CHARTA LEARNING · LESSON NOTES

Portfolio Construction

FRM Part II

1 A story: the optimiser that believed everything

- Harrow Core Equity: active fund, 500-stock benchmark
- Analyst forecasts go straight into an optimiser
- First proposal: 12% of the fund in one small stock
- One forecast was too bold; the optimiser believed it
- Construction turns forecasts into holdings

IN THE LESSON

- Picture Harrow Asset Management.
- It runs Harrow Core Equity, an active fund judged against a benchmark of five hundred stocks.
- Its analysts send their forecasts straight into a new optimiser, which is a program that picks the weights giving the best trade off between expected reward and risk.
- The first portfolio it proposes puts twelve per cent of the fund in one small stock.
- Nobody intended that.
- One forecast was far too bold, and the optimiser believed it completely.
- This lesson is about avoiding that outcome.
- Portfolio construction is the step that turns forecasts into holdings, and its quality depends on what goes in.

2 Four inputs, one objective

Value added = α_P − λ × ψ^2 − TC

Alpha	Expected return beyond benchmark
Lambda	Active risk aversion
Risk model	Covariances; gives active risk ψ
Costs	Transaction costs of the trades

IN THE LESSON

- Portfolio construction takes four main inputs, and the optimiser joins them in one objective.
- Value added equals the portfolio's alpha, minus lambda times active risk squared, minus transaction costs.
- Alpha is the first input: forecasts, for each asset, of return over and above the benchmark.
- Lambda is the active risk aversion, which sets how much risk the manager accepts for each unit of alpha.
- Third is the risk model, which supplies the covariance matrix of returns, showing how assets move together.
- Active risk, psi, is not supplied: it is computed from that matrix and the weights, as the standard deviation of return relative to the benchmark.
- The last input is the estimates of transaction costs.
- The optimiser picks the weights that make value added as large as possible.

3 Why every input is hard

- Alphas: noisy and usually overstated
- Covariances: unstable from short samples
- More assets than observations: singular, so use factors
- Costs: depend on size and markets, known only after
- Optimiser magnifies errors: garbage in, amplified out

IN THE LESSON

- Each input is hard to get right.
- Alphas are noisy, and they tend to be overstated, because forecasters are confident while their signals are weak.
- Covariance estimates from short samples are unstable.
- With more assets than observations, the sample matrix is singular, which means some combinations of assets falsely look riskless, so a factor model is normally used instead.
- Costs depend on trade size and on market conditions, and they are known only after trading.
- And the optimiser magnifies every error, because it takes large positions wherever alpha is overstated or risk is understated.
- The exam phrase to remember is: garbage in, amplified garbage out.

4 Example 1: too many numbers, too little data

Harrow: 500 stocks and 60 months of returns. How many numbers must it estimate?

Numbers to estimate	500 + 124,750 pairs	**125,250**
Data available	60 months × 500 stocks	**30,000**
Observations vs assets	60 < 500	**Singular**
Factor model, 10 factors	5,000 + 55 + 500	**5,555**

Pairs: 500 × 499 ÷ 2. Factors: 10 + 45 = 55

IN THE LESSON

- Harrow has five hundred stocks and sixty months of returns.
- How many numbers must it estimate?
- The matrix needs five hundred variances, plus one covariance for every pair of stocks, one hundred and twenty four thousand, seven hundred and fifty pairs.
- So the total is one hundred and twenty five thousand, two hundred and fifty numbers.
- The data are sixty months times five hundred stocks, only thirty thousand returns.
- Worse, sixty observations is fewer than five hundred assets, so the sample matrix is singular.
- A factor model with ten factors needs five thousand exposures, fifty five factor variances and covariances, and five hundred specific variances, the risk unique to each stock.
- That is just five thousand, five hundred and fifty five numbers.
- The pairs are five hundred times four hundred and ninety nine, divided by two; the fifty five is ten variances plus forty five covariances.

5 Example 2: the optimiser magnifies errors

Stock A: α 3%, residual risk ω 25%, λ = 8. Best weight h = α ÷ (2λω^2)

Denominator 2 × λ × ω^2	2 × 8 × 0.25^2	**1.00**
True alpha 3%	0.03 ÷ 1.00	**3%**
Alpha overstated at 12%	0.12 ÷ 1.00	**12%**
Risk understated at 12.5%	0.03 ÷ (2×8×0.125^2)	**12%**

Quadruple alpha or halve risk: four times the position

IN THE LESSON

- Stock A has an alpha of three per cent and residual risk, written omega, of twenty five per cent.
- Residual risk is the risk not explained by the benchmark.
- Harrow uses a lambda of eight.
- For a stock uncorrelated with the rest, the best active weight, h, is alpha divided by two times lambda times omega squared.
- The denominator, two times eight times point two five squared, is exactly one.
- So with the true alpha, the weight is point zero three divided by one, three per cent.
- If the analyst overstates alpha at twelve per cent, the weight becomes twelve per cent.
- If alpha is right but risk is understated at twelve and a half per cent, the weight is again twelve per cent.
- Quadrupling alpha or halving risk gives four times the position.

