CHARTA LEARNING · LESSON NOTES

Portfolio Performance Evaluation

FRM Part II

1 A story: did the manager add value?

- Kestrel hires Larch and Birch to beat the market
- The question: skill, or just riding the market?
- Step 1: measure the return fairly
- Step 2: adjust the return for risk
- Step 3: skill or luck, and where did it come from?

IN THE LESSON

- The Kestrel Pension Plan pays two outside equity managers, Larch Capital and Birch Partners, to beat the market.
- Each year its committee asks one question: did the manager add value, or simply ride the market?
- Answering it takes three steps.
- First, measure the return in a way that is fair to the manager.
- Second, adjust that return for the risk taken to earn it.
- Third, ask whether the result is skill or luck, and where it came from.
- We take the steps in that order, starting with a surprisingly slippery question: what was the return?

2 Two ways to measure a return

TWR = (1 + r1) × (1 + r2) − 1
DWR = IRR of the cash flows in and out

Return r1	Return before the cash flow
Return r2	Return after the cash flow
Time-weighted return	Ignores size and timing of flows
Dollar-weighted return	Reflects when money came and went

IN THE LESSON

- Two returns answer different questions.
- The time weighted return breaks the period at each external cash flow, meaning money the client adds or withdraws.
- The portfolio is valued just before each flow, and the sub period returns are chained.
- With two sub periods, where r one and r two are the returns before and after the flow, it is one plus r one, times one plus r two, minus one.
- Because each sub period counts equally, the size and timing of the flows have no effect.
- The dollar weighted return is the internal rate of return, I R R, on the portfolio's cash flows, so it also reflects when money came in or went out.

3 Example 1: Larch's time-weighted return

Larch starts with $100m. Year 1: +20%. Kestrel then adds $100m. Year 2: −10%

Value after year 1	100 × 1.20	**$120m**
After Kestrel adds $100m	120 + 100	**$220m**
Value after year 2	220 × 0.90	**$198m**
Time-weighted, two years	1.20 × 0.90 − 1	**8.0%**
Time-weighted, a year	√1.08 − 1	**3.92%**

IN THE LESSON

- Here are Larch's first two years with Kestrel.
- Larch starts with one hundred million dollars and earns twenty per cent in year one.
- Then Kestrel adds another one hundred million, and in year two Larch loses ten per cent.
- After year one the account is worth one hundred and twenty million, the deposit takes it to two hundred and twenty, and after year two it is two hundred and twenty times point nine, one hundred and ninety eight million.
- The time weighted return chains the two sub period returns: one point two times point nine, minus one, is eight per cent over two years.
- The square root of one point zero eight, minus one, gives about three point nine two per cent a year.

4 Example 2: the dollar-weighted return

IRR solves: 100 + 100 ÷ (1 + r) = 198 ÷ (1 + r)^2

Money put in	100 now + 100 after year 1	**$200m**
Value taken out	after year 2	**$198m**
Dollar-weighted return	IRR, a year	**−0.67%**
Time-weighted return	from Example 1, a year	**+3.92%**

Kestrel's timing of the deposit, not Larch's skill, cut the DWR

IN THE LESSON

- Now the dollar weighted return on the same account.
- The I R R is the rate r that solves the equation at the top: one hundred, plus one hundred discounted one year, equals one hundred and ninety eight discounted two years.
- Kestrel put in two hundred million dollars in total, and could take out one hundred and ninety eight million at the end.
- Solving gives a dollar weighted return of about minus point six seven per cent a year.
- Compare the time weighted return of plus three point nine two per cent.
- The gap arises because more than twice as much money, two hundred and twenty million against one hundred, was at work during the bad year.
- The timing of Kestrel's deposit, not Larch's skill, cut the dollar weighted return.

5 Which return, and when?

- Time-weighted: what the manager did with the assets
- So use it to compare and judge managers
- Dollar-weighted: the investor's own experience
- Example: a limited partner in a private fund
- Trap: judging a manager by DWR when clients timed flows

IN THE LESSON

- So which is right?
- The time weighted return measures what the manager did with the assets it was given, untouched by deposits and withdrawals.
- That makes it the right basis for comparing and judging managers.
- The dollar weighted return measures the investor's own experience, including the effect of the investor's own timing.
- It suits a limited partner in a private fund, where contribution timing is part of the result.
- The exam trap is judging a manager by the dollar weighted return when the clients, not the manager, chose when money came in and out.
- Larch's minus point six seven per cent is Kestrel's result, not Larch's.

