CHARTA LEARNING · LESSON NOTES

Building active equity portfolios

Portfolio Management

1 A story: the chef's two choices

- Where to start: stocks or themes?
- How to decide: rules or judgement?
- Bottom up and top down; systematic and discretionary
- Then the third choice: how big is each position?

IN THE LESSON

- Think of a chef planning a banquet.
- The chef makes two choices.
- The first is where to start.
- One chef picks the best ingredients at the market and lets the menu grow from them.
- Another fixes the menu first and then finds ingredients to fit.
- The second choice is how to decide.
- One chef follows a written recipe that anyone could repeat.
- Another cooks by taste and experience.
- An active equity manager faces the same two choices.
- Start from single stocks, which is called bottom up, or start from big themes, which is called top down.
- Decide by rules and models, which is systematic, or by judgement, which is discretionary.
- After that comes a third choice, how much of each stock to hold.

2 Four ways to build a portfolio

Two questions give four styles

Style	Decides by, starts from
Quant stock selection	Rules, stocks
Factor rotation	Rules, themes
Fundamental stock picker	Judgement, stocks
Macro or thematic	Judgement, themes

IN THE LESSON

- Cross the two questions and you get four styles.
- A quant stock selector uses rules and starts from stocks.
- It ranks thousands of companies with a model and builds from the best.
- A factor rotator also uses rules, but starts from themes such as value or momentum, and moves between them.
- A fundamental stock picker uses judgement and starts from stocks.
- The analyst visits companies and forms a conviction, and the sector and style exposures simply fall out of those choices.
- A macro or thematic manager uses judgement and starts from themes, such as a view on the economy, and then finds stocks to express it.

3 Four ways to set position sizes

From idea to position size

Method	Position rule
Equal weight	Same for all
Conviction tiers	By conviction
Limits around the benchmark	Near the index
Optimiser	Alpha less risk

IN THE LESSON

- Once the stocks are chosen, the manager must size each position.
- The simplest rule is equal weight.
- Every stock gets the same slice.
- A step up is conviction tiers.
- The strongest ideas get big weights, the middle ideas medium weights and the weakest small ones.
- A third way is to anchor on the benchmark and allow each stock to stray only so far from its index weight, with similar limits for each sector.
- The most formal way is an optimiser.
- It takes forecasts of alpha, which means expected excess return, along with risk and trading cost, and searches for the weights that maximise alpha less penalties for risk and cost.

4 Example 1: tiers and a limit

Roskin fund: 25 stocks, three conviction tiers

Equal weight, 25 stocks	100% ÷ 25	**4%**
Top tier: 5 stocks	5 × 8%	**40%**
Middle tier: 10 stocks	10 × 4%	**40%**
Bottom tier: 10 stocks	10 × 2%	**20%**
Cap: index 3% + 3 points	top-tier stock	**6% max**

IN THE LESSON

- Ines Cadmoor holds twenty five stocks.
- With equal weights each one gets one hundred divided by twenty five, which is four per cent.
- But she has stronger and weaker ideas, so she uses three tiers.
- Five top ideas at eight per cent make forty.
- Ten middle ideas at four per cent make another forty.
- Ten bottom ideas at two per cent make twenty.
- The tiers add up to the whole fund.
- Now add a limit.
- Her rule says no stock may be more than three points above its weight in the index.
- If a top tier stock has an index weight of three per cent, its cap is six per cent so the plan is cut back to six.

5 How an optimiser works

From forecasts to weights

Forecast alpha, risk and cost →
Maximise alpha less penalties →
Active weights come out

IN THE LESSON

- Now the optimiser.
- Think of it as a very fast and very literal assistant.
- The first box is its input.
- Ines gives it a forecast of alpha for every stock, an estimate of how risky each stock is and how the stocks move together, and the cost of trading.
- The second box is the instruction: choose the weights that maximise alpha, less a penalty for the risk taken, and less the trading cost.
- The third box is the answer, a list of active weights, meaning how far each holding sits above or below the index.
- But it believes every forecast it is handed, as we are about to see.

