CHARTA LEARNING · LESSON NOTES

Factor Theory

FRM Part II

1 A story: the asset that earns less

- Kestrel holds two assets with the same volatility
- Halden Industries: falls hard in recessions
- A shelter asset: rises in recessions
- Halden is expected to earn far more. Why?
- Answer: what matters is when an asset loses

IN THE LESSON

- Kestrel's trustees notice something odd.
- The fund holds two assets that are equally volatile: their returns swing by about the same amount.
- The first is shares in Halden Industries, a made up maker of machinery, which falls hard in recessions.
- The second is a shelter asset that tends to rise in recessions.
- Yet Halden is expected to earn far more.
- Why?
- What matters is not how much an asset swings, but when it loses.
- Losing money in bad times, when money is most precious, is what investors demand to be paid for.

2 Example 1: two levers move a price

A claim pays $110 expected in one year; risk-free rate 3%, risk premium 5%

Discount rate	3% + 5%	**8%**
Price today	110 ÷ 1.08	**$101.85**
Premium up to 7%	110 ÷ 1.10	**$100.00**
Cash flow down to $105	105 ÷ 1.08	**$97.22**

Price = PV of expected cash flows. Two levers: cash flows, discount rate

IN THE LESSON

- A claim is expected to pay one hundred and ten dollars in one year.
- The risk free rate is three per cent and the risk premium five per cent.
- A price is the present value of expected cash flows, discounted at the risk free rate plus a risk premium.
- So the discount rate is eight per cent, and the price today is one hundred and ten divided by one point zero eight, one hundred and one dollars eighty five.
- With a premium of seven per cent instead, the rate is ten per cent and the price falls to one hundred dollars.
- If instead the expected cash flow falls to one hundred and five dollars, at eight per cent the price is ninety seven dollars twenty two.
- So a price has two levers: the expected cash flows, and the discount rate.

3 Which risks earn a premium?

- Premium paid only for systematic factor exposure
- Systematic: growth, inflation, real rates, liquidity
- Idiosyncratic: one firm's own news
- Idiosyncratic risk diversifies away: no premium
- In a large portfolio, firm news cancels out

IN THE LESSON

- A premium is paid only for exposure to systematic factors.
- These are broad sources of risk that hit many assets at once, such as economic growth, inflation, real interest rates and market liquidity.
- Idiosyncratic risk is different.
- It is risk specific to one firm.
- Idiosyncratic risk earns no premium, because it diversifies away.
- When Kestrel holds hundreds of assets, one firm's bad news is offset by another firm's good news, so in a large portfolio firm news cancels out.
- Risk that can be removed for free is not paid for.

4 Bad times decide the premium

- Bad times: an extra unit of wealth is most valuable
- Losing value in bad times: a positive premium
- Paying off in bad times: works as insurance
- Insurance assets can earn less than the risk-free rate
- Halden and shelter: same volatility, opposite premiums

IN THE LESSON

- Why do systematic factors earn a premium?
- Think of bad times, such as a recession.
- In bad times an extra unit of wealth is most valuable.
- A factor earns a positive premium when the assets exposed to it lose value in bad times, just when investors can least afford it, so investors must be paid to hold them.
- Assets that pay off in bad times work like insurance.
- Investors value that protection so much that they accept a lower return, and such assets can even earn less than the risk free rate.
- This solves the trustees' puzzle.
- Halden and the shelter asset have the same volatility but opposite premiums, because one loses in bad times and the other gains.

5 Three explanations, and the common error

- Rational: compensation for bad-times risk
- Behavioural: investor biases
- Structural: institutional constraints
- Common error: rewarding total volatility
- Premiums follow covariance with bad times

IN THE LESSON

- Premiums have three families of explanation.
- Rational explanations say the premium is compensation for bad times risk, exactly the story we just told.
- Behavioural explanations say investor biases, such as overreacting to news, push prices away from fair value.
- Structural explanations point to institutional constraints, for example rules that stop some investors from borrowing or from holding certain assets.
- The common error on the exam is to reward total volatility, as if any swing in returns must be paid for.
- It is not.
- Premiums follow covariance with bad times, meaning how an asset moves when times are bad.

