Arbitrage pricing theory and factor models

Portfolio Management

1 A story: the recipe for a return

- A return is like a cake with a recipe
- Ingredients: factors like growth, size
- Each share mixes them differently
- Same recipe, same expected return

IN THE LESSON

- Imagine you are asked why a cake tastes the way it does.
- You could say it is sweet, or you could give the recipe: so much sugar, so much butter, so much lemon.
- The recipe explains the taste, and two cakes with the same recipe must taste the same.
- Arbitrage pricing theory treats a return like a cake.
- A share does not earn its return from one single force.
- It earns it from a handful of common ingredients, such as economic growth, inflation or company size, each mixed in a different amount.
- Meet Mira Okafor, an analyst at the made up firm Cedar Ridge Capital.
- She wants to price shares by their recipes.

2 The three assumptions behind APT

- 1. Returns come from a factor model
- 2. Well diversified portfolios exist
- 3. No arbitrage opportunity persists
- No market portfolio needed

IN THE LESSON

- So what must be true for the recipe idea to price assets?
- Arbitrage pricing theory makes three assumptions.
- First, returns are generated by a factor model: a few systematic factors drive them, and the rest is specific to each asset.
- Second, investors can form well diversified portfolios, in which the asset specific noise cancels out, so only factor risk is left.
- Third, no arbitrage opportunities persist among such portfolios.
- If a free profit appeared, traders would pounce on it and push prices until it vanished.
- Notice what is missing.
- The theory does not say investors must have a particular taste for risk, and it does not say they all hold the market portfolio.

3 The APT pricing equation

$E(R) = Rf + \beta1 \times \lambda1 + \beta2 \times \lambda2 + \ldots$

Risk-free rate	Earned with no factor risk
Factor sensitivity	How strongly asset responds
Factor risk premium	Reward per unit of factor risk
Expected return	Sum of the rows

IN THE LESSON

- Out of those assumptions comes one equation.
- The expected return on an asset equals the risk free rate, plus, for each factor, the asset's sensitivity to that factor times the factor's risk premium.
- Sensitivity says how strongly the asset responds to the factor: one means it moves one for one, zero means it ignores it.
- The risk premium is the extra return investors earn per unit of that factor's risk.
- With one factor, the equation is a straight line, and with several factors you simply add one more term for each.
- The risk free rate is what you would earn if every sensitivity were zero.
- It is the starting point of the recipe, and the factor terms are the flavour added on top.

4 APT against the CAPM

Both make return a straight line in risk

Feature	Where it stands
CAPM: factors	One: the market
APT: factors	One or several
CAPM needs	Market portfolio
APT needs	No arbitrage only

IN THE LESSON

- Let us put the two models side by side, because the exam loves to contrast them.
- The capital asset pricing model has one factor, the market, and to get there it assumes investors share the same expectations and all hold the market portfolio.
- Arbitrage pricing theory allows one factor or many, and needs only a factor model, diversification and no arbitrage.
- That makes it more flexible and its assumptions lighter.
- But there is a price.
- The theory does not tell you what the factors are, or how many there should be.
- Someone must supply them, and that is the job of the multifactor models you meet later in this lesson.

5 How APT gets to a price

From assumptions to a price

Few common factors → Specific risk fades → No free profit → Linear pricing

IN THE LESSON

- Here is the chain of reasoning in four boxes, read from left to right.
- The first box says returns follow a few common factors.
- The second says that in a well diversified portfolio the specific risk of each asset fades away, because gains and losses on individual companies cancel out.
- What remains is only the factor risk.
- The third box says that if two such portfolios have the same factor risk but different expected returns, a trader could earn a free profit, and so the prices would move until the profit was gone.
- The result in the last box is a linear price: expected return depends on sensitivities and premiums, and nothing else.

