CHARTA LEARNING · LESSON NOTES

Modern Portfolio Theory and the CAPM

FRM Part I

1 A story: do not judge a player alone

- Judge an asset by what it adds to the whole portfolio
- Umbrellas and ice cream: the ups and downs offset
- A portfolio's risk is not the average of its parts' risks
- Making that precise gives Modern Portfolio Theory

IN THE LESSON

- Imagine you are building a football team.
- You would not pick the eleven best strikers, because the team would have no goalkeeper and no defence.
- You judge each player by what he adds to the team.
- Investing works the same way.
- Suppose you own an umbrella company and an ice cream company.
- In wet years one does well, in sunny years the other does.
- Each business alone is risky, yet together your income is steadier.
- The risk of a collection is not simply the average of the risks of its parts.
- That observation, made precise, is the foundation of modern portfolio theory.

2 Markowitz and the idea of MPT

- Investors care about expected return and risk (volatility)
- Portfolio return = weighted average of asset returns
- Portfolio risk also depends on how assets move together
- Measured by covariance, or its scaled form, correlation

IN THE LESSON

- Modern Portfolio Theory comes from Harry Markowitz, who set it out in the early nineteen fifties.
- He proposed that investors care about two things: the expected return of a portfolio, and its risk, measured by the standard deviation of returns, which we call volatility.
- The expected return of a portfolio is easy.
- It is the weighted average of the returns of its assets.
- Risk is different.
- It depends not just on how risky each asset is, but on how the assets move together, which is measured by covariance, or by its scaled version, correlation.
- This is why adding a risky asset can sometimes lower the risk of a portfolio.

3 Portfolio variance for two assets

$$\sigma_p^2 = w_1^2\sigma_1^2 + w_2^2\sigma_2^2 + 2 \cdot w_1 \cdot w_2 \cdot \rho \cdot \sigma_1 \cdot \sigma_2$$

Weights, w	Share of money in each asset
Volatility, sigma	Standalone risk of each asset
Correlation, rho	How closely the two move together
Portfolio variance	Portfolio risk, squared

IN THE LESSON

- Here is the key formula, for two assets.
- Portfolio variance equals the first weight squared times the first variance, plus the second weight squared times the second variance, plus two times the two weights times the correlation times both volatilities.
- The weights say how much of the portfolio sits in each asset.
- The volatilities are the standalone risks.
- The correlation, which runs from minus one to plus one, says how closely the two move together.
- The last term is the whole story.
- With both weights positive and correlation below plus one, it shrinks, so portfolio volatility ends up below the weighted average of the two volatilities.
- That gap is the benefit of diversification.

4 Example 1: a fifty-fifty portfolio

Illustration: A 6% return, 10% risk. B 12% return, 20% risk. Correlation 0

Portfolio return	0.5×6% + 0.5×12%	**9%**
Variance	0.0025 + 0.0100 + 0	**0.0125**
Portfolio volatility	√0.0125	**11.18%**
Average of the two risks	0.5×10% + 0.5×20%	**15%**

IN THE LESSON

- Let's use made up numbers.
- Asset A has an expected return of six per cent and a volatility of ten per cent.
- Asset B has an expected return of twelve per cent and a volatility of twenty per cent.
- Suppose the two are uncorrelated, a correlation of zero.
- Put half in each.
- The return is the average: nine per cent.
- For the variance, half squared is point two five.
- First term: point two five times point zero one is point zero zero two five.
- Second term: point two five times point zero four is point zero one.
- The third term vanishes because the correlation is zero.
- The total is point zero one two five, a volatility of eleven point one eight per cent.
- The weighted average of the two volatilities is fifteen, so diversification has cut nearly four points of risk, with no loss of return.

