund's weight times its volatility, squared; and a term for how they move together, which is two times the two weights, times the correlation, times the two volatilities.
- The correlation sits in the last term.
- When it is low, the last term shrinks, and the whole variance falls below what a simple average would suggest.
- Volatility is the square root of variance.

27 The Sharpe ratio

Sharpe ratio = (portfolio return − risk-free rate) ÷ volatility

Portfolio return	Expected return of the portfolio
Risk-free rate	Return on a safe deposit
Portfolio volatility	Standard deviation of returns

IN THE LESSON

- To compare portfolios fairly, we use the Sharpe ratio.
- It asks: how much extra return do I earn over a safe deposit, for each unit of volatility I put up with?
- Take the portfolio's expected return, subtract the risk-free rate, and divide by the portfolio's volatility.
- A higher Sharpe ratio means more reward for each unit of risk.
- It is also the slope of a line from the risk-free rate to the portfolio on a chart of return against volatility, which is why a strategy can lower return and still help.

28 Example 5: return and risk, step by step

60/40 portfolio: 6.4%, volatility 10%. Fund: 6.0%, 7%, corr 0.3

Expected return	$0.8 \times 6.4 + 0.2 \times 6.0$	**6.32%**
Own-risk terms	$0.8^2 \times 10^2 + 0.2^2 \times 7^2$	**65.96**
Covariance term	$2 \times 0.8 \times 0.2 \times 0.3 \times 10 \times 7$	**6.72**
Variance	65.96 + 6.72	**72.68**

20% of the money moves into the hedge fund

IN THE LESSON

- Here is an illustrative case.
- A portfolio of sixty per cent equities and forty per cent bonds has an expected return of six point four per cent and a volatility of ten per cent.
- A hedge fund strategy offers six per cent with seven per cent volatility, and a correlation of point three with the portfolio.
- The investor moves twenty per cent of the money into the strategy, so the weights are point eight and point two.
- Expected return is point eight times six point four, plus point two times six, which is six point three two per cent, slightly lower.
- Now variance.
- The own-risk terms are point eight squared times ten squared, plus point two squared times seven squared, which is sixty five point nine six.
- The covariance term is two times point eight times point two times point three times ten times seven, which is six point seven two.
- Add them and the variance is seventy two point six eight.

29 Example 5, continued: volatility and Sharpe

Risk-free rate 2%. Compare before and after.

Volatility	$\sqrt{72.68}$	**8.53%**
Sharpe ratio, before	$(6.4 - 2) \div 10$	**0.44**
Sharpe ratio, after	$(6.32 - 2) \div 8.53$	**0.51**

Weighted average of volatilities would be 9.4%

IN THE LESSON

- Now finish the job.
- Volatility is the square root of seventy two point six eight, which is eight point five three per cent.
- Compare that with ten per cent before.
- The investor gave up point zero eight of a point of return and bought one point four seven points less volatility.
- A plain weighted average of the two volatilities would have been nine point four per cent, so the low correlation did real work.
- Now the Sharpe ratio.
- Before the change it was six point four minus two, divided by ten, which is point four four.
- After, it is six point three two minus two, divided by eight point five three, which is point five one.
- The reward per unit of risk has gone up, even though the return went down.

30 Less risk for a little less return

Portfolio volatility (% a year)

60/40 alone	**10.0%**
With 20% fund	**8.53%**

Return only slipped from 6.40% to 6.32%

IN THE LESSON

- Here is the volatility result as bars.
- The left bar is the sixty forty portfolio on its own, ten per cent.
- The right bar is the portfolio with twenty per cent in the hedge fund strategy, eight point five three per cent.
- The right bar is clearly shorter, and the expected return barely moved, from six point four to six point three two.
- Why is the drop so large?
- Because the strategy's ups and downs only partly coincide with the portfolio's.
- At a correlation of point three, some of its losses are covered by gains elsewhere.
- If the correlation had been one, there would have been no benefit beyond a plain average.
- The lower the correlation, the larger the saving.

31 The second test: the worst year

Portfolio falls 20%. Managed futures +5%, merger arbitrage −8%.

Stress year result	Portfolio
No hedge fund	−20.0%
20% in managed futures	−15.0%
20% in merger arbitrage	−17.6%

IN THE LESSON

- The Sharpe ratio is the first test.
- The second is the worst year, because average figures hide tail risk.
- Suppose the traditional portfolio loses twenty per cent in a stress year.
- A strategy that has tended to help in crises, such as managed futures or global macro, might gain five per cent.
- With twenty per cent of the money in it, the portfolio's loss is point eight times minus twenty, plus point two times five, which is minus fifteen per cent.
- Now merger arbitrage.
- It looks diversifying on average, but it tends to fall with equities when deals break.
- If it loses eight per cent, the portfolio loses seventeen point six.
- Two strategies that look alike on average leave very different worst years.
- Judge the effect on the whole portfolio, not the strategy's standalone figures.

32 Recap: the ideas to remember

- Macro = judgement, trend = rules, both diversify
- Specialist: unrelated to markets, not low risk
- Fund of funds: two fee layers and netting risk
- Factor models: use crisis betas, not averages
- Judge a strategy by its effect on the whole portfolio

IN THE LESSON

- Let us recap.
- Opportunistic funds are top-down: global macro uses judgement, and its returns are lumpy; managed futures follow trends by rule, with positive skew, strongest in long sell-offs and weakest when markets reverse or chop.
- Specialist funds, such as volatility traders and life settlement funds, earn returns unrelated to traditional markets, but they carry idiosyncratic tail risks, so low correlation is not low risk.
- Multi-manager funds: a fund of funds has two fee layers and netting risk, a multi-strategy fund has one fee on net but concentrates operational and key-person risk.
- Factor models split return into beta and alpha, and a stress dummy shows crisis betas.
- And when you add a strategy, judge return, volatility, Sharpe ratio and the worst year of the whole portfolio.

33 Try three quick questions

1. Alpha 0.2%, crisis betas 0.5 and 0.3, factors −6% and −2%. Expected return?
   Answer: 0.2 − 3.0 − 0.6 = −3.4%.
2. A gains $10m, B loses $4m, fees 20% of own gains. Fee, versus 20% of the net?
   Answer: $2.0m, against 20% × 6 = $1.2m.
3. Return 7%, volatility 10%, risk-free 2%. Sharpe ratio?
   Answer: (7 − 2) ÷ 10 = 0.5.

IN THE LESSON

- Now try three quick questions.
- Pause after each one.
- Question one.
- A fund has alpha of point two per cent, a crisis equity beta of point five and a crisis credit beta of point three.
- In a crisis month equities fall six per cent and the credit factor falls two.
- What return does the conditional model expect?
- Point two, minus three, minus point six, which is minus three point four per cent.
- Question two.
- Manager A gains ten million and Manager B loses four million, each charging twenty per cent of its own gains.
- What is the total incentive fee, and what would one multi-strategy fee on the net be?
- Two million, against one point two million.
- Question three.
- A portfolio expects seven per cent with ten per cent volatility, and the risk-free rate is two per cent.
- What is the Sharpe ratio?
- Five divided by ten, which is point five.