6 What is an arbitrage opportunity?

- Arbitrage: profit with no risk and no net investment
- Fund the purchase by selling something else short
- APT: among diversified portfolios, it cannot last
- Traders pile in until the gap closes

IN THE LESSON

- Now the key word.
- An arbitrage opportunity is a way to earn a profit that has no risk and needs no net investment.
- Think of a market stall selling an apple for one coin while the stall next door buys apples for two.
- You buy at the first, sell at the second, and keep the difference, with nothing of your own at stake.
- In investing you fund the purchase by selling something else short, so the net money you put in is zero.
- Arbitrage pricing theory says that, among well diversified portfolios, such chances cannot last.
- Every trader would rush in at once, buying the cheap portfolio and selling the dear one, until the gap closed.

7 One portfolio sits above the line

Chart: Expected return (%) vs Factor sensitivity (0 to 2)

- Birch 7%
- Mix 10%
- Cedar 11%
- Alder 13%

IN THE LESSON

- Here is the picture for Mira's case.
- The horizontal axis is the factor sensitivity, from zero to two.
- The vertical axis is expected return in per cent.
- Take two well diversified portfolios, Alder with a sensitivity of one and a half and a return of thirteen, and Birch with a sensitivity of a half and a return of seven.
- Draw the straight line through them.
- It starts at four per cent, the risk free rate, and rises six points per unit of sensitivity.
- Now a half and half mix of Alder and Birch has a sensitivity of one and sits on the line at ten per cent.
- Cedar also has a sensitivity of one, yet offers eleven.
- Cedar floats above the line, and that is the opportunity.

8 Building a copy with the right sensitivity

w × β1 + (1 − w) × β2 = β target

Weight in first	Share put in the first portfolio
Factor sensitivity	Sensitivities of the two blended
Target sensitivity	Sensitivity of portfolio tested

IN THE LESSON

- To test for arbitrage, we build a copy.
- We mix two portfolios so that the blend has the same factor sensitivity as the third one we are testing.
- Let w be the weight in the first portfolio, so the weight in the second is one minus w.
- The blend's sensitivity is the weighted average of the two sensitivities, and we set it equal to the target.
- Solve for w.
- Because the weights add to one, no money is needed beyond the total amount.
- If the weight in one comes out negative, that simply means selling it short.
- The blend's expected return is then the same weighted average of the two expected returns, and that return is what the copy should earn.

9 Example 1: the replicating weights

Alder 1.5, Birch 0.5. Copy Cedar's sensitivity of 1.0

Target sensitivity	wanted	**1.0**
Equation	w×1.5 + (1−w)×0.5 = 1.0	**solve**
Weight in Alder	0.5 ÷ 1.0	**0.5**
Weight in Birch	1 − 0.5	**0.5**
Copy's return	0.5×13% + 0.5×7%	**10.0%**

IN THE LESSON

- Example one.
- Mira wants a copy of Cedar, whose sensitivity is one.
- She blends Alder, at one and a half, with Birch, at a half.
- The equation is w times one point five, plus one minus w times a half, equals one.
- Expanding it gives a half plus w equals one, so w is a half.
- She puts half in Alder and half in Birch.
- Check the sensitivity: half of one and a half is three quarters, half of a half is a quarter, and together that is one.
- The copy's expected return is half of thirteen plus half of seven, which is ten per cent.
- So the copy of Cedar's risk should pay ten, but Cedar itself offers eleven.

10 Example 2: trading the gap

Buy Cedar 10m; short Alder 5m and Birch 5m

Buy Cedar	10m × 11%	**+1,100,000**
Short Alder	5m × 13%	**−650,000**
Short Birch	5m × 7%	**−350,000**
Net investment	10m − 5m − 5m	**0**
Riskless profit	1.10m − 0.65m − 0.35m	**100,000**

IN THE LESSON

- Example two.
- Cedar offers eleven, the copy offers ten, and both carry the same factor risk.
- So Mira buys Cedar and sells the copy short.
- Take ten million in Cedar.
- To fund it she sells five million of Alder and five million of Birch short, so her net investment is zero.
- Her factor exposure also nets to zero: Cedar adds one, the short copy takes away one.
- Now the returns.
- Cedar earns eleven per cent of ten million, one million one hundred thousand.
- She owes thirteen per cent of five million on Alder, six hundred fifty thousand, and seven per cent of five million on Birch, three hundred fifty thousand.
- The profit is one hundred thousand, riskless.

