CHARTA LEARNING · LESSON NOTES

Multifactor models in practice, and a first look at value at risk

Portfolio Management

1 Four jobs for a factor model

- Return attribution: where did the return come from?
- Risk attribution: where does the risk come from?
- Portfolio construction: build the bets you want
- Strategic allocation: set the long run mix
- Always measure against the benchmark

IN THE LESSON

- Level one gave you the idea of a factor model: a return is explained by sensitivities to a few common drivers, plus a leftover part that belongs to the one security.
- Now we put it to work.
- A manager and a client use it in four ways.
- To explain where past return came from, which is return attribution.
- To explain where the risk comes from, which is risk attribution.
- To build a portfolio on purpose, which is portfolio construction.
- And to set the long run mix of a whole plan, which is strategic allocation.
- Everything depends on one habit: always compare the portfolio with its benchmark, because the benchmark's own exposures are not the manager's doing.

2 Splitting active return

Active return = Σ (active sensitivity × factor return) + selection

Active return	Portfolio minus benchmark
Active sensitivity	Portfolio β − benchmark β
Factor return	What the factor earned
Selection return	The leftover, stock picking

IN THE LESSON

- Here is the rule for return attribution.
- Active return is the portfolio's return minus the benchmark's.
- A factor model splits it in two.
- The first part is the factor return.
- For each factor, take the active sensitivity, which is the portfolio's sensitivity minus the benchmark's, and multiply it by what that factor returned.
- Add these up over all the factors.
- The second part is selection return, and it is simply whatever is left over once the factor part is taken out.
- It is the effect of choosing particular securities within each exposure.
- So factor tilts are the manager's style bets, and selection is the manager's stock picking.

3 Example 1: the active sensitivities

Active return 1.2%. Factor returns: 6%, −2%, 3%

Market: 1.10 − 1.00 = 0.10	0.10 × 6%	**+0.6%**
Size: 0.30 − 0.00 = 0.30	0.30 × −2%	**−0.6%**
Value: 0.20 − 0.10 = 0.10	0.10 × 3%	**+0.3%**

IN THE LESSON

- Meet the Heron Capital fund.
- Last year it beat its benchmark by one point two per cent.
- A three factor model is used: market, size and value.
- Heron's sensitivities are one point one, zero point three and zero point two.
- The benchmark's are one point zero, zero and zero point one.
- Subtract, factor by factor, and the active sensitivities are zero point one, zero point three and zero point one.
- The factor returns last year were six per cent for the market, minus two per cent for size, because small companies lagged, and plus three per cent for value.
- Now multiply each active sensitivity by its factor return.

4 Example 2: factor part, selection, and the trap

Total active return 1.2%

Factor part	0.6 − 0.6 + 0.3	**0.3%**
Selection part	1.2% − 0.3%	**0.9%**
Trap: own sensitivities	1.1×6 + 0.3×−2 + 0.2×3	**6.6%**
Trap: selection then	1.2% − 6.6%	**−5.4%**

IN THE LESSON

- Add the three contributions: zero point six, minus zero point six, plus zero point three, for a factor return of just zero point three per cent.
- The active return was one point two, so the selection return is one point two minus zero point three, which is zero point nine per cent.
- Most of the outperformance came from picking stocks, not from the style bets.
- Now the classic trap.
- Suppose you multiply the fund's own sensitivities by the factor returns instead.
- One point one times six, plus zero point three times minus two, plus zero point two times three, makes six point six.
- That credits the manager with the benchmark's exposures, and leaves a silly selection figure of minus five point four.

5 Splitting active risk

Active risk² = Active factor risk² + Active specific risk²

Active risk	Tracking risk, squared is variance
Active factor risk	From sensitivity gaps
Active specific risk	From stock weight gaps

IN THE LESSON

- Risk attribution works the same way, but with one big difference: it has to be done with variances, not standard deviations.
- Active risk, also called tracking risk, is the standard deviation of the portfolio's return minus the benchmark's return.
- Square it and you get active variance.
- In a factor model, active variance equals active factor variance plus active specific variance.
- Active factor risk comes from the gaps between the portfolio's and the benchmark's sensitivities.
- Active specific risk comes from the active weights on individual securities, times their residual risks.
- Each factor's share of the factor part shows exactly where the bets lie.