6 The capital asset pricing model

ER = rf + β × (Rm − rf)
β = cov(i, m) ÷ σm^{2} = ρ × σi ÷ σm

Risk-free rate	rf: return on the risk-free asset
Beta	Exposure to the market factor
Market risk premium	Rm − rf: market return less rf
Correlation	ρ: correlation with the market

IN THE LESSON

- The simplest factor model is the capital asset pricing model, the C A P M.
- Its single priced factor is the return on the market portfolio.
- The expected return on asset i, written E R, equals r f, the risk free rate, plus beta times the market risk premium.
- The market risk premium is R m, the expected market return, minus the risk free rate.
- Beta, the asset's exposure to the market, is computed two ways on the second line: the covariance of the asset with the market, divided by the market variance, or rho, the correlation with the market, times the asset's volatility, sigma i, divided by the market's volatility, sigma m.

7 What the CAPM assumes

- Mean-variance investors with identical expectations
- One period
- Unlimited borrowing and lending at one risk-free rate
- No taxes or transaction costs
- Price-takers who can trade every asset

IN THE LESSON

- The model rests on strong assumptions, and the exam expects you to list them.
- Investors are mean variance investors: they care only about expected return and the variance of return.
- They all share identical expectations about every asset.
- There is one period, so nobody plans beyond it.
- Everyone can borrow and lend as much as they like at one risk free rate.
- There are no taxes and no transaction costs.
- And investors are price takers, too small to move prices, who can trade every asset.
- Under these conditions the only reward is the market risk premium times the asset's beta.
- The model's weaknesses come from this list.

8 Example 2: Halden's beta and expected return

Halden: ρ = 0.6, σ = 30%. Market: σ = 20%, Rm = 9%. Risk-free rate 3%

Covariance with the market	0.6 × 30% × 20%	**0.036**
Beta	0.036 ÷ 0.20^{2}	**0.90**
Expected return	3% + 0.9 × (9% − 3%)	**8.4%**
Wrong: σi ÷ σm as beta	30% ÷ 20%	**1.5**
Wrong: beta times Rm	3% + 0.9 × 9%	**11.1%**

IN THE LESSON

- Halden's correlation with the market is point six and its volatility thirty per cent.
- The market's volatility is twenty per cent and its expected return nine, with a risk free rate of three.
- The covariance is the correlation times the two volatilities, point zero three six.
- Divided by the market variance, point two squared, beta is point nine.
- The expected return is three per cent plus point nine times the premium of nine minus three, eight point four per cent.
- Now the traps.
- Dividing the volatilities without the correlation gives a beta of one point five, which is wrong.
- And applying beta to the market return instead of the premium gives eleven point one per cent, also wrong.

9 Splitting risk: rewarded and unrewarded

σi^{2} = β^{2} × σm^{2} + σe^{2}
Reward = β × (Rm − rf), from the first part only

Systematic variance	First part: systematic, rewarded
Residual variance	Second part: residual, unrewarded
Market risk premium	Reward: beta × premium

IN THE LESSON

- The C A P M also tells us how total risk splits.
- The total variance of an asset, sigma i squared, equals beta squared times the market variance, plus the residual variance, sigma e squared.
- The first part is systematic: it comes from the market and cannot be diversified away, so it is rewarded.
- The second part is residual, or idiosyncratic: it comes from the firm's own news and can be diversified away, so it is unrewarded.
- The second line says the reward is beta times the market risk premium, and it comes from the first part only.
- Two assets with the same total variance can therefore earn very different returns, if their variance splits differently between the two parts.