11 Example 3: a portfolio priced too high

Dune: sensitivity 0.8, offers 8.5%. Test it

Return on the line	4 + 6 × 0.8	**8.8%**
Weight in Alder	(0.8 − 0.5) ÷ 1.0	**0.3**
Copy's return	0.3×13% + 0.7×7%	**8.8%**
Dune offers	below its copy by 0.3	**8.5%**
Short Dune, buy copy	0.3% of 10m	**30,000**

IN THE LESSON

- Example three shows the other direction.
- Dune has a sensitivity of point eight and offers eight and a half per cent.
- Is that an arbitrage?
- Start with the line: four plus six times point eight is eight point eight, so Dune should offer eight point eight.
- To confirm it with trades, build a copy.
- Solving for w gives point three in Alder and point seven in Birch.
- The copy returns point three times thirteen plus point seven times seven, eight point eight.
- Dune offers three tenths of a point less than its copy, so Dune is overpriced.
- Mira buys the copy and sells Dune short.
- On ten million, the gain is thirty thousand.

12 The trap: a gap is not always an arbitrage

- Different returns alone do not prove arbitrage
- Higher sensitivity should earn a higher return
- Step 1: build a copy with the same sensitivity
- Step 2: compare returns; trade only if they differ

IN THE LESSON

- Here is the trap that costs marks.
- Many candidates see two portfolios with different expected returns and shout arbitrage.
- That is wrong.
- Alder offers thirteen and Cedar eleven, yet no free profit exists, because Alder carries one and a half units of factor risk while Cedar carries one.
- A higher sensitivity should be rewarded with a higher return.
- Arbitrage needs two positions with the same factor exposure priced differently, or a blend that copies one of them.
- So the test has two steps.
- First, match the sensitivity, using a copy.
- Second, compare the expected returns.
- Only if the copy and the original differ is there a gain to harvest.

13 A story: the bill with several charges

- Basic charge: the risk free rate
- Each extra: sensitivity × factor risk premium
- A negative sensitivity is a refund, a hedge
- Add all the lines for the expected return

IN THE LESSON

- Think of a restaurant bill.
- There is a basic charge for the table, then a line for each extra: drinks, dessert, the special.
- Each extra costs the number of items you ordered times the price per item.
- A factor model bill works the same way.
- The basic charge is the risk free rate.
- Each extra is a factor, and its cost is your sensitivity, the quantity ordered, times the factor risk premium, the price per item.
- Sensitivity can be negative, like returning an item for a refund.
- A share that does well when a factor does badly acts as a hedge, and investors accept a lower expected return for it.
- Back to Mira, who now needs to fill in each line of the bill.

14 Where a factor risk premium comes from

Factor premium = E(R of factor portfolio) − Rf

Factor risk premium	Reward per unit of one factor
Factor portfolio return	Portfolio: sensitivity 1, others 0
Risk-free rate	Subtract it once only

IN THE LESSON

- Where does the price per item come from?
- Picture a pure factor portfolio: a portfolio built to have a sensitivity of one to a single factor and zero to every other factor.
- Its expected return, less the risk free rate, is that factor's risk premium.
- So the premium is an excess return, never the total return.
- This matters because the price you plug into the equation is the reward above the risk free rate, which is already counted once at the start of the bill.
- If you used the total return of the factor portfolio you would count the risk free rate again.
- Remember that one line: premium equals factor portfolio return minus the risk free rate.