6 Example 2: the optimiser magnifies noise

Two stocks, 25% volatility each, correlation 0.9

A: alpha 2.0 against B: 1.5	2.0 − 0.9 × 1.5	**+0.65**
B: alpha 1.5 against A: 2.0	1.5 − 0.9 × 2.0	**−0.30**
Forecasts swapped, A	1.5 − 0.9 × 2.0	**−0.30**
Forecasts swapped, B	2.0 − 0.9 × 1.5	**+0.65**

A half-point forecast change flips both positions

IN THE LESSON

- Here is why an optimiser is not as objective as it looks.
- Ines has two stocks in the same industry.
- Each has twenty five per cent volatility and they move almost together, with a correlation of point nine.
- Her analysts forecast alpha of two per cent for A and one point five for B.
- When volatilities are equal, the unconstrained weight of each stock is proportional to its alpha minus point nine times the other stock's alpha.
- For A that is plus point six five, so overweight.
- For B it is minus point three, so underweight, even though its alpha is positive.
- Now swap the two forecasts, a change of half a point.
- A turns underweight and B overweight.
- The model has magnified noise.

7 Taming the optimiser

- Cap active weights near the benchmark weight
- Shrink alpha forecasts towards zero first
- Blend the model's answer with judgement
- Hybrids are normal: a model inside limits

IN THE LESSON

- So what do careful managers do?
- They fence the optimiser in.
- They cap every active weight so that no stock can sit more than a few points from its index weight.
- They shrink the alpha forecasts towards zero before optimising, which takes some of the noise out of the input.
- And they blend the machine's answer with human judgement, checking each large position against the analyst's story.
- The result is a hybrid: a model working inside limits.
- This is the point to remember for the exam.
- An optimiser is a good tool for treating risk and cost consistently across hundreds of stocks.
- It is not a source of objective truth.

8 A story: two cooks, one recipe

- Cook one: swaps a little of everything
- Cook two: swaps one thing that matters
- Active Share: how different are the holdings?
- Active risk: how far do the returns wander?

IN THE LESSON

- Imagine two cooks given the same recipe.
- The first cook swaps a little of every ingredient: a different herb here, a different cut there, many small changes.
- The second cook leaves every ingredient alone except one, but that one is the chilli, and it changes the whole dish.
- Which cook changed the recipe more?
- It depends on the question.
- If you ask how many ingredients are different, it is the first cook.
- If you ask how far the taste moved, it is the second.
- Active Share counts how different the holdings are.
- Active risk, also called tracking error, measures how far the returns wander from the index.
- We need both.

9 Two measures of being active

Active Share = ½ × Σ |wp − wb|, from 0 to 1
Active risk = standard deviation of excess return

Portfolio weight	wp: what the fund holds
Benchmark weight	wb: what the index holds
Absolute gap	Size only, ignore the sign
Active risk	Spread of excess returns

IN THE LESSON

- Here are the two measures.
- Active Share is one half times the sum of the absolute gaps between the fund's weight and the index weight in each stock.
- Absolute means we ignore the sign, so an overweight and an underweight both count as a gap.
- We halve the total because every overweight has a matching underweight, and we do not want to count the same difference twice.
- The answer runs from zero, an exact copy of the index, to one, a fund that shares no stock with the index.
- Active risk is different.
- It is the standard deviation of the fund's excess return over the index, measured over time.
- It depends on the size of the bets and on how they move together.

10 Example 3: Active Share, stock by stock

Index 30, 25, 20, 15, 10, 0; fund 15, 25, 35, 0, 15, 10 (%)

Overweights: C, E, F	+15, +5, +10	**30**
Underweights: A, D	−15, −15	**30**
Sum of absolute gaps	30 + 30	**60**
Active Share	½ × 60	**30%**
An index copy scores	every gap is zero	**0%**

IN THE LESSON

- Let us calculate Active Share for the Roskin fund.
- Six stocks, A to F, are in play, with the index weights shown, and the fund holds them in different weights.
- Compare stock by stock.
- The fund is overweight stocks C, E and F, by fifteen, five and ten points: thirty in all.
- It is underweight A and D by fifteen and fifteen: thirty again.
- Overweights always equal underweights, because both portfolios are fully invested.
- Add the absolute gaps and you get sixty.
- Half of that is thirty per cent.
- So Active Share is thirty per cent, which means that share of the fund differs from the index.
- A pure index copy would score zero, because every gap would be nothing.