10 Example 3: more volatile, yet a lower return

Halden: σ 30%, β 0.9. Vantor: σ 40%, ρ 0.15. Market σ 20%, rf 3%, premium 6%

Halden: systematic variance	0.9^{2} × 0.04	**0.0324**
Halden: residual variance	0.09 − 0.0324	**0.0576**
Vantor: beta	0.15 × 40% ÷ 20%	**0.30**
Vantor: systematic variance	0.3^{2} × 0.04; total 0.4^{2}	**0.0036 of 0.16**
Vantor: expected return	3% + 0.3 × 6%	**4.8%**

Vantor is more volatile than Halden, yet expects 4.8% against 8.4%

IN THE LESSON

- Halden has a volatility of thirty per cent and a beta of point nine.
- Vantor Biotech, a made up drug developer, has a volatility of forty per cent and a correlation with the market of only point one five.
- The market's volatility is twenty per cent, the risk free rate three per cent and the market premium six per cent.
- Halden's systematic variance is point nine squared times point zero four, point zero three two four.
- Its total variance is point three squared, point zero nine, so its residual variance is point zero five seven six.
- Vantor's beta is point one five times forty, over twenty, point three.
- Its systematic variance is point three squared times point zero four, only point zero zero three six, out of a total of point one six.
- So its expected return is three per cent plus point three times six, four point eight per cent.
- Vantor is more volatile than Halden, yet expects four point eight per cent against eight point four.

11 Everyone holds the market: the capital market line

ERp = rf + [(Rm − rf) ÷ σm] × σp
Two-fund separation: market portfolio + risk-free

Risk-free rate	rf: return on the risk-free asset
Slope of the line	Reward per unit of market risk
Chosen volatility	Volatility the investor picks

IN THE LESSON

- Now the implications.
- If every investor shares the C A P M assumptions, all hold the same risky portfolio, the market portfolio.
- They differ only in how much they lend or borrow at the risk free rate.
- The choices trace a straight line, the capital market line.
- Read the formula.
- The expected return on the chosen mix, E R p, equals the risk free rate plus the slope, which is the market risk premium divided by sigma m, times sigma p, the volatility of the chosen mix.
- The second line names the result: two fund separation, because every choice is a mix of just the market portfolio and the risk free asset.

12 Example 4: two members on the line

Market: Rm 9%, σm 20%. Risk-free 3%. Targets: σp 10% and 30%

Slope of the line	(9% − 3%) ÷ 20%	**0.30**
Cautious: weight in market	10% ÷ 20%	**50%**
Cautious: expected return	3% + 0.30 × 10%	**6.0%**
Bold: weight in market	30% ÷ 20%; borrow 50%	**150%**
Bold: expected return	3% + 0.30 × 30%	**12.0%**

IN THE LESSON

- The market's expected return is nine per cent and its volatility twenty per cent, with a risk free rate of three per cent.
- Kestrel's members target volatilities of ten and thirty per cent.
- The slope of the line is nine minus three, over twenty, which is point three.
- The cautious member wants ten per cent volatility: ten over twenty, so half the money in the market and half lent at the risk free rate.
- The expected return is three per cent plus point three times ten, six per cent.
- The bold member wants thirty per cent, which needs one hundred and fifty per cent in the market, borrowing fifty per cent.
- The expected return is three per cent plus point three times thirty, twelve per cent.

13 Example 5: alpha against the security market line

Alpha = forecast − [rf + β × (Rm − rf)]. Forecasts: Halden 10%, Vantor 4%

Halden: required return	3% + 0.9 × 6%	**8.4%**
Halden: alpha	10.0% − 8.4%	**+1.6%**
Vantor: required return	3% + 0.3 × 6%	**4.8%**
Vantor: alpha	4.0% − 4.8%	**−0.8%**

Positive alpha: above the line, underpriced. Negative: below, overpriced

IN THE LESSON

- Alpha equals the forecast return minus the return the C A P M requires.
- Kestrel forecasts ten per cent for Halden and four for Vantor.
- Individual assets lie on the security market line, which plots required return against beta.
- Halden's required return is three per cent plus point nine times six per cent, eight point four per cent.
- Its alpha is ten minus eight point four, plus one point six per cent.
- Vantor's required return is three per cent plus point three times six, four point eight per cent.
- Its alpha is four minus four point eight, minus point eight per cent.
- Positive alpha means the asset plots above the line and is underpriced.
- Negative alpha, below the line, means overpriced.