6 Sharpe's and Treynor's measures

Sharpe = (Rp − Rf) ÷ σp
Treynor = (Rp − Rf) ÷ βp

Excess return, Rp − Rf	Return above the risk-free rate
Total risk, σp	Standard deviation of return
Systematic risk, βp	Beta: moves with the market

IN THE LESSON

- Step two adjusts return for risk.
- Write R p for the portfolio return, R f for the risk free rate and R m for the market return.
- Sharpe's measure is the excess return, R p minus R f, divided by sigma p, the portfolio's standard deviation.
- It is excess return per unit of total risk.
- Treynor's measure divides the same excess return by beta p, the portfolio's beta, its sensitivity to the market.
- It is excess return per unit of systematic risk, the risk that diversification cannot remove.
- The two measures differ only in the risk they divide by, and that difference decides when each one fits.

7 Jensen's alpha and the information ratio

Alpha = Rp − [Rf + βp × (Rm − Rf)]
IR = Alpha ÷ σ(e)

CAPM required return	Rf + βp × (Rm − Rf)
Market risk premium	Rm − Rf
Jensen's alpha	Return above what CAPM requires
Residual risk, σ(e)	Risk diversification removes
Information ratio	Or active return ÷ tracking error

IN THE LESSON

- Jensen's measure, usually called Jensen's alpha, is the average return above what the capital asset pricing model, C A P M, requires.
- The required return is the risk free rate plus beta times the market risk premium, which is R m minus R f.
- Alpha is the portfolio return minus that requirement.
- Like Treynor's measure, it charges only for systematic risk.
- The information ratio divides alpha by sigma of e, the standard deviation of the residual return, which is the part of the return the market does not explain.
- That is non systematic risk, which diversification could remove.
- So the information ratio rewards alpha per unit of diversifiable risk taken to earn it.
- Some texts define it as active return over tracking error against a benchmark: same idea, different yardstick.

8 Example 3: Larch's four measures

Larch (16-year averages): Rp 9%, σp 15%, βp 0.6, σ(e) 9%. Rf 3%, Rm 10%

Sharpe	(9 − 3) ÷ 15	**0.40**
Treynor	(9 − 3) ÷ 0.6	**10.0%**
CAPM required return	3 + 0.6 × (10 − 3)	**7.2%**
Jensen's alpha	9 − 7.2	**1.8%**
Information ratio	1.8 ÷ 9	**0.20**

IN THE LESSON

- Let us measure Larch over its sixteen years with Kestrel.
- Its average return is nine per cent, its standard deviation fifteen per cent, its beta point six, and its residual standard deviation nine per cent.
- The risk free rate averaged three per cent and the market ten.
- Sharpe is six over fifteen, point four zero.
- Treynor is six over point six: ten per cent.
- The C A P M required return is three plus point six times seven, which is seven point two per cent.
- So Jensen's alpha is nine minus seven point two, one point eight per cent.
- The information ratio is one point eight over nine, point two zero.
- Larch earned less than the market, yet has a positive alpha, because its beta is low.

9 Example 4: Birch, and the market

Birch (16-year averages): Rp 14%, σp 25%, βp 1.2, σ(e) 7%. Market: σm 20%, β 1

Sharpe	(14 − 3) ÷ 25	**0.44**
Treynor	(14 − 3) ÷ 1.2	**9.17%**
Jensen's alpha	14 − (3 + 1.2 × 7)	**2.6%**
Information ratio	2.6 ÷ 7	**0.37**
Market: Sharpe, Treynor	7 ÷ 20 and 7 ÷ 1	**0.35, 7.0%**

IN THE LESSON

- Now Birch over the same sixteen years: average return fourteen per cent, standard deviation twenty five, beta one point two, and residual standard deviation seven.
- The market's standard deviation was twenty per cent, and its beta is one.
- Birch's Sharpe is eleven over twenty five, point four four.
- Its Treynor is eleven over one point two, about nine point one seven per cent.
- The required return is three plus one point two times seven, eleven point four, so alpha is fourteen minus eleven point four, two point six per cent.
- The information ratio is two point six over seven, about point three seven.
- The market's Sharpe is seven over twenty, point three five, and its Treynor seven over one, seven per cent.

10 Four measures, two winners

Both beat the market on a risk-adjusted basis (Sharpe 0.35, Treynor 7%)

Measure	Larch	Birch	Higher
Sharpe	0.40	0.44	Birch
Treynor	10.0%	9.17%	Larch
Jensen's alpha	1.8%	2.6%	Birch
Information ratio	0.20	0.37	Birch

IN THE LESSON

- Put the two side by side.
- On a risk adjusted basis both beat the market, whose Sharpe is point three five and Treynor seven per cent, but which is better depends on how the fund is held.
- On Sharpe, Larch has point four zero and Birch point four four, so Birch is higher.
- On Treynor, Larch has ten per cent and Birch nine point one seven, so Larch is higher, thanks to its low beta.
- On Jensen's alpha, Larch has one point eight per cent and Birch two point six, so Birch is higher.
- On the information ratio, Larch has point two zero and Birch point three seven, so Birch is higher again.
- The ranking depends on the measure.