5 Correlation drives the benefit

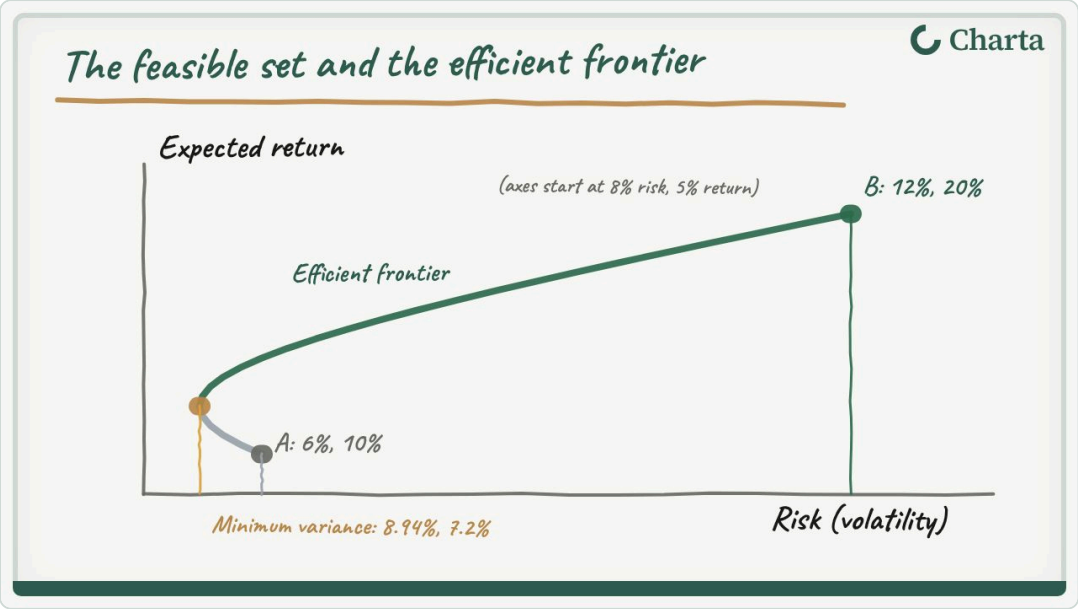
Same 50/50 mix, same 9% return, different correlation

Correlation	Portfolio volatility
+1.0	15.00%
+0.5	13.23%
0.0	11.18%
−1.0	5.00%

IN THE LESSON

- Now watch what correlation does to the same fifty fifty mix.
- If the two assets move in perfect lockstep, a correlation of plus one, there is no benefit at all: volatility is the simple average, fifteen per cent.
- At a correlation of plus point five it falls to about thirteen point two.
- At zero, where the assets move independently, it is eleven point one eight.
- And at minus one, where they move in perfect opposition, the risks cancel and volatility drops to just five per cent.
- The lower the correlation, the greater the gain.
- The return is nine per cent in every row.
- Diversification lowers risk without lowering expected return, and that is why it is called the only free lunch in finance.

6 The feasible set and the efficient frontier



IN THE LESSON

- Now draw every possible mix of our two assets.
- Along the bottom is risk, the volatility.
- Up the side is expected return.
- At the lower left sits asset A, at six per cent return; at the upper right, asset B, at twelve.
- Mixes lie along the curve between them, and because the correlation is below one, the curve bulges to the left.
- The leftmost point is the global minimum variance portfolio, the lowest risk you can achieve.
- From there the upper branch, in green, is the efficient frontier: the highest return available for each level of risk.
- The short lower branch, in grey, is inefficient, because a point directly above it offers more return for the same risk.

7 Reading the frontier

- Dominated: more return or less risk is available elsewhere
- Efficient set = upper-left edge, from minimum variance up
- Your risk appetite picks the point on the frontier
- Diversification does not remove all risk

IN THE LESSON

- Let's read the picture in plain words.
- Any portfolio inside the region, or on the lower branch, is dominated: some other portfolio gives more return for the same risk, or less risk for the same return, so no sensible investor would hold it.
- Only the upper left boundary, starting at the global minimum variance portfolio and running upward, is efficient.
- Which point on it you choose depends on your own appetite for risk.
- One common error to avoid.
- Diversification does not remove all risk.
- A well diversified portfolio sheds the risk that is specific to individual assets, but keeps the broad market part, and real portfolios can keep leftover risk from concentration, shifting correlations and errors in the estimates themselves.