6 Why alphas are refined

- Raw alphas: noise, overconfidence, unintended biases
- Optimiser: large or wrong alpha, large position
- Scaling: size alphas to the information
- Trimming: cap extreme values
- Neutralising: remove unintended bets

IN THE LESSON

- Raw alphas are rarely fit to feed straight into an optimiser.
- They carry noise, overconfidence and unintended biases, such as a hidden bet on the whole market or on one industry.
- And as we have just seen, an optimiser turns a large or mistaken alpha into a large position.
- So alphas are refined during implementation, in three common steps.
- Scaling sets the size of the alphas to what the information can support.
- Trimming caps extreme values, so that no single forecast dominates.
- Neutralising removes unintended bets.
- Each step makes the positions match what the manager really knows.

7 Scaling: the rule of thumb

Alpha = volatility × IC × score

Volatility	Volatility, usually residual
IC	Information coefficient: skill
Score	Standardised signal, SD of one

IN THE LESSON

- Scaling uses a rule of thumb.
- Alpha equals volatility times I C times score.
- Volatility is the stock's volatility, in practice usually its residual volatility.
- I C, the information coefficient, measures skill: the correlation between the forecasts and the returns that follow.
- The score is a standardised signal, with a cross sectional standard deviation of one.
- So a typical stock scores between minus one and plus one, and a score of two is unusually strong.
- Because the scores have a spread of one, the spread of alphas across stocks is volatility times I C.
- Weak skill means small alphas, and that keeps the optimiser modest.

8 Example 3: scaling Harrow's alphas

Residual volatility 25%, IC 0.06. The raw forecast for Stock A was 12%

Spread of alphas	25% × 0.06	**1.5%**
Stock A, score +2	25% × 0.06 × 2	**3.0%**
Stock B, score −1	25% × 0.06 × (−1)	**−1.5%**
Raw forecast vs scaled	12% ÷ 3%	**4 times too big**
Every alpha doubled	weights, risk × 2; A: 2 × 3%	**6%**

IN THE LESSON

- Residual volatility is twenty five per cent, I C is point zero six, and the analyst's raw forecast for Stock A was twelve per cent.
- First, the spread of alphas across stocks is twenty five per cent times point zero six, one and a half per cent.
- Stock A has a score of plus two, so its alpha is twenty five per cent times point zero six times two, three per cent.
- Stock B scores minus one, so its alpha is minus one and a half per cent.
- The raw twelve per cent was four times too big.
- And with no costs or constraints, doubling every alpha doubles every active position and the active risk, so Stock A's weight goes from three to six per cent.

9 Example 4: trimming an extreme forecast

Stock C: residual risk 25%, score +5. Harrow caps scores at 3 standard deviations

Untrimmed alpha	25% × 0.06 × 5	**7.5%**
Untrimmed weight, λ = 8	0.075 ÷ 1.00	**7.5%**
Trimmed score	capped at 3	**3**
Trimmed alpha and weight	25% × 0.06 × 3	**4.5%**

First check an extreme score for a data error

IN THE LESSON

- Stock C has residual risk of twenty five per cent, but its score is plus five, which is extreme, and Harrow caps scores at three standard deviations.
- Untrimmed, its alpha is twenty five per cent times point zero six times five, seven and a half per cent.
- With lambda of eight, the denominator is still one, so its weight would be point zero seven five divided by one, also seven and a half per cent.
- Trimming caps the score at three.
- The trimmed alpha is twenty five per cent times point zero six times three, four and a half per cent, and so is the weight.
- One forecast no longer dominates.
- But first check whether an extreme score is simply a data error.