11 Which measure fits which investor

- Sharpe: the investor's whole risky wealth is in this one
- Treynor or alpha: one part of a well-diversified total
- Information ratio: alpha per unit of diversifiable risk
- An only holding: rank on Sharpe, so Birch
- Trap: Treynor or alpha for an investor's only holding

IN THE LESSON

- When does each fit?
- Sharpe's measure suits an investor whose whole risky wealth sits in this one portfolio, because total risk is then what the investor bears.
- Treynor's measure and alpha suit a portfolio that is one component of a well diversified total portfolio, where residual risk is largely diversified away and only beta matters.
- The information ratio measures alpha per unit of the diversifiable risk taken to earn it.
- So, if one of these funds were an investor's only holding, rank on Sharpe and choose Birch.
- The common error is ranking an only holding by Treynor or alpha.

12 M squared: Sharpe in percentage points

Weight in P = σm ÷ σp (rest in T-bills, or borrowed)
M^2 = return of adjusted P − Rm

Weight in P	Brings risk to the market's σm
Adjusted portfolio	P mixed with T-bills or borrowing
M squared	Sharpe ranking in % points

IN THE LESSON

- Modigliani squared, or M squared, puts the Sharpe ranking into return units.
- Mix the portfolio with T bills, or borrow at the risk free rate, until its standard deviation equals the market's.
- The weight in the portfolio is sigma m over sigma p, and the rest sits in T bills, or is borrowed if the weight is above one.
- That gives the adjusted portfolio, which has exactly the market's risk.
- M squared is the adjusted portfolio's return minus the market return.
- Because both now carry the same risk, the difference is a fair comparison, and it ranks portfolios exactly as the Sharpe ratio does, but in percentage points.

13 Example 5: M squared for Larch and Birch

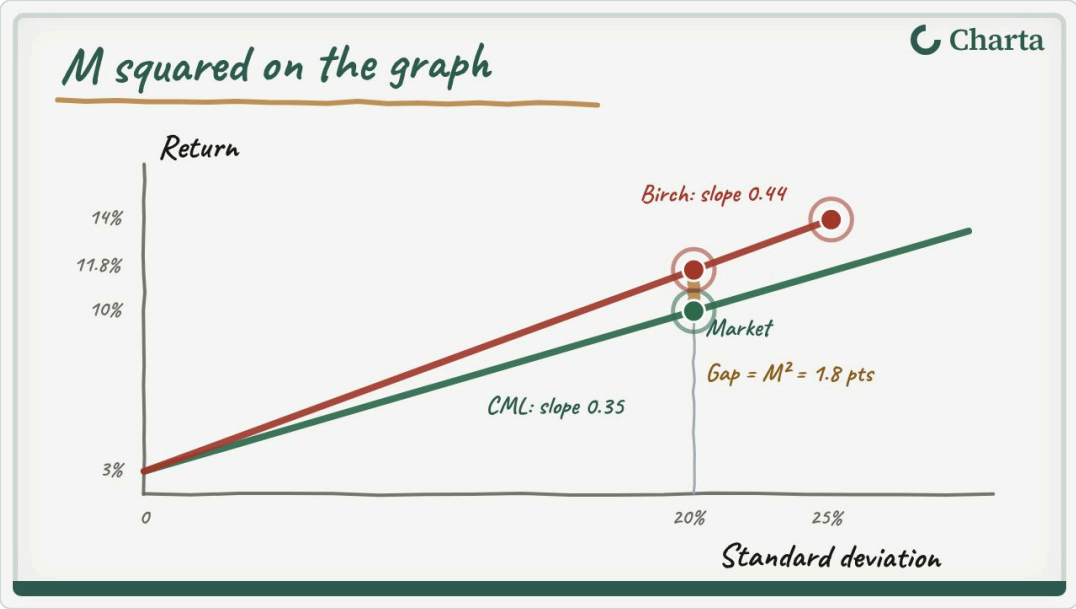
σm 20%, Rm 10%, Rf 3%. Larch: σp 15%, Rp 9%. Birch: σp 25%, Rp 14%

Larch: weight	20 ÷ 15	**1.33**
Larch: adjusted return	1.333 × 9 − 0.333 × 3	**11.0%**
Birch: weight	20 ÷ 25	**0.80**
Birch: adjusted return	0.8 × 14 + 0.2 × 3	**11.8%**
M²: Larch, Birch	11.0 − 10, 11.8 − 10	**1.0, 1.8 pts**

IN THE LESSON

- Here is M squared for both managers.
- The market has a standard deviation of twenty per cent and a return of ten.
- Larch's standard deviation is fifteen, below the market's, so its weight is twenty over fifteen, one and a third.
- The adjusted return is four thirds of nine, which is twelve, minus a third of three, which is one: eleven per cent.
- Birch is riskier, so its weight is twenty over twenty five, point eight, with point two in T bills.
- Its adjusted return is eleven point two plus point six, eleven point eight per cent.
- Subtract the market's ten: M squared is one point zero percentage point for Larch and one point eight for Birch.
- Birch is ahead, just as on Sharpe.