11 Four kinds of active manager

Chart: Active risk (% a year) vs Active Share (0 to 1)

- Closet indexer
- Diversified stock picker
- Factor or sector bettor
- Concentrated stock picker

IN THE LESSON

- Now put the two measures on one picture.
- Along the bottom is Active Share, from zero to one.
- Up the side is active risk, in per cent a year.
- A closet indexer holds nearly the index, so both numbers are low, yet it may still charge an active fee.
- At the top right is the concentrated stock picker, with few stocks, mostly unlike the index, so both numbers are high.
- Now the two interesting corners.
- The diversified stock picker, bottom right, holds many stocks very different from the index, but its many small bets are unrelated and partly cancel, so its active risk is modest.
- The factor or sector bettor, upper left, holds mostly index stocks but leans hard in one direction, so its Active Share is low and its active risk is high.

12 Example 4: why correlation matters

Four active bets, each adds 1.5% of active risk

Bets unrelated, correlation 0	$\sqrt{(4 \times 1.5^2)}$	**3.0%**
Partly linked, correlation 0.5	$\sqrt{(4 \times 1.5^2 \times 2.5)}$	**4.7%**
Moving as one, correlation 1	$\sqrt{(4 \times 1.5^2 \times 4)}$	**6.0%**
Same bets, risk compared	6.0 ÷ 3.0	**twice**

IN THE LESSON

- Why does a fund with many bets have modest active risk?
- Because of correlation.
- Ines has four active bets, and each one on its own adds one point five per cent of active risk.
- If the bets are unrelated, the risks partly cancel: square each, add, and take the square root, which gives three per cent.
- If the bets are partly linked, with a correlation of point five, the answer climbs to four point seven.
- If the bets all move as one, nothing cancels, and the answer is six per cent.
- Same four bets, same size, double the risk.
- This is why a diversified stock picker can be very different from the index and still stay close to it in returns.

13 Example 5: the sector bettor

Index: tech 30, 20; banks 20, 15, 15. Fund: 45, 30; 10, 7.5, 7.5

Absolute gaps	15+10+10+7.5+7.5	**50 pts**
Active Share	½ × 50	**25%**
Technology weight	75% against 50% in index	**+25 pts**
Tracking error	stated for the fund	**6%**
Reading	low Active Share, high risk	**Sector bettor**

IN THE LESSON

- Now the sector bettor.
- Suppose the index has five stocks.
- Two are technology, at thirty and twenty per cent of the index, and three are banks, at twenty, fifteen and fifteen.
- The fund holds the same five stocks, but technology at forty five and thirty, and the banks at ten, seven and a half and seven and a half.
- The gaps are fifteen, ten, ten, seven and a half and seven and a half, which add to fifty points.
- Half of fifty is twenty five per cent, so Active Share is low.
- But technology is seventy five per cent of the fund against fifty in the index, a twenty five point sector bet, and the fund's tracking error is six per cent a year.
- Low Active Share, high active risk: a sector bettor, whatever the stock list suggests.

14 Matching the numbers to the strategy

Active Share and active risk, read together

Manager type	Share / risk
Closet indexer	Low / low
Diversified stock picker	High / modest
Factor or sector bettor	Low / high
Concentrated stock picker	High / high

IN THE LESSON

- The table gives the reading rule.
- Low and low is the closet indexer.
- High Active Share with modest active risk is the diversified stock picker.
- That manager claims many small, independent insights and no skill in forecasting sectors, so this pattern is exactly what the claim predicts.
- If such a manager showed high active risk, we would ask what hidden common bet was driving it.
- Low Active Share with high active risk is the factor or sector bettor, whose claim is skill at calling themes.
- High and high is the concentrated stock picker.
- The test is always the same: does the pattern of the two measures fit what the manager says it is good at?

15 The trap: one number is not enough

- High Active Share does not mean high risk
- Low Active Share can still carry big bets
- Neither shows where the bets come from
- Read both against what the manager claims

IN THE LESSON

- The trap here is to treat either measure as a full description.
- High Active Share does not mean high risk, as the diversified stock picker shows.
- Low Active Share does not mean a quiet fund, as the sector bettor shows.
- Tracking error cannot tell you whether the holdings differ from the index.
- Active Share cannot tell you whether the differences move together.
- Neither shows where the bets come from, and neither tells you whether the manager has any skill.
- So when a committee asks whether a manager is genuinely active, the answer is a pair of numbers read against the manager's own claims.
- Then the fee question follows.
- Is the client paying for the kind of activity it wanted?