10 Neutralisation: removing unintended bets

- Neutralise: do not reward unintended exposures
- Cheerful on the market? A hidden beta bet
- Benchmark neutral: benchmark's weighted alpha is zero
- Cash neutral (long-short): weighted average alpha zero
- Factor neutral: strip alpha of each unwanted factor

IN THE LESSON

- The third step is neutralisation.
- It means adjusting alphas so that the optimiser is not rewarded for taking exposures the manager did not intend to take.
- For example, suppose Harrow's analysts are cheerful about the market as a whole.
- Then stocks with a high beta, meaning a high sensitivity to the benchmark, get high alphas, and the optimiser quietly builds a hidden beta bet.
- There are three common forms.
- Benchmark neutral alphas give the benchmark a weighted average alpha of zero.
- Cash neutral alphas, for a long short portfolio, are shifted so that their weighted average is zero.
- Factor neutral alphas strip out the alpha linked to each unwanted factor.

11 Benchmark-neutral alphas

Adjusted α = α − β × α_B
α_B = Σ w_B × α

Alpha	Raw alpha of the asset
Beta	Asset's beta to the benchmark
Benchmark alpha	Benchmark's weighted alpha

IN THE LESSON

- Here is the benchmark neutral adjustment.
- Each adjusted alpha equals the raw alpha minus the asset's beta times the benchmark's alpha.
- The benchmark's alpha is the sum, over all stocks, of benchmark weight times alpha.
- Alpha is the raw forecast for the asset.
- Beta is the asset's sensitivity to the benchmark.
- And alpha B is the benchmark's weighted alpha.
- Because the benchmark has a beta of one, the adjusted alphas give it a weighted alpha of exactly zero.
- A useful result follows.
- With no other constraints, benchmark neutral alphas lead to a portfolio with zero active beta, so there is no hidden market bet.
- But constraints, such as a ban on short sales, can bring exposures back, so always check the result.

12 Example 5: making alphas benchmark-neutral

Weights X 50%, Y 30%, Z 20%. Alphas 2.0, 1.0, −0.5%. Betas 1.2, 1.0, 0.5

Benchmark alpha	0.5×2.0 + 0.3×1.0 − 0.2×0.5	**1.2%**
Stock X	2.0 − 1.2 × 1.2	**0.56%**
Stock Y	1.0 − 1.0 × 1.2	**−0.20%**
Stock Z	−0.5 − 0.5 × 1.2	**−1.10%**
Check: benchmark alpha	0.28 − 0.06 − 0.22	**0.00%**

Benchmark beta: 0.6 + 0.3 + 0.1 = 1.0

IN THE LESSON

- Stocks X, Y and Z have benchmark weights of fifty, thirty and twenty per cent, alphas of two, one and minus a half per cent, and betas of one point two, one point zero and point five.
- The benchmark's alpha is point five times two, plus point three times one, minus point two times a half: one point two per cent.
- That is the hidden market bet.
- Stock X's adjusted alpha is two minus one point two times one point two, point five six per cent.
- Stock Y's is one minus one point two, minus point two per cent.
- Stock Z's is minus a half minus point six, minus one point one per cent.
- Check: weighted, these give point two eight, minus point zero six and minus point two two, which sum to zero.
- It works because the benchmark's beta is point six plus point three plus point one, exactly one.

13 Example 6: cash-neutral and factor-neutral

Long-short alphas average +0.4%. Value factor portfolio alpha 1.5%: unwanted

Cash neutral: a 1.0% alpha	1.0 − 0.4	**0.6%**
Factor neutral: the rule	α − exposure × factor alpha	**per factor**
Stock P: value exposure 0.8	2.0 − 0.8 × 1.5	**0.80%**
Stock Q: value exposure −0.4	0.5 − (−0.4) × 1.5	**1.10%**
Then check the portfolio	constraints can add bets back	**Check**

IN THE LESSON

- Harrow also runs a long short book whose alphas have a weighted average of plus point four per cent.
- And Harrow does not want a bet on value, whose factor portfolio has an alpha of one and a half per cent.
- Cash neutral first: subtract the average from every alpha, so a one per cent alpha becomes point six per cent, and the weighted average becomes zero.
- Factor neutral next: for each unwanted factor, subtract the stock's exposure times that factor portfolio's alpha.
- Stock P, with exposure point eight, goes from two per cent to two minus one point two, point eight per cent.
- Stock Q, with exposure minus point four, goes from point five to point five plus point six, one point one per cent.
- Last, check the portfolio, because constraints can add the bets back.