6 Example 3: Heron's active risk

Active factor risk 2.4%, active specific risk 3.2%

Squared factor part	2.4^2	**5.76**
Squared specific part	3.2^2	**10.24**
Active variance, then risk	√(5.76 + 10.24)	**4.0%**
Specific share	10.24 ÷ 16	**64%**
Factor share	5.76 ÷ 16	**36%**

IN THE LESSON

- Heron's risk model reports active factor risk of two point four per cent and active specific risk of three point two.
- Square each.
- Two point four squared is five point seven six.
- Three point two squared is ten point two four.
- Add them: sixteen.
- The square root of sixteen is four.
- So Heron's active risk is four per cent, not five point six, which is what you would get by adding the two numbers directly.
- And the shares are worked out on the variances: ten point two four over sixteen is sixty four per cent from stock specific bets, leaving thirty six per cent from the factors.
- That matches the return story, where selection dominated.

7 Building with a factor model

Match the benchmark

List benchmark sensitivities → Match each one → Only stock picking differs

Tilt one factor

Choose the factor → Raise its sensitivity → Neutralise the rest

IN THE LESSON

- The model is also a drawing board.
- Follow the two rows.
- In the top row, a manager wants to track the benchmark closely.
- She sets out the benchmark's factor sensitivities, builds a portfolio whose sensitivities match them, and then lets stock selection be the only difference.
- Active factor risk falls close to zero.
- In the bottom row, a manager wants one deliberate bet.
- She picks a single factor, raises the sensitivity to it, and holds every other sensitivity at the benchmark's level so that the other factors are neutralised.
- Strategic allocators use the same tool in a larger way, choosing the factor exposures they want before choosing any manager.

8 Same beta, different risks

- Equal market beta does not mean equal risk
- Recession, inflation and credit shocks hit differently
- A single beta hides these exposures
- The model shows what you really own

IN THE LESSON

- Why bother with several risk dimensions?
- Think of two cars with the same top speed.
- In a single factor world they look identical.
- But one is built for motorways and the other for muddy tracks, and you only find out when the weather changes.
- Two portfolios can share a market beta of one, and still respond very differently to a recession, to a jump in inflation or to a credit scare.
- A single beta hides that.
- A multifactor model lays it out, so the investor sees what they really own, and can pick the portfolio whose weak spots do not line up with their own.

9 How two portfolios react

Return effect (%)

X in recession	**−3.0%**
Y in recession	**−1.0%**
X, inflation up	**+0.2%**
Y, inflation up	**+1.2%**

Growth surprise −2 points; inflation surprise +2 points (illustration)

IN THE LESSON

- Here are Portfolios X and Y, which both have a market beta of one.
- The vertical axis is the effect on the portfolio's return, in per cent.
- The first two bars show a recession, with growth surprising low by two percentage points.
- X is highly sensitive to growth, with a sensitivity of one point five, so it loses three per cent.
- Y's sensitivity is only zero point five, so it loses one per cent.
- The last two bars show an inflation surprise of plus two points.
- X, with a sensitivity of just zero point one, gains about zero point two per cent.
- Y, with a sensitivity of zero point six, gains one point two.
- Same beta, quite different behaviour.

10 Example 4: which portfolio suits her?

X expects 9%, Y expects 8%; both have market beta 1.0

Recession, X	1.5 × −2%	**−3.0%**
Recession, Y	0.5 × −2%	**−1.0%**
Inflation surprise	X +0.2%, Y +1.2%	**see chart**
Expected return given up	9.0% − 8.0%	**1.0%**
Her choice	if she accepts the cost	**Y**

IN THE LESSON

- Now the investor.
- Thandi owns a cyclical manufacturing business whose profits fall in every recession.
- Portfolio X expects nine per cent a year, and portfolio Y expects eight.
- A single factor view says: same beta, take the higher expected return, so choose X.
- But look at her whole position.
- In a recession her business suffers, and X would lose three per cent on top, while Y would lose only one.
- Choosing Y costs her one point of expected return, nine minus eight.
- If she is willing to give up one per cent a year to cut her growth exposure, Y fits her better, because it has less of the risk she already carries.