14 The security market line



IN THE LESSON

- Along the bottom is beta, and up the side is expected return.
- The straight line rises from left to right.
- The ticks mark betas of zero, one and two.
- The line starts at three per cent, the risk free rate, and reaches fifteen per cent at a beta of two, rising by the six per cent premium per unit of beta.
- This is the security market line.
- The market itself sits on the line at a beta of one and nine per cent.
- Halden, with a beta of point nine and a forecast of ten per cent, sits above the line by one point six per cent.
- Vantor, with a beta of point three and a forecast of four per cent, sits below it by point eight per cent.

15 Example 6: diversification removes unrewarded risk

Equal stakes in n stocks: each β 1 (3% + 1 × 6% = 9%), residual σ 40%, independent

1 stock	√(0.04 + 0.16)	**44.7%**
10 stocks	√(0.04 + 0.016)	**23.7%**
100 stocks	√(0.04 + 0.0016)	**20.4%**
Very many stocks	√0.04	**20.0%**

Expected return stays 9% at every size: only beta is paid

IN THE LESSON

- Kestrel builds equally weighted portfolios of n stocks, each with a beta of one, so an expected return of nine per cent, and a residual volatility of forty per cent, the residuals independent.
- The portfolio variance is the market variance, point zero four, plus the residual variance, point one six, divided by n.
- With one stock, volatility is the square root of point two, forty four point seven per cent.
- With ten stocks it is the square root of point zero five six, twenty three point seven per cent.
- With one hundred stocks it is twenty point four per cent.
- With very many stocks it settles at the market's twenty per cent.
- Residual risk is diversified away, but systematic risk stays.
- The expected return is nine per cent at every size, because only beta is paid.

16 Where the CAPM falls short

- Roll's critique: the true market portfolio is unobservable
- So the model cannot be tested cleanly
- The assumptions are unrealistic
- Evidence: the line is flatter than predicted
- Low-beta stocks earn more than the model says

IN THE LESSON

- The first is Roll's critique.
- The true market portfolio contains every risky asset, not just listed shares, so it is unobservable.
- Any index we use is only a stand in.
- So the model cannot be tested cleanly: a failed test may just mean the wrong proxy.
- Second, the assumptions are unrealistic.
- Investors do not share identical expectations, and they cannot borrow without limit at the risk free rate.
- Third, the evidence.
- The security market line is flatter than predicted, and low beta stocks have earned more than the model says.
- These gaps are one reason for using more than one factor, which is where we go next.

17 Multifactor models

ER = rf + Σ βk × λk

CAPM: the special case with one factor, the market

Risk-free rate	rf: return on the risk-free asset
Exposure to factor k	Exposure to k; can be negative
Premium of factor k	Risk premium of factor k

IN THE LESSON

- A multifactor model builds expected return from several priced factors: the risk free rate plus, for each factor, beta k, the asset's exposure to that factor, times lambda k, the factor's risk premium, all summed together.
- Notice that exposures can be negative.
- An asset that tends to do well when a factor does badly has a negative exposure, and that term then subtracts from its expected return.
- The second line makes the link to what we have done.
- The C A P M is the special case with just one factor, the market, where the exposure is beta and the premium is the market risk premium.

18 Which factors?

- Macroeconomic: growth, inflation
- Statistical: drawn out of return data
- Characteristic: size, value, momentum
- Like the CAPM: linear, pays only systematic exposure
- Idiosyncratic risk unpriced in both

IN THE LESSON

- Factors come in three kinds.
- Macroeconomic factors are surprises in the economy, such as growth and inflation.
- Statistical factors are drawn out of the return data itself.
- Characteristic based factors are built from features of firms: size, small firms against large; value, cheap shares against expensive ones; and momentum, recent winners against recent losers.
- What do multifactor models share with the C A P M?
- Both are linear, and both pay only for systematic exposure.
- In both, idiosyncratic risk is left unpriced, because it can be diversified away.