14 How transaction costs enter

- Costs: a penalty on trades, buys and sells
- Purchase and sale costs usually differ
- Alpha is yearly; a cost is paid once: amortise it
- Trade only if gain in alpha less risk beats the cost
- Higher costs: lower turnover. Impact rises with size

IN THE LESSON

- Now transaction costs.
- They enter the optimisation as a penalty on trades, charged on purchases and on sales.
- The purchase cost and the sale cost usually differ.
- The key point is timing.
- Alpha is quoted per year, while a cost is paid once, so the cost must be amortised, meaning spread, over the expected holding period.
- A trade is worthwhile only if the improvement in alpha less risk exceeds that amortised cost.
- So higher costs lower turnover, and keep the portfolio nearer its current holdings.
- And market impact, which is the price move caused by the trade itself, usually rises with trade size.

15 Example 7: amortising the cost

Gain in alpha less risk: 2.0% a year. Buy cost 1.2%, sell 1.8%. Holding 2 years

Round-trip cost	1.2% + 1.8%	**3.0%**
Amortised over 2 years	3.0% ÷ 2	**1.5% a year**
Compare	2.0% > 1.5%	**Trade**
Wrong: the full cost	2.0% < 3.0%	**Reject**
If held only 1 year	3.0% ÷ 1 > 2.0%	**Do not trade**

Cost a year = round-trip cost ÷ holding years

IN THE LESSON

- A proposed trade improves alpha less risk by two per cent a year.
- Buying costs one point two per cent, selling later costs one point eight per cent, and Harrow expects to hold for two years.
- The round trip cost is one point two plus one point eight, three per cent.
- Amortised over two years, it is one and a half per cent a year.
- Two per cent beats one and a half, so the trade is worthwhile.
- The common mistake compares the annual two per cent with the full three per cent, and wrongly rejects the trade.
- If Harrow held for only one year, the cost a year would be the full three per cent, more than two, so it would not trade.
- In short, the cost a year equals the round trip cost divided by the holding years.

16 Example 8: splitting trades, and keeping holdings

Impact: 1.0% on a $10m trade, 0.6% on $5m. Amortised purchase cost: 1.0% a year

One $10m trade	1.0% × $10m	**$100,000**
Two $5m trades	2 × 0.6% × $5m	**$60,000**
Saving from splitting	100,000 − 60,000	**$40,000**
Held, worth 0.5% a year	no purchase cost to keep it	**Keep**
Same stock, not held	0.5% < 1.0% to buy	**Do not buy**

IN THE LESSON

- Market impact is one per cent on a ten million dollar trade, but only point six per cent on a five million dollar trade.
- Separately, the amortised purchase cost of a stock is one per cent a year.
- Trading ten million at once costs one per cent of ten million, one hundred thousand dollars.
- Two trades of five million cost two times point six per cent times five million, sixty thousand dollars.
- So splitting saves forty thousand dollars.
- Next, a stock Harrow already holds is worth point five per cent a year.
- Keeping it needs no purchase, so no cost is paid, and Harrow keeps it.
- The same stock, if not held, would need more than one per cent a year to be worth buying, so Harrow does not buy it.
- A held asset is kept longer than an equivalent unheld one is bought.

17 Setting risk aversion from a target

Unconstrained: active risk = IR ÷ (2 × λ)
λ = IR ÷ (2 × target risk)
Alphas × k = λ ÷ k

Active risk	Active risk, in decimals
IR	Information ratio: alpha per risk
Lambda	Active risk aversion

IN THE LESSON

- Now some practical choices, starting with risk aversion.
- It should be set from a target, not guessed.
- With unconstrained optimisation, active risk equals the information ratio divided by two times lambda.
- Turned round, lambda equals the information ratio divided by two times the target risk.
- Active risk is written in decimals, so four per cent is point zero four.
- I R, the information ratio, is the alpha a manager expects per unit of active risk, a measure of skill.
- Lambda is the active risk aversion from our objective.
- One more link: scaling all alphas up by a factor has the same effect on positions as dividing lambda by that factor.
- So alphas that are too large, and aversion that is too low, are the same mistake.

18 Example 9: Harrow's risk aversion

Harrow: IR 0.4 and λ = 8 today. The board sets a target active risk of 4%

Active risk today	0.4 ÷ (2 × 8)	**2.5%**
λ for a 4% target	0.4 ÷ (2 × 0.04)	**5**
Check	0.4 ÷ (2 × 5)	**4.0%**
Same as alphas × 1.6	8 ÷ 1.6	**λ = 5**
Trap: target typed as 4	0.4 ÷ 8; risk 400%	**λ = 0.05**

IN THE LESSON

- Harrow's information ratio is point four, its lambda today is eight, and its board sets a target active risk of four per cent.
- Today's active risk is point four divided by two times eight, which is point zero two five, or two and a half per cent.
- To hit four per cent, lambda is point four divided by two times point zero four, which is five.
- Check: point four divided by two times five is point zero four, four per cent.
- Lowering lambda from eight to five has the same effect as scaling every alpha up by one point six, since eight divided by one point six is five.
- Now the trap.
- Type the target as four instead of point zero four, and lambda comes out at point four over eight, point zero five, which would mean active risk of four hundred per cent.
- Write risk in decimals.