8 Example 2: the minimum variance mix

Minimum variance weight in A = (σB² − ρσAσB) ÷ (σA² + σB² − 2ρσAσB)

Weight in A	0.04 ÷ 0.05	**80%**
Weight in B	1 − 80%	**20%**
Return	0.8×6% + 0.2×12%	**7.2%**
Volatility	√0.008	**8.94%**

IN THE LESSON

- Which mix of our two assets carries the lowest risk?
- With two assets there is a standard formula.
- The weight in A equals B's variance, minus the correlation times both volatilities, all divided by the two variances added together, minus twice the correlation times both volatilities.
- With a correlation of zero, it is simple.
- B's variance is point zero four.
- The two variances add to point zero five.
- So the weight in A is point zero four over point zero five, which is eighty per cent, and twenty per cent goes in B.
- The return is eighty per cent of six plus twenty per cent of twelve, which is seven point two per cent.
- The variance is point zero zero eight, a volatility of eight point nine four per cent.
- That is lower than either asset alone, and with a higher return than A.

9 From MPT to the CAPM: the big step

- Same forecasts and the same risk-free rate for all
- So everyone picks the same best risky portfolio (tangency)
- Cautious or bold, they differ only in how much cash to hold
- Every asset is held, so tangency = market portfolio

IN THE LESSON

- The Capital Asset Pricing Model asks what happens if every investor follows Markowitz.
- Suppose all of them have the same forecasts, and all can lend and borrow at the same risk free rate.
- Then every investor looks at the same efficient frontier and picks the same best risky portfolio, the one that gives the highest reward per unit of risk.
- This is called the tangency portfolio.
- Each person just mixes it with the risk free asset, more of one for the cautious and more of the other for the bold.
- But every risky asset has to be owned by someone.
- So the tangency portfolio must be the market portfolio, which holds everything in proportion to its value.

10 The CAPM and its parts

$E(R_i) = Rf + \beta_i \times [E(Rm) - Rf]$

Expected return	Return the asset should offer
Risk-free rate	Price of time
Market risk premium	Price of risk
Beta	Quantity of systematic risk

IN THE LESSON

- The result is a single line.
- Expected return equals the risk free rate plus beta times the market risk premium.
- Take the parts in turn.
- The risk free rate is the price of time: what you earn for waiting, with no risk.
- The market risk premium is the expected market return minus the risk free rate, the price of risk.
- And beta is the quantity of systematic risk the asset carries, defined as the covariance of the asset with the market divided by the variance of the market.
- The key message: only systematic risk is rewarded.
- The market pays you for the risk you cannot diversify away, not for total volatility.

11 The assumptions behind the CAPM

- Rational, risk-averse: only mean and variance matter
- One common period; normal returns or quadratic utility
- Homogeneous expectations: everyone sees the same frontier
- Unlimited lending and borrowing at one risk-free rate
- Frictionless: no taxes or costs, price takers, same info

IN THE LESSON

- The CAPM is an equilibrium model, and its tidy answer rests on strong assumptions.
- Investors are rational, risk averse, and care only about expected return and variance over one common period.
- That focus is justified if returns are normal, or if utility is quadratic.
- They hold homogeneous expectations, so they all see the same efficient frontier.
- They can lend and borrow any amount at one risk free rate.
- And markets are frictionless: no taxes, no trading costs, assets can be divided and traded freely, everyone is a price taker, and everyone has the same information.
- Together these make every investor hold the same risky portfolio.