19 Example 7: Halden in a three-factor model

Premiums: market 6%, size 2%, value 3%. Halden's β: 0.9, 0.5, −0.4

Market	0.9 × 6%	**5.4%**
Size	0.5 × 2%	**1.0%**
Value: negative exposure	−0.4 × 3%	**−1.2%**
Expected return	3% + 5.4% + 1.0% − 1.2%	**8.2%**

IN THE LESSON

- Kestrel now uses a three factor model with assumed premiums of six per cent for the market, two per cent for size and three per cent for value.
- Halden's exposures are point nine to the market, point five to size and minus point four to value.
- The market term is point nine times six per cent, five point four per cent.
- The size term is point five times two per cent, one per cent.
- The value term is minus point four times three per cent, minus one point two per cent, so the negative exposure subtracts.
- The expected return is three per cent plus five point four plus one, minus one point two, which is eight point two per cent.

20 Example 8: an alpha that disappears

Small Value Fund: return 11%; β market 1.0, size 0.4, value 0.4

CAPM required return	3% + 1.0 × 6%	**9.0%**
CAPM alpha	11% − 9%	**+2.0%**
Size and value premiums	0.4 × 2% + 0.4 × 3%	**2.0%**
Multifactor required return	9% + 2%	**11.0%**
Multifactor alpha	11% − 11%	**0.0%**

The CAPM alpha was pay for omitted factors

IN THE LESSON

- The Small Value Fund, an external manager, has earned eleven per cent a year with a market beta of one, and exposures of point four to size and to value.
- Under the C A P M, its required return is three per cent plus one times six, nine per cent, so its alpha is eleven minus nine, plus two per cent.
- But the size and value premiums add point four times two, plus point four times three, which is two per cent.
- The multifactor required return is nine plus two, eleven per cent, and the alpha is eleven minus eleven, zero.
- The apparent skill was simply pay for omitted factors.

21 The CAPM and the APT side by side

Both are linear; both price only systematic risk

	CAPM	APT
Factors	One: the market	Several
Names its factors?	Yes	No: user chooses
Foundation	Equilibrium, mean-variance	No arbitrage, diversified
Idiosyncratic risk	Unpriced	Unpriced

IN THE LESSON

- Both models are linear, and both price only systematic risk.
- The table has a column for the C A P M and a column for the arbitrage pricing theory, the A P T, the classic multifactor model.
- On factors, the C A P M has one, the market, while the A P T has several.
- Does the model name its factors?
- The C A P M does.
- The A P T does not say which factors to use, so the user must choose them.
- On foundations, the C A P M is an equilibrium result for mean variance investors.
- The A P T rests on no arbitrage, applied to well diversified portfolios.
- And on idiosyncratic risk, both agree: it is unpriced.

22 One price for every payoff: the discount factor

- m: one random variable that prices every payoff
- m takes a value in each state of the world
- m is high in bad states: marginal utility is high
- A dollar in a bad state is worth more today
- Paying little when m is high: a premium

IN THE LESSON

- Now a more general tool that contains all these models.
- A stochastic discount factor, written m, is a single random variable that prices every payoff.
- Think of the future as a set of possible states of the world, such as a boom or a recession.
- The discount factor takes a value in each state.
- It is high in bad states, because that is when marginal utility is high, meaning an extra dollar adds the most to well being.
- So a dollar paid in a bad state is worth more today than a dollar paid in a good state.
- It follows that assets which pay little when m is high, that is, which pay little in bad times, must offer a premium to be held.