11 Hedge some risks, harvest others

- Liabilities and income already carry factor risks
- Hedge unwanted exposures, even at a lower expected return
- Harvest premiums the investor is able to bear
- The market portfolio is not best for everyone

IN THE LESSON

- The same idea works for institutions.
- A pension fund whose liabilities rise with inflation already has a big inflation exposure on the other side of its balance sheet, so it may want assets that rise with inflation, and it may accept a lower expected return to get them.
- A family that owns a cyclical business may want less growth risk in its shares.
- On the other hand, an investor with no such exposures may deliberately hold more of a factor the market pays a premium for.
- So a multifactor view lets investors hedge exposures they do not want, and harvest premiums they can bear, instead of simply holding the market portfolio.

12 Mid-lesson: four common mistakes

- Own sensitivities, not active sensitivities
- Selection is the leftover that makes the total add up
- Equal beta does not mean equal risk
- The market portfolio is not best for everyone

IN THE LESSON

- Let us pause on the slips that cost marks so far.
- One: using the portfolio's own sensitivities instead of the active sensitivities, which hands the benchmark's exposures to the manager.
- Two: forgetting that the selection return is a leftover, so it must make the parts add up to the total active return.
- Three: assuming two portfolios with equal market betas have equal risk, when growth, inflation or credit exposures may differ widely.
- And four: assuming the market portfolio is the best risky portfolio for every investor, whatever else they own.
- It is a sensible starting point, but an investor's other exposures can move the best answer away from it.

13 Where active risk comes from

- Active risk = standard deviation of (portfolio − benchmark)
- Also called tracking risk
- Source one: factor tilts versus the benchmark
- Source two: active weights in single securities

IN THE LESSON

- Picture two drivers who both stray about the same distance from the road each year.
- One drifts because of a steady crosswind, the other because of random bumps.
- The distance looks the same, but the cause is very different.
- Active risk is like this.
- It measures how far the portfolio's return strays from the benchmark's, as a standard deviation of the difference, and it is also called tracking risk.
- A factor model says the straying has two sources.
- One is deliberate tilts, where the portfolio's factor sensitivities differ from the benchmark's.
- The other is stock specific, where the manager holds more or less of single securities than the index does.

14 Two managers, two kinds of active variance

Active variance (squared %)

A: factor	**9**
A: specific	**16**
B: factor	**36**
B: specific	**4**

A total 25, risk 5.0%. B total 40, risk 6.3%

IN THE LESSON

- Let us compare two managers.
- The vertical axis here is active variance, in squared percentage points, because variance is the thing that adds.
- Manager A has factor variance of nine and specific variance of sixteen.
- Add them to get twenty five, and its active risk is the square root, five per cent.
- Most of A's variance comes from stock picking.
- Manager B has a factor variance of thirty six and a specific variance of just four.
- The total is forty, and the active risk is the square root of forty, about six point three per cent.
- B looks only a little riskier, but nearly all of its risk is a factor tilt, which is quite a different product.

15 Example 5: adding variances, not deviations

Active factor risk 3%, active specific risk 4%

Active variance	$3^2 + 4^2$	**25**
Active risk	√25	**5.0%**
Factor share of variance	9 ÷ 25	**36%**
Trap: adding deviations	3 + 4	**7%**
Trap: factor share	3 ÷ 7	**43%**

IN THE LESSON

- Example five uses a different manager.
- The factor model reports active factor risk of three per cent and active specific risk of four per cent.
- Square them: nine and sixteen.
- The sum, twenty five, is the active variance, and its square root, five per cent, is the active risk.
- Factor risk is nine out of twenty five, thirty six per cent of the variance, and specific is sixty four per cent.
- The tempting mistake is to add three and four, getting seven per cent.
- That overstates the active risk by two points, and it also gives a factor share of three over seven, forty three per cent, which is wrong.
- The numbers are standard deviations, and standard deviations do not add.
- Variances do.

16 The information ratio

Information ratio = Mean active return ÷ Active risk

Information ratio	Reward per unit of active risk
Mean active return	Average of portfolio − benchmark
Active risk	Standard deviation of that gap

IN THE LESSON

- Now, how do we compare managers fairly?
- A manager with a large active return may have taken enormous risk to get it.
- The information ratio fixes that.
- It is the mean active return divided by the active risk.
- It tells you how much benchmark relative return you earned for each unit of benchmark relative risk.
- Notice the similarity with the Sharpe ratio, except that the reference point is the benchmark and not the risk free rate.
- A higher information ratio is better.
- And because it measures reward per unit of risk, it can be used to compare managers of different sizes and styles, and to judge whether skill is being rewarded by the risk taken.