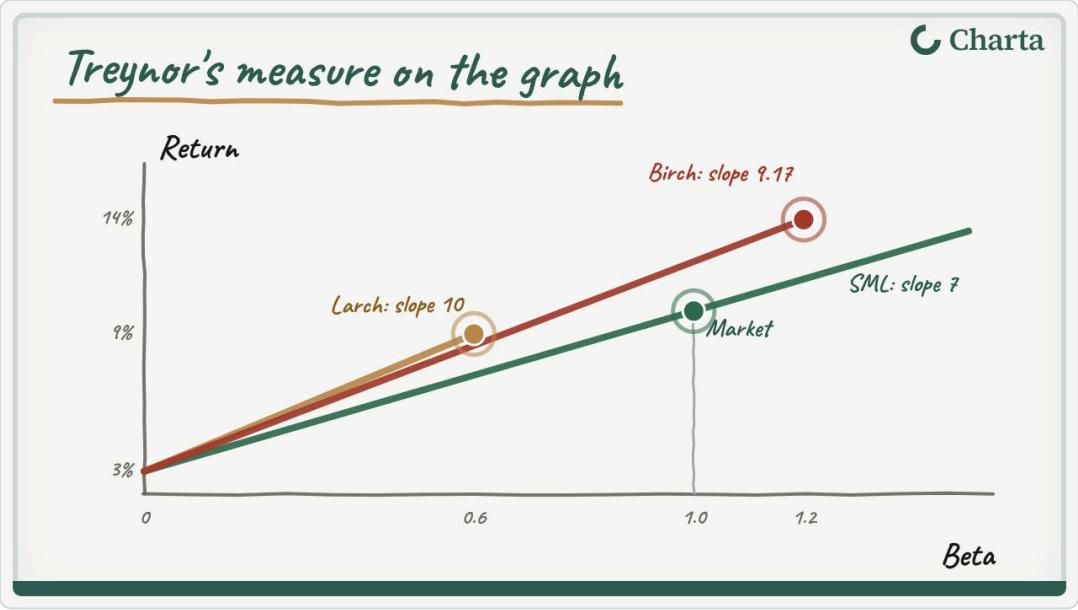
14 M squared on the graph



IN THE LESSON

- Here is M squared on a graph.
- Standard deviation runs along the bottom and return up the side.
- The green line is the capital market line, C M L.
- It runs from the risk free rate of three per cent through the market, at twenty per cent risk and ten per cent return, and its slope is the market's Sharpe ratio, point three five.
- The red line runs from the risk free rate through Birch, at twenty five per cent risk and fourteen per cent return, and its slope is Birch's Sharpe ratio, point four four.
- At the market's risk of twenty per cent, the red line reaches eleven point eight.
- M squared is the vertical gap there between the two lines: one point eight percentage points.

15 Treynor's measure on the graph



IN THE LESSON

- Treynor's measure has its own graph.
- Beta runs along the bottom.
- The green line is the security market line, S M L.
- It starts at the risk free rate of three per cent, its slope is R m minus R f, seven per cent, and it passes the market at a beta of one.
- Treynor's measure is the slope from the risk free rate to the portfolio.
- The gold line runs to Larch, at a beta of point six and a return of nine per cent, so its slope is ten.
- The red line runs to Birch, at a beta of one point two and fourteen per cent, so its slope is nine point one seven.
- Both lines are steeper than the S M L, so both portfolios plot above it, and Larch's is steeper.

16 Skill or luck? The t-statistic

t = estimate ÷ SE, SE = σ ÷ √n
n needed = (t × σ ÷ estimate)^2

Estimate	Alpha or average excess return
Standard error	Noise in the estimate
t-statistic	Above about 2: reject zero
Observations, n	Grows fast as noise rises

IN THE LESSON

- Step three asks whether a result is skill or luck.
- An alpha or an average excess return is a sample figure, measured with noise, so we test whether it differs from zero.
- The t statistic is the estimate divided by its standard error.
- For a mean of n returns with standard deviation sigma, the standard error is sigma over the square root of n.
- For a regression alpha, the standard error comes from the regression output.
- With enough observations, an absolute t above about two rejects zero, and one point nine six is the exact five per cent two sided level.
- Rearranging gives the number of observations needed: t times sigma over the estimate, all squared.
- It grows very fast as noise rises relative to signal.