15 Example 4: finding the premiums

Rf = 3%. Pure factor portfolios' returns are given

Market premium	8% − 3%	**5.0%**
Size premium	5% − 3%	**2.0%**
Value premium	6% − 3%	**3.0%**

IN THE LESSON

- Example four.
- The risk free rate is three per cent.
- Mira has pure factor portfolios for three factors.
- The market factor portfolio, with a sensitivity of one to the market and zero to the others, returns eight per cent.
- Subtract three and the market premium is five.
- The size factor portfolio returns five, so the size premium is two.
- The value factor portfolio returns six, so the value premium is three.
- These premiums are what she will plug into the equation.
- Notice they are all measured above the risk free rate.
- Had she multiplied by eight, five or six instead, every answer would be inflated by the risk free rate once for every factor.

16 Example 5: expected return of Maple

Maple: market 1.1, size 0.5, value −0.4

Risk-free rate		**3.0%**
Market	1.1 × 5.0%	**+5.5%**
Size	0.5 × 2.0%	**+1.0%**
Value	−0.4 × 3.0%	**−1.2%**
Expected return	3.0 + 5.5 + 1.0 − 1.2	**8.3%**

IN THE LESSON

- Example five.
- Maple is a share with a sensitivity of one point one to the market, a half to size, and minus point four to value.
- The risk free rate is three, and the premiums are the ones we just found.
- Line one is the risk free rate, three.
- The market line is one point one times five, five and a half.
- The size line is a half times two, one.
- The value line is minus point four times three, minus one point two.
- Add them: three plus five and a half plus one minus one point two is eight point three per cent.
- The negative value sensitivity pulls the return down, because Maple behaves a little like a growth share and so offers protection when value shares suffer.

17 Example 6: Quill and the premium trap

Rf 2.5%. Growth premium 3.5%, inflation premium 1.5%

Growth	1.3 × 3.5%	**+4.55%**
Inflation	−0.6 × 1.5%	**−0.90%**
Right answer	2.5 + 4.55 − 0.90	**6.15%**
Wrong: total returns	2.5 + 1.3×6 − 0.6×4	**7.90%**

IN THE LESSON

- Example six shows the trap.
- The risk free rate is two and a half per cent.
- Quill has a sensitivity of one point three to a growth factor and minus point six to an inflation factor.
- The growth factor portfolio returns six, so its premium is three and a half.
- The inflation factor portfolio returns four, so its premium is one and a half.
- The right answer: two and a half, plus one point three times three and a half, four point five five, plus minus point six times one and a half, minus point nine.
- Total six point one five.
- The wrong route multiplies by six and four instead, and gives seven point nine.
- That is almost two points too high.

18 Mistakes to avoid so far

- APT does not need the market portfolio
- A return gap alone is not arbitrage
- Use premiums, not total factor returns
- Keep the sign of a negative sensitivity
- Match sensitivity first, then compare returns

IN THE LESSON

- Let us stop and collect the common mistakes.
- One: believing arbitrage pricing theory needs everyone to hold the market portfolio.
- It does not, and that is its main advantage over the capital asset pricing model.
- Two: calling any gap in returns an arbitrage.
- Different sensitivities can fairly earn different returns.
- Three: multiplying sensitivities by the total returns of factor portfolios instead of by premiums.
- Four: dropping the minus sign on a negative sensitivity, which flips the effect on the answer.
- Five: forgetting that a replicating blend must have the same sensitivity as the target before you compare returns.
- A slow, steady check of each step prevents all five.

19 A story: three ways to find the ingredients

- Macroeconomic: factors are surprises in economic news
- Fundamental: factors are company attributes
- Statistical: factors are patterns found in the data
- All three: factor exposure plus a specific part

IN THE LESSON

- APT told us there are factors but did not name them.
- So where do analysts find them?
- Imagine three detectives trying to explain why share prices move.
- The first looks at the news: growth came in above forecast, inflation surprised on the upside.
- That is the macroeconomic model.
- The second looks at the companies themselves: how cheap, how big, how strong the recent run.
- That is the fundamental model.
- The third knows nothing about the world and simply studies the data, searching for patterns that explain how returns move together.
- That is the statistical model.
- All three give returns the same shape: exposure to a few factors, plus a specific part.