16 A story: the household budget

- Fix a total: for example 4% active risk a year
- Divide it among the sources of risk
- Each source should earn its share of reward
- Weight is not risk: size misleads

IN THE LESSON

- Think of a family budget.
- The family fixes a total for the month, then divides it among rent, food, heating and fun.
- A risk budget works the same way.
- The total might be an active risk of four per cent a year.
- The manager divides it among the places risk comes from, such as stock picking, sector tilts and style tilts.
- Then comes the test.
- Each source should use a share of the risk that matches the share of reward it is expected to bring.
- And here is the surprise.
- The size of a holding does not tell us its share of the risk.
- A small holding in a wild stock can use more of the budget than a big holding in a sleepy one.

17 Example 6: splitting a budget of 4%

Kinloss allows 4% active risk; the sources are unrelated

Total budget	4% squared	**16**
Stock selection	3.0% squared	**9**
Sector tilts	2.0% squared	**4**
Style tilts	1.73% squared	**3**
Check: add, then square root	9 + 4 + 3 = 16	**4.0%**

Shares of the risk: 56%, 25%, 19%

IN THE LESSON

- Here is a budget in numbers.
- Kinloss Pension Trust allows four per cent active risk.
- Variances of unrelated sources simply add, so we split in variance terms.
- Four per cent squared is sixteen.
- Ines splits it into stock selection at three per cent, which is nine; sector tilts at two per cent, which is four; and style tilts at one point seven three per cent, which is three.
- Nine plus four plus three is sixteen, and the square root of sixteen is four per cent.
- Stock selection takes fifty six per cent of the risk, sector tilts twenty five and style tilts nineteen.
- Now ask whether stock selection brings about fifty six per cent of her expected alpha.
- If her real edge is in sectors, she is spending risk where it earns the least.

18 Contribution to risk, and the efficient budget

Contribution = weight × covariance with portfolio
Efficient when marginal alpha ÷ marginal risk is equal for all

Weight	Money in the stock
Covariance	Moves with the portfolio
Contribution	Share of the variance
Marginal alpha	Extra alpha per extra weight
Marginal risk	Extra risk per extra weight

IN THE LESSON

- Now the rule behind a risk budget.
- A position's contribution to portfolio variance is its weight times its covariance with the portfolio.
- Covariance here measures how much the stock moves with the whole portfolio, so it captures both the stock's own volatility and its correlation.
- The contributions of all the positions add up exactly to the portfolio variance.
- For active risk, we do the same with active weights and the covariance of each stock's active return with the portfolio's active return.
- The budget is efficient when the marginal alpha, divided by the marginal contribution to risk, is the same for every position.
- If one position offers less alpha per unit of risk than another, move weight from it to the other.

19 Example 7: where the risk really sits

D 60%, vol 15%; V 40%, vol 40%; correlation 0.3

Covariance of D with portfolio	0.0135 + 0.0072	**0.0207**
Covariance of V with portfolio	0.0640 + 0.0108	**0.0748**
Contributions, D and V	0.6 × 0.0207; 0.4 × 0.0748	**0.0124; 0.0299**
Portfolio variance	0.0124 + 0.0299; vol 20.6%	**0.0423**
V's share of the risk	0.0299 ÷ 0.0423	**71%**

IN THE LESSON

- D is a defensive stock with fifteen per cent volatility, and V is a volatile stock with forty per cent.
- Their correlation is point three.
- D has sixty per cent of the money and V has forty.
- D's covariance with the portfolio is point zero two zero seven, and V's is point zero seven four eight.
- Multiply each by its weight and the contributions are point zero one two four for D and point zero two nine nine for V.
- They add to point zero four two three, the portfolio variance, which is a volatility of about twenty point six per cent.
- V's share of the risk is point zero two nine nine divided by point zero four two three, which is seventy one per cent.
- Forty per cent of the money, seventy one per cent of the risk.

20 Stock V: money, alpha and risk

Share of portfolio (%)

By weight	**40%**
By alpha	**50%**
By risk	**71%**

Stock V in the Roskin example

IN THE LESSON

- Here is the same finding as a picture.
- The vertical axis is the share of the portfolio, in per cent, and all three bars are for stock V.
- The first bar is its share of the money: forty per cent.
- The second bar is its share of the expected alpha.
- The manager expects each stock to supply about half of the alpha, so this bar is fifty per cent.
- The third bar is its share of the risk, and it towers above the other two at seventy one per cent.
- Stock V is using far more of the risk budget than it earns in alpha.
- The budget is out of balance, and the weight in V should come down until risk and reward line up.