12 When an assumption fails

Relax one assumption and the single tangency portfolio breaks

If this fails	Then
One risk-free rate	No single tangency line
Same expectations	Frontiers differ by investor
No taxes or costs	Best mix differs by investor
Normal returns	Variance misses tail risk
Price takers, full access	Prices can be pushed around

IN THE LESSON

- Now a typical scenario question: if one assumption is relaxed, what happens?
- Suppose borrowing costs more than lending.
- Then there is no single tangency line, and different investors can hold different risky portfolios.
- Suppose expectations differ.
- Then investors see different frontiers, and again the market portfolio need not be the one tangency portfolio.
- Suppose there are taxes or trading costs.
- Then the best portfolio differs by investor.
- Suppose returns are not normal, with fat tails.
- Then variance alone does not capture risk, and beta can miss the tail risk.
- Suppose some investors are large enough to move prices, or have better information.
- Then prices can be pushed around.
- In each case the model does not collapse, but its neat line becomes an approximation, and you should say which assumption fails.

13 The capital market line

- CML: expected return against total risk (σ)
- Rp = Rf + [(Rm − Rf) ÷ σm] × σp
- Slope = the market's Sharpe ratio
- Efficient portfolios only: risk-free asset + market

IN THE LESSON

- Now the first of two lines.
- Mix the risk free asset with the market portfolio and you trace out a straight line, the capital market line.
- Put total risk, the standard deviation, on the horizontal axis and expected return on the vertical axis.
- The line starts at the risk free rate, with no risk, and rises through the market portfolio.
- Its equation is the risk free rate plus the market risk premium divided by the market volatility, times the portfolio's own volatility.
- The slope is therefore the market's Sharpe ratio: reward per unit of total risk.
- It describes only efficient portfolios, those that are combinations of the risk free asset and the market.
- Hold less than all in the market and you are lending; hold more than all and you are borrowing.

14 Example 3: two points on the CML

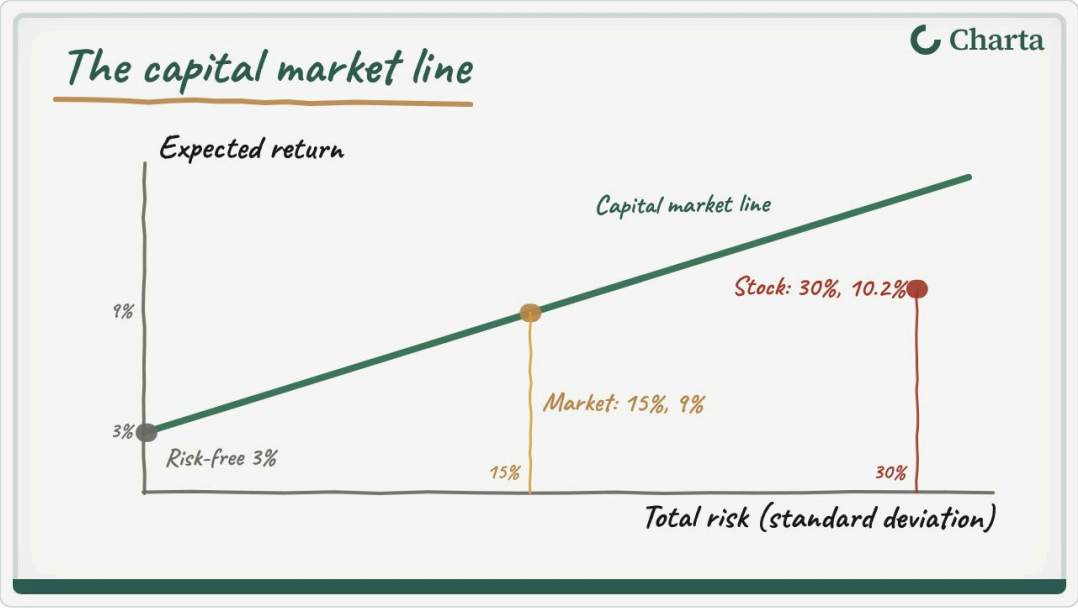
Rf 3%, market return 9%, market volatility 15%. Slope = 6 ÷ 15 = 0.4