17 Example 6: are the alphas significant?

16 years of yearly residual returns. Alpha is their mean: SE = σ(e) ÷ √16

Larch: standard error	9 ÷ 4	**2.25%**
Larch: t-statistic	1.8 ÷ 2.25	**0.80**
Birch: standard error	7 ÷ 4	**1.75%**
Birch: t-statistic	2.6 ÷ 1.75	**1.49**
Significant at about 2?	neither clears it	**No**

IN THE LESSON

- Is either alpha real?
- Kestrel has sixteen years of yearly residual returns for each manager.
- We treat each alpha as their mean, so the standard error is the residual standard deviation over the square root of sixteen, which is four.
- Larch's standard error is nine over four, two point two five per cent, and its t statistic is one point eight over two point two five, point eight zero.
- Birch's standard error is seven over four, one point seven five per cent, and its t statistic is two point six over one point seven five, about one point four nine.
- Is either significant at about two?
- Neither clears the bar.
- After sixteen years, Kestrel still cannot reject the idea that both alphas are luck.

18 Example 7: how many years of data?

Years needed for t = 2: n = (2 × σ(e) ÷ alpha)^2

Birch: years needed	(2 × 7 ÷ 2.6)^2	**29 years**
Larch: years needed	(2 × 9 ÷ 1.8)^2	**100 years**
Ranking by size of alpha	2.6 beats 1.8	**can mislead**
Ranking by t-statistic	1.49 against 0.80	**the right test**

IN THE LESSON

- The premise gives the rule: set t to two and solve for the number of years, two times the residual standard deviation over the alpha, all squared.
- For Birch, two times seven over two point six is about five point three eight, and squared that is about twenty nine years.
- For Larch, two times nine over one point eight is ten, and squared that is one hundred years.
- Small alphas in large noise need decades of data.
- Here the two rankings agree, Birch first, but only the t statistic, one point four nine against point eight zero, tells us how much to trust the size of alpha.
- A larger alpha with larger residual noise could have the lower t, so rank by t, not by alpha.

19 Style analysis

- Regress fund returns on asset-class or style index returns
- Weights ≥ 0, summing to 1; no unconstrained intercept
- Fitted weights = style exposures; mix = custom benchmark
- R^2: share of return variation the style mix explains
- Residual: credited to security selection and timing

IN THE LESSON

- Next, what does a manager really hold?
- Style analysis, developed by William Sharpe, infers a fund's effective asset mix from its returns alone.
- Regress the fund's returns on the returns of a set of asset class or style indices.
- The weights are constrained to be non negative and to sum to one, which mirrors a long only portfolio, with no unconstrained intercept.
- The fitted weights are the fund's style exposures, and the style weighted return of the indices is a customised benchmark.
- R squared shows how much of the fund's return variation the style mix explains.
- The residual, the part left over, is credited to security selection and any timing.

20 Example 8: Birch's style

Birch's monthly returns, regressed on three indices. R^2 = 0.90

Large-cap growth	60% × 16%	**9.6%**
Small-cap	25% × 12%	**3.0%**
Bonds	15% × 6%	**0.9%**
Style benchmark	60% + 25% + 15% = 100%	**13.5%**
Selection and timing	14.0 − 13.5	**0.5%**

IN THE LESSON

- Kestrel runs style analysis on Birch, regressing its monthly returns on three indices.
- The fit has an R squared of point nine, so the style mix explains ninety per cent of Birch's return variation.
- The fitted weights are sixty per cent large cap growth, whose index averaged sixteen per cent a year, contributing nine point six.
- Next, twenty five per cent small cap, at twelve per cent, contributing three.
- Then fifteen per cent bonds, at six per cent, contributing point nine.
- The weights sum to one hundred per cent, and this style benchmark returns thirteen point five per cent a year.
- Birch earned fourteen, so only point five of a percentage point is left for selection and timing.

21 What style analysis tells Kestrel

- Birch beats the market by 4.0 pts, its style by 0.5
- Rolling windows reveal style drift
- Needs only returns, not holdings
- Weights are estimates: exposures, not positions

IN THE LESSON

- What does this tell Kestrel?
- Against the market Birch earned four points more, but against a benchmark that matches its style only half a point remains.
- Much of the apparent skill was a tilt toward growth stocks, which a cheap index fund could have provided.
- Run over rolling windows the regression also reveals style drift, meaning a mix that moves away from the style the manager was hired for.
- Style analysis needs only returns, not holdings.
- But the weights are estimates that carry error, so treat them as exposures, not as reported positions.