23 Pricing with m

p = E(m × x)
Rf = 1 ÷ E(m)
p = E(m) × E(x) + cov(m, x)
Excess = −cov(m, R) ÷ E(m)

Discount factor	Discount factor: high in bad times
Payoff	Payoff; R is its gross return
Covariance	Covariance of m with the return

IN THE LESSON

- Three equations do the work.
- First, the price p of any payoff x is the expected value of m times x.
- Second, the risk free gross return, capital R f, is one divided by the average discount factor, E of m.
- Third, since the price p equals E of m times E of x, plus the covariance of m with x, an asset's expected excess return over the risk free rate is minus the covariance of m with the return R, divided by E of m.
- The sign is the key.
- A payoff that covaries negatively with m, paying little when m is high, has a positive excess return.
- A payoff that covaries positively with m, paying most in bad times, has a negative one.

24 Example 9: pricing a cyclical claim

Two equal-chance states: m = 0.80 in a boom, 1.10 in a recession

Average discount factor	0.5 × 0.80 + 0.5 × 1.10	**0.95**
Risk-free gross return	1 ÷ 0.95	**1.0526**
Cyclical claim: price	0.5×0.80×130 + 0.5×1.10×80	**$96.00**
Expected payoff	0.5 × 130 + 0.5 × 80	**$105**
Expected return	105 ÷ 96 − 1	**9.375%**

IN THE LESSON

- A boom and a recession are equally likely.
- The discount factor is point eight in a boom and one point one in a recession, higher in the bad state.
- Its average is point nine five.
- The risk free gross return is one over point nine five, one point zero five two six, so the risk free rate in this toy world is five point two six per cent.
- A cyclical claim pays one hundred and thirty dollars in a boom and eighty in a recession.
- Its price is half times point eight times one hundred and thirty, plus half times one point one times eighty: fifty two plus forty four, ninety six dollars.
- Its expected payoff is one hundred and five dollars, so its expected return is nine point three seven five per cent.

25 Example 10: the shelter claim, and the classic mistake

A shelter claim pays $80 in a boom and $130 in a recession. Same m

Cyclical: cov(m, x)	96 − 0.95 × 105	**−3.75**
Cyclical: excess return	(3.75 ÷ 96 price) ÷ 0.95	**4.11%**
Shelter claim: price	0.5×0.80×80 + 0.5×1.10×130	**$103.50**
Shelter: expected return	105 ÷ 103.50 − 1	**1.45%**
Wrong: discount at risk-free	105 ÷ 1.0526	**$99.75**

Check: 9.375% − 5.26% = 4.11%

IN THE LESSON

- A shelter claim pays eighty dollars in a boom and one hundred and thirty in a recession, under the same m.
- For the cyclical claim, the covariance of m with its payoff is the price, ninety six, minus point nine five times one hundred and five, which is minus three point seven five.
- Dividing the payoff covariance by the price, ninety six, turns it into the covariance with the return, so the excess return is three point seven five over ninety six, divided by point nine five, four point one one per cent.
- The shelter claim's price is thirty two plus seventy one point five, one hundred and three dollars fifty.
- Its expected return is only one point four five per cent, below the risk free rate, because it is insurance.
- The classic mistake discounts both at the risk free rate, giving ninety nine dollars seventy five for each.
- As a check, nine point three seven five minus five point two six is four point one one per cent.

26 Building m from preferences

m = β × u′(c1) ÷ u′(c0)
Power utility: m = β × (c1 ÷ c0)^{−γ}

Time discount	β: time discount, not CAPM beta
Marginal utility	u′: marginal utility
Risk aversion	γ: risk aversion
Consumption growth	c1 ÷ c0: consumption growth

IN THE LESSON

- Where does m come from?
- One way is from preferences.
- The discount factor is beta, here the time discount that captures impatience and not the C A P M beta, times marginal utility next period, u prime of c one, divided by marginal utility today, u prime of c zero.
- Marginal utility is how much an extra dollar of consumption adds to well being.
- With power utility, where gamma is risk aversion, this becomes beta times consumption growth, c one over c zero, raised to the power minus gamma.
- Because of that minus sign, m is high when consumption growth is low.
- That is exactly the bad state, when people are poorer and each dollar matters more.