19 Particular risks, and full coverage

- Dislike one risk? Add a penalty on that risk alone
- Raising overall λ would shrink every bet
- Alphas must cover the whole universe
- No forecast: a neutral, benchmark-consistent alpha
- An arbitrary zero can create an unintended bet

IN THE LESSON

- Two more practical choices.
- First, a manager may dislike one particular risk more than others, such as a sector or a common factor.
- The answer is to add a separate penalty on that risk alone.
- Simply raising the overall risk aversion would shrink every bet, including the good ones.
- Second, alphas should cover the whole universe.
- Analysts rarely forecast every stock in the benchmark.
- Assets with no forecast must be given a neutral, benchmark consistent alpha, because an arbitrary value, such as a careless zero, can create an unintended bet between the stocks that are covered and those that are not.

20 Example 10: the stocks nobody covers

Covered: 80% of benchmark weight, alphas average +1.0%. The rest are given zero

Benchmark alpha before	80% × 1.0 + 20% × 0	**0.8%**
Uncovered stocks look	0 against 1.0	**1.0% worse**
Optimiser's response	sells uncovered, buys covered	**Unintended bet**
Fix: neutralise covered	each covered alpha − 1.0	**average 0**
Benchmark alpha after	80% × 0 + 20% × 0	**0.0%**

IN THE LESSON

- Analysts at Harrow cover stocks making up eighty per cent of the benchmark's weight, and their alphas have a benchmark weighted average of plus one per cent.
- The other stocks are simply given zero.
- The benchmark alpha before any fix is eighty per cent times one, plus twenty per cent times zero, point eight per cent.
- Against the covered average of one, the uncovered stocks at zero look one per cent worse.
- So the optimiser sells the uncovered stocks and buys covered ones: an unintended bet.
- One fix is to neutralise the covered alphas by subtracting one per cent from each, so they average zero.
- Then zero is a neutral alpha for the uncovered stocks, and the benchmark alpha after the fix is zero.
- Some sources instead give uncovered stocks the covered average; either way they match covered ones.

21 Revising and rebalancing

- Revising: update alphas and risk as news arrives
- Rebalancing: trade toward the new optimum
- Also corrects drift from price moves
- Trade only if the annual gain beats the amortised cost
- Noisier alphas, higher costs: lower turnover

IN THE LESSON

- Now what happens over time.
- Revising a portfolio means updating the alphas and risk estimates when new information arrives.
- Rebalancing means trading the existing portfolio toward the new optimum.
- Rebalancing also corrects drift, when price moves push weights and risk away from target.
- The rule is the one we met for costs: trade only when the annual gain from the revision exceeds the amortised cost.
- And note this.
- Noisier alphas and higher costs call for lower turnover, not higher, because a noisy revision is worth less than it appears.

22 Four forces that pull against each other

Each force pulls rebalancing a different way

Force	Pulls toward
Alpha	Trade quickly, above all if signals decay fast
Risk	Trade back when drift lifts risk above target
Transaction costs	Trade less
Horizon	Longer: cost spread over more years, more trading

IN THE LESSON

- Each force pulls rebalancing a different way.
- Alpha argues for trading quickly, above all when the signal decays fast, because a forecast that fades within months must be captured soon or it is lost.
- Risk argues for trading back when drift has lifted risk above target.
- Transaction costs argue for trading less.
- The investment horizon decides how many years of benefit repay a cost.
- A longer horizon spreads the cost over more years, and so supports more trading.

23 Example 11: does the revision pay?

A revision adds 0.8% a year. The round-trip cost is 2.4%. Should Harrow trade?

Horizon 2 years	2.4 ÷ 2 = 1.2 > 0.8	**Do not trade**
Horizon 4 years	2.4 ÷ 4 = 0.6 < 0.8	**Trade**
Noisier alpha, IC halved	gain 0.8 ÷ 2	**0.4% a year**
Break-even horizon	2.4 ÷ 0.4	**6 years**

IN THE LESSON

- A revision would add point eight per cent a year, and the round trip cost of the trades is two point four per cent.
- Should Harrow trade?
- With a two year horizon, the cost a year is two point four divided by two, one point two per cent.
- That is more than point eight, so do not trade.
- With a four year horizon, the cost a year is point six, less than point eight, so trade.
- Now suppose the signal is noisier and the I C is halved.
- Since alpha is proportional to I C, the gain halves to point four per cent a year.
- The break even horizon is two point four divided by point four, six years.
- So noisier alphas mean lower turnover.