22 Why active performance is hard to judge

- Alpha of 1 to 2 points a year, volatility far larger
- SE falls only with √n: decades of data needed
- Risk and holdings change, so no one beta fits
- Wrong benchmark or missing style factors fake alpha
- Thousands of managers: some win years running by luck

IN THE LESSON

- Why is it so hard to judge an active manager?
- First, the signal is small.
- Abnormal returns are often one or two percentage points a year, while annual return volatility is far larger.
- Second, the standard error falls only with the square root of the number of observations, so decades of data are needed.
- Third, a manager changes the portfolio's risk and holdings over time, so a beta or volatility estimated over a long window describes no single period, and a shorter window makes the estimate noisier.
- Fourth, a wrong benchmark, or omitted style factors, can create apparent alpha, as Birch's growth tilt did.
- Fifth, with thousands of managers, some will win several years running by luck alone.

23 Example 9: lucky streaks

1,000 managers, no skill: 50% chance to beat the market each year, years independent

Beat it in year 1	1,000 × 0.5	**500**
Beat it 3 years running	1,000 × 0.5^3	**125**
Beat it 5 years running	1,000 × 0.5^5	**about 31**
Proof of skill?	needs a significance test	**No**

IN THE LESSON

- Here is luck at work.
- Imagine one thousand managers with no skill at all, each with a fifty per cent chance of beating the market in any year, and the years independent of each other.
- After one year, one thousand times a half, about five hundred have beaten it.
- Three years running, one thousand times a half cubed, about one hundred and twenty five.
- Five years running, one thousand times a half to the fifth power, about thirty one.
- None of them has any skill.
- Is a record like that proof of skill?
- No.
- A good record needs a significance test, and the more managers there are, the more lucky streaks appear.

24 How reported performance is gamed

- Window dressing: hold recent winners at the report date
- Stale or smoothed prices: low volatility and correlation
- Selling OTM options: steady gains, rare crashes
- Survivorship and backfill: databases keep the winners
- Measures assume normal, honestly priced returns

IN THE LESSON

- Window dressing means trading just before the reporting date, so that the portfolio shows recent winners.
- Valuing illiquid assets at stale or smoothed prices understates reported volatility and correlation with the market, which inflates Sharpe and Treynor figures.
- Selling out of the money options earns a steady premium with rare large losses, so a short record can show a high Sharpe ratio before the crash arrives.
- Survivorship bias and backfilling mean databases hold only the funds that survived or chose to report, which overstates average returns.
- Sharpe, alpha and beta assume roughly normal, honestly priced returns, so they can be raised without adding any value.

25 Example 10: the option seller's Sharpe ratio

Osprey sells out-of-the-money puts. Rf 3%. Returns 10, 9, 11, 10%, then −30%

First 4 years: mean	(10 + 9 + 11 + 10) ÷ 4	**10%**
First 4 years: SD	squares sum 2; √(2 ÷ 3)	**0.816%**
First 4 years: Sharpe	(10 − 3) ÷ 0.8165	**8.57**
All 5 years: mean 2%, SD	squares sum 1282; √(1282 ÷ 4)	**17.9%**
All 5 years: Sharpe	(2 − 3) ÷ 17.9	**−0.06**

IN THE LESSON

- A fund we will call Osprey sells out of the money put options on the market and collects the premium.
- The risk free rate is three per cent.
- Osprey returns ten, nine, eleven and ten per cent in the first four years, then minus thirty in the fifth.
- Over the first four years the mean is forty over four, ten per cent.
- The deviations are zero, minus one, plus one and zero, so the squares sum to two.
- Divide by three, one fewer than the number of years: the standard deviation is the square root of two thirds, about point eight one six per cent.
- So the Sharpe ratio is seven over that, a dazzling eight point five seven.
- Add the crash year, and the mean falls to two per cent.
- The squares of the new deviations sum to one thousand two hundred and eighty two.
- Divided by four, then rooted, the standard deviation jumps to about seventeen point nine.
- Sharpe becomes minus one over seventeen point nine, about minus point zero six.

26 Example 11: survivorship bias

10 funds launch: 7 survive, averaging 9% a year; 3 close after averaging −8%

Database average	survivors only	**9.0%**
True average	(7 × 9 + 3 × (−8)) ÷ 10	**3.9%**
Overstatement	9.0 − 3.9	**5.1 pts**

Backfill adds to it. Treat smooth, high-Sharpe records with suspicion

IN THE LESSON

- Now survivorship bias.
- Ten funds launch.
- Seven survive, averaging nine per cent a year.
- Three close after averaging minus eight per cent, and their records drop out of the database.
- The database average, using survivors only, is nine per cent.
- The true average of all ten is seven times nine, sixty three, plus three times minus eight, minus twenty four, which makes thirty nine, over ten: three point nine per cent.
- The database overstates the average by five point one percentage points.
- Backfilled records add to the bias.
- The lesson of this slide and the last: treat unusually smooth, high Sharpe records with suspicion.