27 Example 11: two ways to build m

β 0.98, γ 2. Growth +4% boom, −2% recession. State prices 0.40, 0.55

Boom: from preferences	0.98 × 1.04^{−2}	**0.906**
Recession: from preferences	0.98 × 0.98^{−2}	**1.020**
Bad state has the higher m	1.020 against 0.906	**bad times**
Boom: state price ÷ prob.	0.40 ÷ 0.5	**0.80**
Recession: state price ÷ prob.	0.55 ÷ 0.5	**1.10**

Separate illustrations; Examples 9 and 10 use 0.80 and 1.10

IN THE LESSON

- Take a time discount of point nine eight and a risk aversion of two.
- Consumption grows four per cent in a boom and falls two per cent in a recession.
- State prices are point four zero and point five five.
- In the boom, m from preferences is point nine eight times one point zero four to the power minus two, about point nine zero six.
- In the recession it is point nine eight times point nine eight to the power minus two, about one point zero two zero.
- So the bad state has the higher m.
- A second construction divides each state price, the price today of a dollar paid only in that state, by the state's probability.
- With equal probabilities, m is point four over point five, point eight, in the boom, and point five five over point five, one point one, in the recession.
- These are separate illustrations with different assumptions, so m differs.
- Examples nine and ten use the state price version, point eight and one point one.

28 Factor models as discount factors

- Linear in factors: m = a − b × factor returns
- Market return as the factor: the CAPM
- Several factors: the multifactor model
- No arbitrage implies a positive m exists
- Trap: ignoring covariance, discounting at rf

IN THE LESSON

- A third construction makes m a linear function of factors: m equals ay minus b times the factor returns.
- When the only factor is the market return, this reproduces the C A P M.
- With several factors it reproduces the multifactor model.
- So the discount factor is not a rival theory.
- It is the general frame that holds the factor models in this lesson.
- A deep result supports it: if there is no arbitrage, meaning no way to make a riskless profit with no outlay, then a positive discount factor exists.
- Remember the usual mistake on the exam.
- It is ignoring covariance and discounting every expected payoff at the risk free rate, which prices the cyclical claim and the shelter claim the same.

29 Efficient markets: three forms

Efficient: prices fully reflect information; no abnormal return after risk and costs

Form	Prices reflect	So this cannot win
Weak	Past prices and volumes	Technical analysis
Semi-strong	All public, incl. past prices	Fundamental analysis
Strong	Private too, incl. all public	Even insiders

IN THE LESSON

- An efficient market is one in which prices fully reflect available information, so no strategy earns abnormal returns after adjusting for risk and costs.
- There are three forms, and each includes the one before it.
- In the weak form, prices reflect past prices and volumes, so technical analysis, the study of price charts, cannot win.
- In the semi strong form, prices reflect all public information, past prices included, so fundamental analysis, the study of accounts and the business, cannot win.
- In the strong form, prices reflect private information too, as well as all public information, so even insiders cannot profit.
- A market that is strong form efficient is therefore also semi strong and weak form efficient.

30 How markets fail to be efficient

- Behavioural biases and noise traders
- Limits to arbitrage: short-sale constraints
- Limits to arbitrage: financing limits
- Limits to arbitrage: the mispricing can widen
- Information costs

IN THE LESSON

- Markets can fail in three ways.
- The first is behavioural biases and noise traders, investors who trade on rumour or emotion rather than information.
- Their trades can push prices away from value.
- Arbitrageurs, who buy what is cheap and sell what is dear, should correct the mispricing.
- But the second source of failure is limits to arbitrage.
- Short sale constraints make it costly or impossible to sell an overpriced asset you do not own.
- Financing limits cap how much an arbitrageur can borrow.
- And there is the risk that the mispricing widens before it closes, forcing the arbitrageur out at a loss.
- The third source is information costs, which we come to shortly.