17 Example 6: comparing two managers

Her: active return 1.5%, risk 5%. Rival: 2.0% and 8%

Her information ratio	1.5 ÷ 5.0	**0.30**
Rival information ratio	2.0 ÷ 8.0	**0.25**
Rival's active return	higher, at 2.0%	**but costlier**
Better reward per risk	0.30 beats 0.25	**Her**

IN THE LESSON

- Back to the manager with active risk of five per cent.
- Her mean active return over the period was one point five per cent.
- So her information ratio is one point five over five, zero point three.
- A rival manager earned a higher mean active return of two per cent, but with active risk of eight per cent.
- His information ratio is two over eight, zero point two five.
- So she wins, even though he earned more in absolute terms.
- Each unit of tracking risk bought her three tenths of a percentage point of active return, and only a quarter for him.
- Higher active return alone is not the right test.
- It must be judged against the risk.

18 Reading active risk the right way

- Specific risk falls as positions are diversified
- Factor risk stays however many names carry the tilt
- The split shows what the client is paying for
- The ratio asks if the reward was worth the risk

IN THE LESSON

- Two lessons from the split.
- First, the two sources behave differently.
- Specific risk is built from many separate stock bets.
- If the manager spreads her active weight over more names, the residual risks partly cancel, and active specific risk falls.
- Factor risk does not behave that way.
- If every stock carries the same tilt towards small companies, spreading the bets across more names leaves the tilt exactly as it was.
- Second, the split tells a client what they are paying for.
- A fund with mostly factor risk is delivering a style tilt, which could often be bought more cheaply, while a fund with mostly specific risk is selling stock picking skill.
- The information ratio then asks whether the reward justified the risk.

19 Value at risk in plain words

- A loss level with a probability and a horizon
- One day, five per cent, two million
- Loss of at least two million on about one day in twenty
- Used for limits, allocation and reporting

IN THE LESSON

- Now a different question.
- A board does not want a page of sensitivities.
- It wants one number: how bad could a bad day be?
- That is the idea behind value at risk, which people shorten to VaR.
- It is a loss figure with two labels attached: a small probability and a time horizon.
- Suppose the one day five per cent VaR is two million.
- That means that, under normal market conditions, a loss of at least two million is expected on about one day in twenty.
- Said the other way, with ninety five per cent confidence the loss on a day will be smaller than two million.
- Banks and funds use it to set limits, to share out risk between desks, and to report to boards and regulators.

20 A cut-off in the left tail

Chart: Likelihood vs Profit or loss on the day

- VaR
- Worst 5% of days
- Average day

IN THE LESSON

- Here is the picture.
- The horizontal axis is the profit or loss on the day, with losses on the left.
- The vertical axis is how likely each outcome is.
- The green curve is the familiar bell shape.
- Most days sit near the middle, where the portfolio gains or loses a little, and that is the average day.
- Now look at the left end.
- The marker is the value at risk, and the thick red line below the curve to its left stands for the worst five per cent of days.
- Notice that the tail carries on to the left of the marker.
- VaR says days in this region lose at least that much, and it says nothing about how far beyond the line a very bad day can go.

21 Parametric value at risk

VaR = z × σ × Value
and for n days: VaR × √n

Value at risk	Loss level at that chance
z value for the probability	1.65 at 5%; 2.33 at 1%
Standard deviation of return	Per day, for a daily VaR
Square root of days	Scales the horizon

IN THE LESSON

- There are several ways to estimate VaR.
- The simplest is the parametric method, which assumes returns are normally distributed.
- With a mean of zero, the VaR is a z value times the standard deviation, times the value of the portfolio.
- The z value depends on the probability.
- For five per cent it is about one point six five, and for one per cent about two point three three, which is bigger, because a rarer day is a worse day.
- To change the horizon, if daily returns are independent with a zero mean, multiply the one day figure by the square root of the number of days.
- That is the square root of time rule, which is why ten days is not ten times riskier, but about three point two times.