20 From clear meaning to best fit

- Macroeconomic
- Fundamental
- Statistical

IN THE LESSON

- Here are the three on one line.
- On the left, the end with clear meaning.
- The macroeconomic model sits there, because its factors are things you can name and read in a newspaper, such as growth or inflation surprises.
- In the middle sits the fundamental model, whose factors are company characteristics that investors know well, such as value, size and momentum, and which are easy to interpret.
- On the right is the statistical model.
- It fits past returns very well, because it is built from them, but its factors have no obvious economic meaning.
- You can see that something drives returns, but you cannot easily say what it is.

21 The macroeconomic factor model

R = a + b1 × F1 + b2 × F2 + ε

Expected return	Return if all surprises are zero
Sensitivity	Estimated by regression over time
Factor surprise	Actual minus expected value
Specific error	Part not explained by factors

IN THE LESSON

- Now the equation for the macroeconomic model.
- A stock's return equals its expected return, plus its sensitivity to each factor times that factor's surprise, plus an error that is specific to the stock.
- The word surprise is the key.
- The factor is not the level of growth, it is growth minus what was expected.
- If growth was forecast to be two and a half and came in at three point one, the surprise is point six.
- Expected growth is already in the expected return and in prices, so only the news moves the return.
- The sensitivities are estimated by a time series regression of the stock's returns on past surprises.

22 Example 7: Reed's return in a macro model

Reed: a 8.0%, GDP 1.2, inflation −0.5, error +0.4

Growth surprise	3.1 − 2.5	**0.6**
Inflation surprise	2.8 − 2.0	**0.8**
Growth effect	1.2 × 0.6	**+0.72**
Inflation effect	−0.5 × 0.8	**−0.40**
Return	8.0 + 0.72 − 0.40 + 0.4	**8.72%**

IN THE LESSON

- Example seven.
- Reed has an expected return of eight per cent, a sensitivity of one point two to growth, and minus a half to inflation.
- The year began with a growth forecast of two and a half and an inflation forecast of two.
- Growth came in at three point one, a surprise of point six.
- Inflation came in at two point eight, a surprise of point eight.
- The growth effect is one point two times point six, point seven two.
- The inflation effect is minus a half times point eight, minus point four.
- Add a specific error of plus point four and the return is eight point seven two.
- Using the level of growth instead of the surprise would give an absurd result.

23 Example 8: Fenn in a fundamental model

Fenn: book-to-price 0.9; universe mean 0.6, SD 0.3

Value sensitivity	(0.9 − 0.6) ÷ 0.3	**1.0**
Size sensitivity	large firm	**−0.5**
Value effect	1.0 × 1.2	**+1.20**
Size effect	−0.5 × 0.4	**−0.20**
Return	6.0 + 1.2 − 0.2 + 0.3	**7.3%**

IN THE LESSON

- Example eight turns things round.
- In a fundamental model the sensitivities are observed, not estimated.
- They are the standardised attributes of the company.
- Fenn has a book to price ratio of point nine, the universe averages point six with a standard deviation of point three.
- Its value sensitivity is point nine minus point six over point three, which is one.
- Say its size sensitivity is minus a half, because it is a large company.
- Each period, the factor returns are estimated by a cross sectional regression across all the stocks.
- This period value returned one point two and size point four.
- With an expected return of six and an error of point three, Fenn returns seven point three.

24 The statistical model: factors from the data

- Factors are extracted from return covariances
- Factor analysis or principal components analysis
- Strength: very good fit to the data
- Weakness: factors have no clear economic meaning

IN THE LESSON

- Now the third kind.
- A statistical factor model takes a large set of past returns and extracts factors from the way they vary together, their covariances.
- Two techniques do this.
- Factor analysis finds the factors that best explain the covariances between the returns.
- Principal components analysis finds combinations of the returns that explain as much of the total variance as possible.
- Neither starts with any economic idea.
- The big advantage is a close fit to the data, and no need to guess what the factors should be.
- The disadvantage is that the factors can be hard to interpret, and may change when the data changes.
- A factor with no meaning is hard to trust or to explain to a client.