27 Testing for market timing

Rp − Rf = a + b(Rm − Rf) + c(Rm − Rf)^2 + e
Rp − Rf = a + b1(Rm − Rf) + b2 × D × (Rm − Rf) + e

Squared term, c	Positive: consistent with timing
Dummy, D	1 if market beats T-bills, else 0
Up-market extra beta, b2	Positive: higher beta when up
Intercept, a	Does not measure timing

IN THE LESSON

- Now market timing, which means shifting between the market and the risk free asset according to forecasts.
- A good timer holds a higher beta when the market is about to rise.
- The Treynor Mazuy regression, on the first line, adds a squared market term to the usual market model.
- If beta rises in up markets, the line curves upward, so a positive, statistically significant c is consistent with timing skill.
- The Henriksson Merton regression, on the second line, uses a dummy variable D, equal to one when the market beats T bills and zero otherwise.
- Its extra term, b two times D times the market excess return, lets beta be higher in up markets, so a positive, significant b two is consistent with timing ability.
- Beware: the intercept ay does not measure timing.

28 Example 12: testing Wren for timing

Wren, a market timer: b1 = 0.5, b2 = 0.4, SE of b2 = 0.16, a = 0

Down-market beta	b1	**0.5**
Up-market beta	b1 + b2 = 0.5 + 0.4	**0.9**
Market excess +8%	0.9 × 8%	**+7.2%**
Market excess −6%	0.5 × (−6%)	**−3.0%**
t-statistic of b2	0.4 ÷ 0.16	**2.5**

IN THE LESSON

- Kestrel also tests Wren, a manager hired to time the market.
- The Henriksson Merton regression gives b one of point five, b two of point four, a standard error on b two of point one six, and an intercept of zero.
- So Wren's down market beta is b one, point five.
- In up markets the dummy switches on, and the beta is point five plus point four, point nine.
- If the market's excess return is plus eight per cent, Wren's predicted excess return is point nine times eight, seven point two per cent.
- If it is minus six, Wren loses only point five times six, three per cent.
- The t statistic on b two is point four over point one six, two point five.
- That is above two, so the result is consistent with timing ability, though noise or misspecification can also produce it.

29 Example 13: what perfect timing is worth

Rf 3%, market σ 20%, one year. Price per $1 with Black-Scholes

Perfect timer earns	max(Rm, Rf)	**T-bills + call**
Strike, present value	1.03 ÷ 1.03	**$1.00**
d1 and d2	(0 ± 0.20^2 ÷ 2) ÷ 0.20	**±0.10**
Call value	N(0.10) − N(−0.10)	**$0.0797**
Value of perfect timing	per dollar, per year	**about 8%**

IN THE LESSON

- Now value perfect timing, with a risk free rate of three per cent, market volatility of twenty per cent, a one year horizon, and Black Scholes to price it per dollar.
- A perfect timer earns the larger of R m and R f.
- That is the same as T bills plus a call option on the market, struck at the risk free payoff.
- The strike of one point zero three, discounted at three per cent, has a present value of exactly one dollar, equal to the dollar invested, so the log term is zero.
- So d one is point two squared over two, divided by point two, which is point one, and d two is minus point one.
- The call is N of point one minus N of minus point one, about point zero seven nine seven.
- Perfect timing is worth about eight per cent a year.

30 Performance attribution

Allocation = Σ (wP − wB) × rB
Sector = wP(class) × Σ (wP − wB) × rB(sector)
Security = wP × (rP − rB)

Asset allocation	Class weights against benchmark
Sector selection	Sector weights inside a class
Security selection	wP: class × sector, 70%×50%=35%

IN THE LESSON

- Finally, where did the result come from?
- Performance attribution explains why a portfolio's return differs from its benchmark's by stepping through the investment decisions.
- Asset allocation compares the portfolio's weight in each class with the benchmark's, and values each difference at the benchmark's return for that class: the sum of w P minus w B, times r B.
- Within a class, sector selection is the sector weight difference times the benchmark sector return, scaled by the portfolio's weight in that class.
- Security selection is measured within each sector, using portfolio weights: w P times r P minus r B.
- In the security line, w P is the sector's weight in the whole portfolio: class weight times sector weight, for example seventy per cent times fifty per cent, thirty five per cent.
- The three effects add exactly to the total excess return.