24 The no-trade rule

MCVA = α − 2 × λ × cov
Buy if MCVA > PC; sell if MCVA < −SC

Alpha	Annual alpha of the asset
Lambda	Active risk aversion
Covariance	Covariance with active portfolio
Purchase cost	Purchase cost, amortised
Sale cost	Sale cost, amortised

IN THE LESSON

- Now the no trade region.
- With transaction costs, there is a range in which doing nothing is best.
- For each asset, the marginal contribution to value added, M C V A, is alpha less a marginal risk penalty: two times lambda times the covariance of the asset with the active portfolio.
- Buy if M C V A exceeds the purchase cost, P C.
- Sell if it is below minus the sale cost, S C.
- Otherwise, hold.
- Alpha is the asset's annual alpha.
- Lambda is the active risk aversion.
- The covariance is with the active portfolio, the fund's bets against the benchmark.
- P C and S C are the purchase and sale costs, amortised to the same annual basis as alpha.

25 Example 12: Stock K and the no-trade region

Stock K: α 2.0%, cov 0.0015, λ = 5. Buy cost 1.2%, sell 1.8%, horizon 2 years

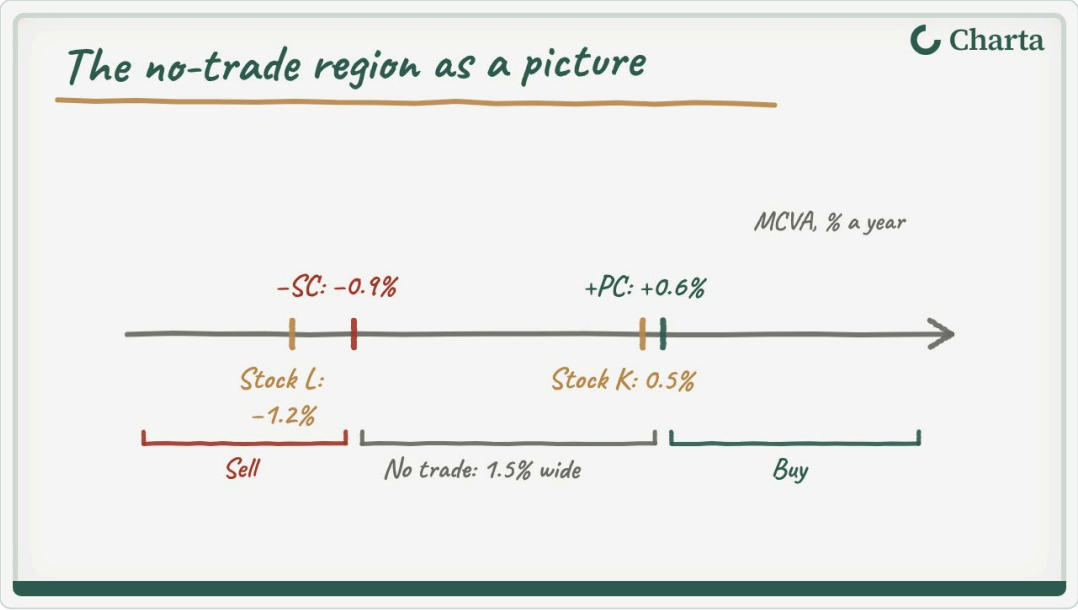
Risk penalty	2 × 5 × 0.0015	**1.5%**
MCVA	2.0% − 1.5%	**0.5%**
Amortised PC and SC	1.2 ÷ 2 and 1.8 ÷ 2	**0.6% and 0.9%**
No-trade region	from −SC to +PC	**−0.9% to 0.6%**
Verdict	−0.9 < 0.5 < 0.6	**Do nothing**

Wrong: alpha 2.0% > 0.6% says buy. Use MCVA, not alpha

IN THE LESSON

- Harrow now runs at a lambda of five, from its target.
- Stock K, which it does not hold, has an alpha of two per cent and a covariance of point zero zero one five with the active portfolio.
- Buying costs one point two per cent and selling one point eight, over a two year horizon.
- The risk penalty is two times five times point zero zero one five, point zero one five, or one and a half per cent.
- M C V A is two minus one and a half, point five per cent.
- Amortised, the purchase cost is point six per cent and the sale cost point nine.
- So the no trade region runs from minus point nine to plus point six.
- Point five lies inside it, so Harrow does nothing.
- The usual mistake uses alpha instead: two per cent beats point six, and says buy.