25 What is observed and what is estimated

Which quantity is known and which is estimated

Model	Factors | sensitivities
Macro	Observed | estimated
Fundamental	Estimated | observed
Statistical	Estimated | estimated

IN THE LESSON

- This table is the one to memorise, because the exam loves to swap the roles.
- In the macroeconomic model, the factors are observed, since the surprises are measured from published data.
- The sensitivities are estimated, by regressing each stock's returns on those surprises through time.
- In the fundamental model it is the other way round.
- The sensitivities are observed, because they are the standardised attributes of each company.
- The factor returns are estimated, by a cross sectional regression of all stocks' returns on those attributes in each period.
- In a statistical model neither is observed.
- Both the factors and the sensitivities are estimated from the returns themselves.

26 Mistakes with factor models

- Macro model: surprises, never the level of the variable
- Fundamental: sensitivities are standardised attributes
- Statistical: factors from covariances, no clear meaning
- Read the clue words to name the model

IN THE LESSON

- Two last traps before the recap.
- First, the macroeconomic model uses surprises, not levels.
- If growth is high but exactly as expected, the surprise is zero and the factor adds nothing to the return.
- Mixing up the level and the surprise is the most common error in vignettes.
- Second, reversing what is observed and what is estimated.
- If a question says the sensitivities are standardised attributes, you are in a fundamental model.
- If the factors are economic surprises, you are in a macroeconomic one.
- If nothing economic is named and factors come from covariances, it is statistical.
- Read the clue words slowly, and the model announces itself.

27 Recap: the lesson in four lines

Rf + Σ β × λ
Premium = factor return − Rf
Copy, then compare

Risk-free rate	Start of every APT return
Factor sensitivity	Observed or estimated by model
Factor risk premium	Factor portfolio less Rf
Target sensitivity	Copy matches this first
Return gap	Trade it only if present

IN THE LESSON

- Let us pull the lesson together.
- Arbitrage pricing theory says expected return is the risk free rate plus the sum of each sensitivity times its factor risk premium.
- It needs a factor model, diversification and no arbitrage, but not the market portfolio, and it does not name the factors.
- An arbitrage is a riskless profit with no net investment.
- To test for one, build a blend with the same sensitivity as the target and compare expected returns.
- A premium is a factor portfolio's return minus the risk free rate.
- Macroeconomic models use surprises, fundamental models use standardised attributes, and statistical models extract factors from covariances.
- Keep these four lines in mind.

28 Try three quick questions

1. Rf 3%; sensitivities 1.0 and 0.5; premiums 4% and 2%. E(R)?
   Answer: 3 + 1.0×4 + 0.5×2 = 8.0%.
2. Same sensitivities; returns 9% and 10%; 5m. Profit?
   Answer: Buy 10%, short 9%: 1% × 5m = 50,000.
3. Factors observed, sensitivities from time series?
   Answer: The macroeconomic model.

IN THE LESSON

- Now try three quick questions yourself.
- Pause after each one.
- Question one.
- The risk free rate is three per cent.
- An asset has a sensitivity of one to a factor with a premium of four per cent, and a half to a second factor with a premium of two per cent.
- What is its expected return?
- Three plus four plus one is eight per cent.
- Question two.
- Two well diversified portfolios have the same factor sensitivities.
- One is expected to return nine, the other ten.
- With five million, what is the riskless profit?
- Buy the ten, sell the nine short, and one per cent of five million is fifty thousand.
- Question three.
- Which model has observed factors and sensitivities estimated by time series regression?
- The macroeconomic model.
- Well done, and keep practising.