31 Example 14: excess return and allocation

Equity: P 70% earning 12%, B 60% earning 10%.
Bonds: P 30% at 3.5%, B 40% at 4%

Portfolio return	0.7 × 12 + 0.3 × 3.5	**9.45%**
Benchmark return	0.6 × 10 + 0.4 × 4	**7.60%**
Excess to explain	9.45 − 7.60	**1.85%**
Allocation: equity	(70% − 60%) × 10%	**+1.00%**
Allocation: bonds	(30% − 40%) × 4%	**−0.40%**

Allocation effect = 1.00 − 0.40 = +0.60%

IN THE LESSON

- Now we split the results of Kestrel's own total portfolio, not of one manager.
- Kestrel holds seventy per cent equity, which earned twelve per cent, against a benchmark weight of sixty, where the equity index earned ten.
- It holds thirty per cent bonds, which earned three point five, against a benchmark weight of forty, where the bond index earned four.
- Kestrel's return is point seven times twelve plus point three times three point five: nine point four five per cent.
- The benchmark's is point six times ten plus point four times four: seven point six.
- So there are one point eight five points to explain.
- Allocation to equity adds ten points of extra weight times ten per cent, plus one point zero.
- Allocation to bonds is minus ten points of weight times four, minus point four.
- The allocation effect is plus point six.

32 Example 15: sector selection inside equity

Inside equity: tech P 50%, B 40%, index 16%. Other P 50%, B 60%, index 6%

Check: benchmark equity	40% × 16 + 60% × 6	**10.0%**
Sector: tech	70% × (50% − 40%) × 16%	**+1.12%**
Sector: other	70% × (50% − 60%) × 6%	**−0.42%**
Sector selection effect	1.12 − 0.42	**+0.70%**

IN THE LESSON

- Now look inside equity, which has two sectors.
- Kestrel holds fifty per cent of its equity in technology, against a benchmark weight of forty, and the benchmark technology index earned sixteen per cent.
- It holds fifty per cent in other stocks, against sixty, and that index earned six.
- As a check, forty per cent of sixteen plus sixty per cent of six gives ten per cent.
- Sector selection for technology is the class weight of seventy per cent, times the ten point overweight, times sixteen: plus one point one two.
- For other stocks it is seventy per cent, times the ten point underweight, times six: minus point four two.
- The sector selection effect is plus point seven zero.

33 Example 16: security selection and the total

Kestrel's returns: tech 17%, other 7%. Bonds 3.5% against the index's 4%

Security: tech	70% × 50% × (17 − 16)	**+0.35%**
Security: other	70% × 50% × (7 − 6)	**+0.35%**
Security: bonds	30% × (3.5 − 4)	**−0.15%**
Security selection effect	0.35 + 0.35 − 0.15	**+0.55%**
Total of three effects	0.60 + 0.70 + 0.55	**1.85%**

IN THE LESSON

- Last, security selection, using Kestrel's own weights.
- Its technology stocks earned seventeen per cent against the index's sixteen, its other stocks seven against six, and its bonds three point five against four.
- In technology the weight is seventy per cent times fifty per cent, times one point: plus point three five.
- Other stocks give the same, plus point three five.
- Bonds give thirty per cent times minus half a point: minus point one five.
- The security selection effect is plus point five five.
- Adding the three effects, point six plus point seven plus point five five gives one point eight five, exactly the excess return.
- The effects can differ in sign: underweighting a falling class gives a positive allocation effect even when selection inside it is negative.
- The traps: using benchmark weights in the selection term, or valuing weight differences at portfolio returns.

34 Try three quick questions

1. A fund earns +10%, a client adds money, then it earns −5%. Time-weighted return?
   Answer: 1.10 × 0.95 − 1 = 4.5%.
2. Rp 12%, σp 25%, Rf 2%, Rm 9%, σm 20%. M squared?
   Answer: Weight 0.8: 0.8 × 12 + 0.2 × 2 = 10%. M^2 = 10 − 9 = 1 pt.
3. Alpha 3%, residual σ 12%, 9 years of data. Significant?
   Answer: SE = 12 ÷ 3 = 4. t = 0.75, below 2: not significant.

IN THE LESSON

- Now try three quick questions.
- Pause, think, then listen.
- Question one.
- A fund earns ten per cent, a client then adds money, and the fund earns minus five per cent.
- What is the time weighted return?
- One point one times point nine five, minus one: four point five per cent.
- Question two.
- A portfolio returns twelve per cent with a standard deviation of twenty five.
- The risk free rate is two per cent, and the market returns nine with a standard deviation of twenty.
- What is M squared?
- The weight is point eight, so the adjusted return is nine point six plus point four, ten per cent, and M squared is one percentage point.
- Question three.
- An alpha of three per cent has a residual standard deviation of twelve over nine years.
- Is it significant?
- The standard error is twelve over three, four, so t is point seven five, below two.
- It is not significant.