21 Example 8: moving ten points from V to D

Alpha 3%, half from each stock; shift 10 points to D

Efficiency ratio, D	2.5 ÷ 0.0207	**121**
Efficiency ratio, V	3.75 ÷ 0.0748	**50**
New weights, D and V	move 10 points	**70 / 30**
Volatility after the move	20.6% to 18.2%	**18.2%**
Alpha after the move	3.00% to 2.88%	**2.88%**

Risk falls 12%, alpha falls only 4%

IN THE LESSON

- With a portfolio alpha of three per cent, split equally, D earns two point five per cent per unit of weight, and V earns three point seven five.
- Divide each by its covariance with the portfolio, and the efficiency ratio is one hundred and twenty one for D and fifty for V.
- D gives far more alpha per unit of risk.
- So move ten points from V to D, making the weights seventy and thirty.
- The portfolio volatility falls from twenty point six to eighteen point two per cent, a fall of twelve per cent.
- The alpha falls only from three per cent to two point eight eight, a fall of four per cent.
- Risk fell by much more than alpha, so each unit of risk now earns more.

22 Common mistakes so far

- Treating the optimiser's output as objective
- Judging a manager by Active Share alone
- Judging a manager by tracking error alone
- Budgeting risk by weight, not by contribution

IN THE LESSON

- Let us collect the traps from the first three topics.
- One: treating an optimiser's output as objective.
- It magnifies errors in the forecasts, so it is fenced in with limits or blended with judgement.
- Two: judging a manager by Active Share alone, which says nothing about whether the bets move together.
- Three: judging a manager by tracking error alone, which cannot tell a sector call from a spread of small stock bets.
- Four: budgeting risk by weight.
- A modest holding in a volatile stock used most of the risk.
- The cure for all four is the same: ask what the number measures, and what it leaves out, before you rely on it.

23 A story: guard rails on a mountain road

- Absolute: volatility, value at risk, conditional VaR
- Relative: active risk against the benchmark
- Caps: sector, country, factor and stock weights
- Also leverage and liquidity; rules or constraints

IN THE LESSON

- Guard rails on a mountain road do not steer the car.
- They stop one mistake from becoming a disaster.
- Limits in an equity portfolio do the same job, and they come in families.
- Absolute limits cap the total risk of the fund.
- They use volatility, value at risk, which is the loss that should not be exceeded on all but a few days, or conditional value at risk, which is the average loss on the days when that line is crossed.
- Relative limits cap active risk against the benchmark.
- Then come caps on sector, country and factor gaps, on single stocks, on leverage and on liquidity.
- A limit may be a simple rule or a formal constraint inside the optimiser.

24 Limits and what they control

Typical limits in a mandate

Limit	Controls
Volatility, VaR, CVaR	Absolute risk
Active risk	Against index
Sector, country, factor gaps	Theme bets
Stock weight, leverage, days to sell	Position size

IN THE LESSON

- This table sorts the limits by what they protect.
- Volatility, value at risk and conditional value at risk guard the absolute outcome: how much the client could lose.
- Active risk guards the relative outcome: how far the fund may drift from the index.
- Caps on sector, country and factor gaps stop the fund from taking one huge theme bet, even if each stock looks modest on its own.
- Caps on stock weight, leverage and days to sell guard the position sizes and the ability to get out.
- A good mandate uses several of these together, because each one covers a gap that the others leave open.

25 Value at risk, in one line

One-day VaR = z × (annual volatility ÷ √250) × portfolio value, with z = 1.65 for 95%

z value	1.65 for 95% confidence
Annual volatility	How much it swings a year
Square root of days	√250: trading days a year
Portfolio value	The money exposed

IN THE LESSON

- Value at risk is the most common absolute limit, so here is its formula in the simple normal form.
- Take the annual volatility and divide it by the square root of the number of trading days in a year, two hundred and fifty, to get daily volatility.
- Multiply by the z value for the confidence level, which is one point six five for ninety five per cent.
- Multiply by the value of the portfolio.
- The answer is a loss that should be exceeded on only a small share of days.
- Two things follow.
- First, measured from recent data, the number rises when markets get rough, so a limit on it bites at the wrong time.
- Second, the normal form understates very large losses, because real returns have fatter tails than the normal curve.