22 Example 7: reading a two million VaR

One-day 5% VaR of 2.0 million; mean zero; 250 trading days

Expected breaches a year	5% × 250	**12.5 days**
Implied daily σ	2.0 ÷ 1.65	**1.21 million**
Check	1.65 × 1.21	**2.0 million**

IN THE LESSON

- Let us put the two million VaR through its paces.
- The portfolio reports a one day, five per cent VaR of two million, with normally distributed returns and a mean of zero.
- First, how often should it be beaten?
- Five per cent of two hundred and fifty trading days is twelve and a half days a year.
- That is the testable promise.
- Second, what daily standard deviation is hidden inside?
- Two divided by one point six five is one point two one million, because the VaR equals z times the standard deviation.
- So the daily standard deviation is about one point two one million.
- These two answers let us build every other version of the figure.

23 Example 8: other probabilities and horizons

From Example 7: σ = 1.21 million a day

1% one-day VaR	2.33 × 1.21	**2.82 million**
5% one-day VaR	1.65 × 1.21	**2.00 million**
5% ten-day VaR	2.0 × √10	**6.32 million**
Larger when	rarer or longer	**yes**

IN THE LESSON

- Now change the labels.
- At the one per cent level, the z value is two point three three, so the one day VaR is two point three three times one point two one million, which is two point eight two million.
- It is bigger than the five per cent figure, because the one per cent cut off lies further out in the tail.
- Next, stretch the horizon to ten days at five per cent.
- Multiply the one day figure by the square root of ten, about three point one six.
- Two million times three point one six is six point three two million.
- The lesson: a VaR figure means nothing until you know the probability and the horizon behind it.

24 VaR is a threshold, not a worst case

- VaR is a minimum loss on the tail days, not a maximum
- A loss above VaR is expected on some days
- Too many breaches suggest the model understates risk
- Always state the probability and the horizon

IN THE LESSON

- The board then asks: last Tuesday we lost three and a half million.
- Does that show the model is wrong?
- Not by itself.
- VaR is a minimum loss on the worst days, so a loss above two million is exactly what the five per cent tail contains, and about twelve or thirteen such days a year are expected.
- The model comes under suspicion if the count of breaches runs well above that rate over many days.
- So remember the trap: VaR is not the maximum possible loss.
- It is the loss level that is expected to be reached or beaten with the stated small probability, and it says nothing about how bad the loss can be beyond that line.

25 Recap: what to take away

- Active return = factor part (active sensitivities) + selection
- Active variance = factor variance + specific variance
- Information ratio = mean active return ÷ active risk
- VaR: a loss level with a probability and a horizon
- Not a worst case

IN THE LESSON

- Let us pull it together.
- A multifactor model splits active return into the factor part, built from active sensitivities times factor returns, and a leftover selection part.
- It splits active variance, not active risk, into a factor part and a specific part.
- Looking at several risk dimensions shows what an investor really holds, and lets them hedge exposures they already carry and harvest others on purpose.
- The information ratio is mean active return over active risk, used to compare managers fairly.
- And value at risk is a loss level with a probability and a horizon, larger for rarer events and longer periods, scaled by the square root of time, and never the worst case.

26 Try three quick questions

1. Active sens 0.2, factor return 4%, active return 1.5%. Selection?
   Answer: 0.2 × 4 = 0.8%; selection 0.7%.
2. Active factor risk 6%, active specific risk 8%. Active risk?
   Answer: √(36 + 64) = 10%, not 14%.
3. One-day 5% VaR is 3 million. Ten-day VaR?
   Answer: 3 × √10 = about 9.5 million.

IN THE LESSON

- Now try three questions.
- Pause after each one.
- Question one.
- A fund has an active sensitivity of zero point two to a factor that returned four per cent, and its active return was one point five per cent.
- How much is selection?
- Zero point two times four is zero point eight, so selection is one point five minus zero point eight, zero point seven per cent.
- Question two.
- Active factor risk is six per cent and active specific risk is eight per cent.
- What is active risk?
- Thirty six plus sixty four is one hundred, so ten per cent, not fourteen.
- Question three.
- A one day five per cent VaR is three million.
- What is the ten day figure, with independent days and zero mean?
- Three times the square root of ten is about nine point five million, so that is the answer.