31 Example 12: right, but out of money first

Corvid: $10m equity; shorts 5 × equity; lender wants equity ≥ 20% of the short

Short position	5 × $10m	**$50m**
Mispricing widens: price +12%	50 × 12%	**−$6m**
Equity left	10 − 6	**$4m**
Equity the lender requires	20% × (50 + 6)	**$11.2m**
Forced to buy back	loss 6 ÷ 10	**60% of equity**

IN THE LESSON

- Corvid Capital, a made up hedge fund, has ten million dollars of equity.
- Its lender allows a short of five times equity, but wants equity of at least twenty per cent of the short.
- Corvid shorts an overpriced share: five times ten million, fifty million dollars.
- Before the price corrects, the mispricing widens and the price rises a further twelve per cent.
- Corvid loses twelve per cent of fifty million, six million dollars, leaving equity of ten minus six, four million.
- The short is now worth fifty six million, so the lender requires twenty per cent of that, eleven point two million.
- Corvid cannot meet it and is forced to buy back, locking in a loss of six out of ten, sixty per cent of its equity.
- Right in the end, it ran out of money first.

32 Information costs and the joint hypothesis

- Grossman and Stiglitz: perfect efficiency is impossible
- If prices told all, nobody would pay to gather information
- Every test is joint: efficiency plus a pricing model
- Apparent alpha may only reflect a faulty model
- Small Value Fund: its 2% CAPM alpha was size and value

IN THE LESSON

- The third source of failure is information costs.
- Grossman and Stiglitz argued that prices cannot be perfectly efficient.
- If prices already told everyone everything, nobody would pay to gather information, and then prices would stop reflecting it.
- So some inefficiency must remain to reward the people who do the research.
- There is also a problem with testing.
- Every test of efficiency is a joint test of efficiency and an asset pricing model, because an abnormal return can be measured only against a model of normal return.
- Apparent alpha may only reflect a faulty model.
- Recall the Small Value Fund.
- Its two per cent alpha under the C A P M was really pay for size and value.

33 Mistakes to avoid

- Paying a premium for total volatility
- Beta times Rm instead of times (Rm − rf)
- σi ÷ σm as beta, without the correlation
- Discounting risky payoffs at the risk-free rate
- Reading alpha as inefficiency without checking the model

IN THE LESSON

- First, paying a premium for total volatility.
- Premiums follow covariance with bad times, so a volatile stock like Vantor can expect less than a calmer one.
- Second, applying beta to the market return instead of the market risk premium.
- Third, using sigma i over sigma m as beta and forgetting the correlation.
- Fourth, discounting risky payoffs at the risk free rate, which ignores their covariance with the discount factor.
- Fifth, reading alpha as proof of inefficiency without checking the model, when the alpha may be pay for an omitted factor.

34 Try three quick questions

1. Risk-free 2%, Rm 8%, ρ 0.4, σi 36%, σm 18%. Expected return?
   Answer: β = 0.4 × 36 ÷ 18 = 0.8. 2% + 0.8 × 6% = 6.8%.
2. m = 0.90 in a boom, 1.06 in a recession, each 50%. Value of $100 paid in a recession only?
   Answer: 0.5 × 1.06 × 100 = $53.
3. rf 3%. Market β 1.2, premium 5%. Value β −0.5, premium 4%. Expected return?
   Answer: 3% + 6% − 2% = 7%.

IN THE LESSON

- Now try three quick questions.
- Pause, think, then listen to the answers.
- Question one.
- The risk free rate is two per cent and the market's expected return eight per cent.
- A stock's correlation with the market is point four, its volatility thirty six per cent, and the market's eighteen.
- Expected return?
- Beta is point four times thirty six over eighteen, point eight, so the return is two plus point eight times six, six point eight per cent.
- Question two.
- The discount factor is point nine in a boom and one point zero six in a recession, each fifty per cent likely.
- What is one hundred dollars paid only in a recession worth?
- Half times one point zero six times one hundred, fifty three dollars.
- Question three.
- The risk free rate is three per cent.
- A stock has a market beta of one point two at a five per cent premium, and a value beta of minus point five at four per cent.
- Expected return?
- Three plus six minus two, seven per cent.