26 Example 9: a limit that binds

Fund $500m; limit: one-day 95% VaR of $10m

Daily volatility, calm	15% ÷ √250	**0.95%**
VaR, calm	1.65 × 0.95% × 500m	**$7.8m**
Daily volatility, sell-off	25% ÷ √250	**1.58%**
VaR, sell-off	1.65 × 1.58% × 500m	**$13.0m**
Exposure the limit allows	500m × 10m ÷ 13.04m	**$383m**

Forced sale: about $117m, into a falling market

IN THE LESSON

- Roskin's fund holds five hundred million dollars, and its mandate caps one day ninety five per cent value at risk at ten million.
- In calm markets volatility is fifteen per cent.
- Daily volatility is point nine five per cent.
- Value at risk is one point six five times point nine five per cent times five hundred million, which is seven point eight million, inside the limit.
- Then markets sell off and volatility rises to twenty five per cent.
- Daily volatility is one point five eight per cent, and value at risk is thirteen point zero million, over the limit.
- The exposure the limit allows is five hundred million times ten over thirteen point zero four, about three hundred and eighty three million.
- Nothing about Ines's views changed, yet she must sell.

27 Calm market against sell-off

One-day VaR ($m)

Calm market	**$7.8m**
Sell-off	**$13.0m**
The limit	**$10m**

Same holdings in the first two bars

IN THE LESSON

- Here is the same story as bars.
- The vertical axis is one day value at risk, in millions of dollars.
- The first bar is the calm market, and it stands at seven point eight million, comfortably under the limit.
- The second bar is the sell-off.
- Value at risk has jumped to thirteen point zero million, and it towers over the third bar, which is the limit of ten million.
- Only the measured volatility changed.
- This is why a risk cap needs a plan.
- Holding more liquid stocks, or using index futures to cut exposure cheaply, lets a manager obey the cap without dumping its best ideas at poor prices.

28 What a tighter active risk limit does

Tighten the active risk limit

Smaller active bets → Closer to the index → Lower Active Share

What the client then gets

Less expected alpha → Target may be missed

IN THE LESSON

- Relative limits have a quieter effect.
- Suppose the client tightens the active risk limit.
- The manager must make smaller bets, so the first box is smaller active bets.
- Smaller bets mean the portfolio sits closer to the index, the second box, and since the active weights are smaller, Active Share falls, the third box.
- On the second row is what the client then gets.
- With smaller bets, the manager earns less alpha.
- And if the client still wants the old return target, the uncomfortable result is in the last box: the target may be missed.
- Tight limits and high targets can quietly contradict each other, and the adviser's job is to notice before the money is invested.

29 Example 10: the limit and the target clash

Active risk 5%, Active Share 0.70, information ratio 0.5

Alpha at a 5% limit	0.5 × 5%	**2.5%**
Alpha at a 3% limit	0.5 × 3%	**1.5%**
Active Share after scaling	0.70 × 3 ÷ 5	**0.42**
Ratio needed for 2.5% alpha	2.5 ÷ 3	**0.83**
Verdict	limit and target clash	**Inconsistent**

IN THE LESSON

- Kinloss's mandate allows five per cent active risk, and Ines has an information ratio of point five.
- That ratio is alpha per unit of active risk.
- At five per cent, her expected alpha is point five times five, or two point five per cent, which is the target.
- After a poor year, the trust cuts the limit to three per cent.
- If she shrinks every active weight by three over five, expected alpha becomes point five times three, which is one point five per cent.
- Active Share falls from point seven zero to point four two.
- To still earn two point five per cent, she would need an information ratio of two point five over three, which is point eight three.
- The limit and the target are inconsistent.

30 Alpha is bought with active risk

Chart: Expected alpha (% a year) vs Active risk limit (% a year)

- Ratio 0.5
- Target 2.5%
- 3%: 1.5% alpha
- 5% needed

IN THE LESSON

- Let us draw it.
- Along the bottom is the active risk limit, in per cent a year.
- Up the side is expected alpha, also in per cent a year.
- The rising line is Ines with an information ratio of point five: every point of active risk she is allowed buys half a point of alpha.
- The flat line is Kinloss's target of two point five per cent.
- They meet at five per cent of active risk, where the old mandate was set.
- Now look at the three per cent limit on the bottom axis.
- Going up from it to the line, expected alpha is only one point five.
- The gap between that point and the target line is the shortfall created by the tighter limit, and no amount of hoping closes it.