26 The no-trade region as a picture



IN THE LESSON

- The line measures M C V A, in per cent a year, rising to the right.
- The first mark is minus S C, at minus point nine per cent.
- The second is plus P C, at plus point six.
- Stock K sits at point five, just inside.
- A second stock, Stock L, which Harrow holds, has an M C V A of minus one point two per cent.
- Below the line are the three zones.
- Left of minus point nine, sell.
- Between the marks is the no trade region, point nine plus point six, so one and a half per cent wide.
- Right of plus point six, buy.
- So Harrow sells Stock L, and leaves Stock K alone.

27 Example 13: what widens the region

Buy cost 1.2%, sell cost 1.8%. What if the horizon or the costs change?

Width, 2-year horizon	0.6 + 0.9	**1.5%**
1-year horizon: region	PC 1.2, SC 1.8	**−1.8% to 1.2%**
Width, 1-year horizon	1.2 + 1.8	**3.0%**
Stock L, MCVA −1.2%	1 year: −1.2 > −1.8	**Hold**
Zero costs: Stock K	region is a point: 0.5 > 0	**Buy**

Traps: alpha for MCVA, no amortisation, sides swapped

IN THE LESSON

- Now change the inputs, keeping purchase cost of one point two per cent and a sale cost of one point eight.
- With the two year horizon, the width is point six plus point nine, one and a half per cent.
- With a one year horizon nothing is spread, so P C is one point two and S C is one point eight, and the region runs from minus one point eight to plus one point two.
- Its width is one point two plus one point eight, three per cent: a shorter horizon widens the region.
- Stock L, at minus one point two, is now above minus one point eight, so Harrow holds it.
- With zero costs, the region shrinks to a point at zero, and Stock K, at point five, is bought.
- Remember the traps: using alpha instead of M C V A, forgetting amortisation, and swapping the buy and sell sides.

28 Four ways to build a portfolio

Four methods trade simplicity against rigour

Method	Does well	Does badly
Screens	Simple, transparent	Ignore correlation, size, exposures
Stratification	Controls exposures	Discards alpha size
Linear programming	Alpha under linear limits	No variance penalty
Quadratic programming	Covariance, costs, limits	Exploits input errors

IN THE LESSON

- Four construction methods trade simplicity against rigour.
- Screens rank assets on alpha and hold the top group, usually equal weighted.
- They are simple and transparent, but they ignore correlation, the size of alpha and benchmark exposures, so bets can cluster.
- Stratification splits the universe into cells, such as sector by size, matches the benchmark's weight in each cell and picks stocks within it.
- It controls unintended exposures, but discards information about alpha size.
- Linear programming maximises alpha under linear constraints on weights, turnover or exposures, but has no variance penalty for risk.
- Quadratic programming maximises alpha less a risk penalty, using the full covariance matrix, costs and constraints.
- It is the most complete, but it exploits input errors, so inputs must be refined.

29 Example 14: a screen, then stratification

Harrow screens its 500 stocks and holds the top 50, equal-weighted

Weight of each stock	100% ÷ 50	**2% each**
Alphas of 4.0% and 0.5%	both get the same weight	**2%**
30 of the 50 are banks	30 × 2%; benchmark 15%	**60% in banks**
Stratify: sector by size	10 sectors × 2 sizes	**20 cells**
Banks sector, its cells together	held at benchmark weight	**15%**

IN THE LESSON

- Harrow screens its five hundred stocks and holds the top fifty, equal weighted.
- Each stock gets one hundred per cent divided by fifty, two per cent.
- A stock with an alpha of four per cent and one with point five per cent both get the same two per cent, so the size of alpha is thrown away.
- Worse, suppose thirty of the fifty are banks, because one signal favours them.
- That is thirty times two, sixty per cent in banks, against fifteen per cent in the benchmark: the bets have clustered.
- Stratification fixes the clustering.
- With ten sectors and two size groups there are twenty cells, and the banks sector, its cells together, is held at its benchmark weight of fifteen per cent.

30 Example 15: linear against quadratic

Two bank stocks: α 3% each, residual risk 25%, λ = 8. LP limit: 5% a stock

LP: alpha pushed to the limit	ignores correlation	**5% each**
QP, correlation 0	0.03 ÷ (2×8×0.0625)	**3.00% each**
QP, correlation 0.9	0.03 ÷ (1.0 × 1.9)	**1.58% each**
QP, alpha mistyped as 12%	0.12 ÷ 1.9	**6.32% each**

IN THE LESSON

- Harrow likes two bank stocks equally: each has an alpha of three per cent and residual risk of twenty five per cent, lambda is eight, and the linear program allows at most five per cent in a stock.
- The linear program simply pushes both to the limit, five per cent each, whatever their correlation.
- The quadratic program weighs risk.
- If their residual returns are uncorrelated, each gets point zero three divided by two times eight times point zero six two five, which is three per cent.
- If the correlation is point nine, the denominator is multiplied by one point nine, and each gets one point five eight per cent, because the two bets are really one.
- But feed it an alpha mistyped as twelve per cent, and each weight jumps to six point three two per cent.

31 Dispersion and its causes

- Dispersion: return gaps among accounts on one strategy
- Not tracking error against a benchmark
- Causes: starting holdings, inception dates, cash flows
- Client restrictions; trade timing and execution
- Separate optimisations react to small input changes

IN THE LESSON

- Our last topic is dispersion.
- Harrow runs many separate client accounts to one model portfolio.
- Dispersion is the difference in returns among portfolios that follow the same strategy.
- It is not tracking error against a benchmark.
- What causes it?
- Different starting holdings and inception dates, since an account opened last year bought at different prices, and different cash flows, as clients add or withdraw money.
- Then client specific restrictions, such as a ban on tobacco stocks, and different trade timing and execution.
- And separate optimisations, which can react differently to small changes in their inputs.

32 Keeping dispersion under control

- Manage every account toward one model portfolio
- Trade accounts together
- Rebalance more often
- Standardise guidelines
- Costs: more trading; restrictions cannot always go

IN THE LESSON

- How is dispersion kept under control?
- First, manage every account toward one common model portfolio.
- Second, trade accounts together, so they get the same prices at the same time.
- Third, rebalance more often, so that differences do not build up.
- Fourth, standardise guidelines, so accounts differ only where they must.
- But each step has a cost.
- More trading costs more, and a client's restrictions cannot always be removed, so some dispersion is unavoidable.
- The aim is that the dispersion that remains comes from genuine client differences, not from careless implementation.

33 Example 16: how far apart two accounts drift

Two accounts track the model with independent tracking errors of 1.5% each

Each account vs the model	tracking error s	**1.5%**
Gap between the two	√(1.5^2 + 1.5^2)	**2.12%**
Shortcut	1.5 × √2 = 1.5 × 1.414	**2.12%**
After trading together	s = 0.5: 0.5 × √2	**0.71%**

Dispersion = s × √2: more than either account's tracking error s

IN THE LESSON

- Two Harrow accounts each track the model portfolio with a tracking error of one and a half per cent, and their deviations are independent.
- Each account's tracking error against the model, written s, is one and a half per cent.
- The gap between the two accounts has a standard deviation equal to the square root of one point five squared plus one point five squared, two point one two per cent.
- The shortcut is s times the square root of two, which is about one point four one four times s, giving the same two point one two.
- If trading together cut s to point five, dispersion would fall to point seven one per cent.
- The two accounts differ by more than either does from the model.

34 Try three quick questions

1. IC 0.05, volatility 20%, score +1.5. Scaled alpha?
   Answer: 20% × 0.05 × 1.5 = 1.5%.
2. IR 0.6, target active risk 3%. Risk aversion λ?
   Answer: 0.6 ÷ (2 × 0.03) = 10.
3. α 1.8%, λ = 4, cov 0.001. PC 0.7%, SC 0.5%. Trade?
   Answer: MCVA = 1.8 − 0.8 = 1.0% > 0.7%: buy.

IN THE LESSON

- Pause the video, think, then listen to the answers.
- Question one.
- A stock has an I C of point zero five, a volatility of twenty per cent and a score of plus one and a half.
- What is its scaled alpha?
- Twenty per cent times point zero five times one and a half is one and a half per cent.
- Question two.
- A manager has an information ratio of point six and a target active risk of three per cent.
- What risk aversion should it use?
- Lambda is point six divided by two times point zero three, which is ten.
- Question three.
- An asset has an alpha of one point eight per cent, lambda is four, and its covariance with the active portfolio is point zero zero one.
- Amortised costs are point seven to buy and point five to sell.
- Trade?
- M C V A is one point eight minus point eight, one per cent, above point seven, so buy.