31 Example 11: caps on a stock and on liquidity

Roskin fund $2,000m; sell within 5 days at 20% of volume

Optimiser asks for stock Q	index weight 3%	**9%**
Cap on stock Q	3% + 3 points	**6%**
Stock R: sold per day	20% × 8m a day	**1.6m**
Most held if sold in 5 days	5 × 1.6m ÷ 2,000m	**0.4%**
Days to sell a 1% position	20m ÷ 1.6m	**12.5 days**

IN THE LESSON

- Two more caps show how limits reshape the portfolio.
- First, a stock cap.
- The optimiser asks for nine per cent in stock Q, whose index weight is three per cent.
- Ines's rule is the index weight plus three points, so the cap is six per cent.
- Second, a liquidity cap.
- Her rule is that every position must be sellable within five days, trading no more than twenty per cent of the stock's daily volume.
- Stock R trades eight million dollars a day, so she can sell one point six million a day.
- In five days, that is eight million, which is point four per cent of the two thousand million fund.
- A one per cent position, twenty million, would take twelve and a half days to sell.

32 The trap: limits that fight each other

- A VaR cap can bind exactly when markets fall
- Tight active risk limits starve the alpha target
- Test the limits together, against the target
- Plan for the day a limit binds

IN THE LESSON

- The last trap is limits that disagree with each other, or with the return target.
- A value at risk cap can bind in volatile markets and force selling at the worst moment.
- A tight active risk limit shrinks Active Share and expected alpha, as we just saw.
- None of these is wrong on its own.
- The mistake is setting them one at a time.
- So test the limits together, against the alpha target, and ask whether the mandate can realistically meet its objective.
- Then ask what the manager will do on the day a limit binds.
- A written plan is cheaper than a forced sale.

33 Recap: the lesson in five lines

- Style: rules or judgement, stocks or themes
- Active Share and active risk: read the pair
- Contribution = weight × covariance; budget by risk
- Limits fence the portfolio and shrink alpha
- Check limits against the return target

IN THE LESSON

- Let us pull the lesson together.
- First, active portfolios differ in how they decide and where they start, and positions are sized by tiers, limits or an optimiser, which magnifies forecast errors and so needs fences.
- Second, Active Share is half the sum of absolute gaps, active risk depends on correlation, and the pair of numbers must fit the manager's claims.
- Third, a risk budget assigns each position a contribution of weight times covariance, and is efficient when alpha per unit of marginal risk is equal.
- Fourth, limits on volatility, value at risk and active risk fence the portfolio.
- Fifth, tight limits lower Active Share and alpha, and they must be consistent with the return target.

34 Try three quick questions

1. Index 50, 30, 20; fund 40, 20, 40. Active Share?
   Answer: Gaps 10 + 10 + 20 = 40, and half of 40 is 20%.
2. Weight 30%, covariance 0.05, portfolio variance 0.04. Share of risk?
   Answer: 0.3 × 0.05 = 0.015; 0.015 ÷ 0.04 = 37.5%.
3. $200m, volatility 20%, 250 days, 95% (z = 1.65). One-day VaR?
   Answer: 20% ÷ √250 = 1.26%; × 1.65 × 200m ≈ $4.2m.

IN THE LESSON

- Now try three quick questions yourself.
- Pause after each one.
- Question one.
- A benchmark holds three stocks at fifty, thirty and twenty per cent, and a fund holds them at forty, twenty and forty.
- What is the Active Share?
- The gaps are ten, ten and twenty, which add to forty, and half of forty is twenty per cent.
- Question two.
- A stock has a weight of thirty per cent and a covariance with the portfolio of point zero five, and the portfolio variance is point zero four.
- What share of the risk does it use?
- Point three times point zero five is point zero one five, which is thirty seven point five per cent of the variance.
- Question three.
- A fund of two hundred million dollars has annual volatility of twenty per cent.
- What is its one day value at risk at ninety five per cent?
- Daily volatility is one point two six per cent and the answer is about four point two million dollars.