Slope (market Sharpe)	(9% − 3%) ÷ 15%	**0.40**
Cautious: 10% risk	2/3 market, 1/3 cash	**7%**
Check on the line	3% + 0.4 × 10%	**7%**
Bold: 20% risk	133% market, borrow 33%	**11%**

IN THE LESSON

- Let's use made up market numbers.
- The risk free rate is three per cent.
- The market is expected to return nine per cent, with a volatility of fifteen per cent.
- The premium is six per cent, and the slope of the line is six over fifteen, point four.
- A cautious investor wants a volatility of ten per cent.
- Two thirds in the market gives two thirds of fifteen, which is ten.
- The rest, one third, sits in the risk free asset.
- Return: two thirds of nine plus one third of three, equal to seven per cent.
- The line agrees: three plus point four times ten is seven.
- A bold investor wants twenty per cent.
- That needs one and one third times the market, borrowing one third at three per cent.
- Return: twelve minus one, eleven per cent.
- The line gives three plus eight, eleven.

15 The capital market line



IN THE LESSON

- Here is the line.
- Along the bottom is total risk, the standard deviation, and up the side is expected return.
- The line starts at three per cent on the vertical axis, the risk free rate, where there is no risk at all.
- It passes through the market portfolio, at fifteen per cent risk and nine per cent return.
- To the left of the market you are lending, mixing in the risk free asset; to the right you are borrowing to invest more than your own money.
- Now look at an individual stock with thirty per cent volatility and an expected return of ten point two per cent.
- It sits well below the line.
- It carries a lot of total risk, but much of that risk is diversifiable, and the market does not pay for it.

16 The CML and the SML side by side

Both lines come from the CAPM; they answer different questions

	CML | SML
Horizontal axis	Total risk σ | Beta
Applies to	Efficient | Everything
Slope	Sharpe ratio | Premium
Equation	Rf + slope×σ | Rf + β×premium
Mispriced asset	n/a | Alpha

IN THE LESSON

- Now the second line, the security market line.
- It also comes from the CAPM, but answers a different question.
- The capital market line puts total risk on the horizontal axis and applies only to efficient portfolios.
- The security market line puts beta on the horizontal axis, and applies to every asset and every portfolio, efficient or not.
- Its slope is the market risk premium, not the Sharpe ratio.
- A fairly priced asset sits exactly on the security market line.
- One above it has a positive alpha and looks underpriced.
- One below looks overpriced.
- Under the CAPM, an individual stock plots on the security market line but below the capital market line.
- Do not mix up the axes or the slopes.

17 Example 4: required return and alpha

Rf 3%, market 9%, premium 6%. Work in percentages throughout

Required, beta 1.2	3% + 1.2 × 6%	**10.2%**
Forecast 12%: alpha	12% − 10.2%	**+1.8%**
Required, beta 0.8	3% + 0.8 × 6%	**7.8%**
Forecast 6.5%: alpha	6.5% − 7.8%	**−1.3%**

IN THE LESSON

- Now use the CAPM to find what an asset should earn.
- Take the same market: risk free three per cent, market nine per cent, so the premium is six per cent.
- Do not use the market return itself, use the premium.
- A stock with beta one point two should return three plus one point two times six, which is ten point two per cent.
- That is the required return.
- A beta of one would give nine, the market return.
- A beta of zero would give three, the risk free rate, however volatile the asset is.
- Now suppose analysts forecast twelve per cent for our stock.
- Forecast minus required is plus one point eight per cent, its alpha.
- It looks cheap.
- Another stock with beta point eight requires seven point eight, but is forecast at six point five.
- Alpha is minus one point three, so it looks expensive.

18 Example 5: a fair price from the CAPM

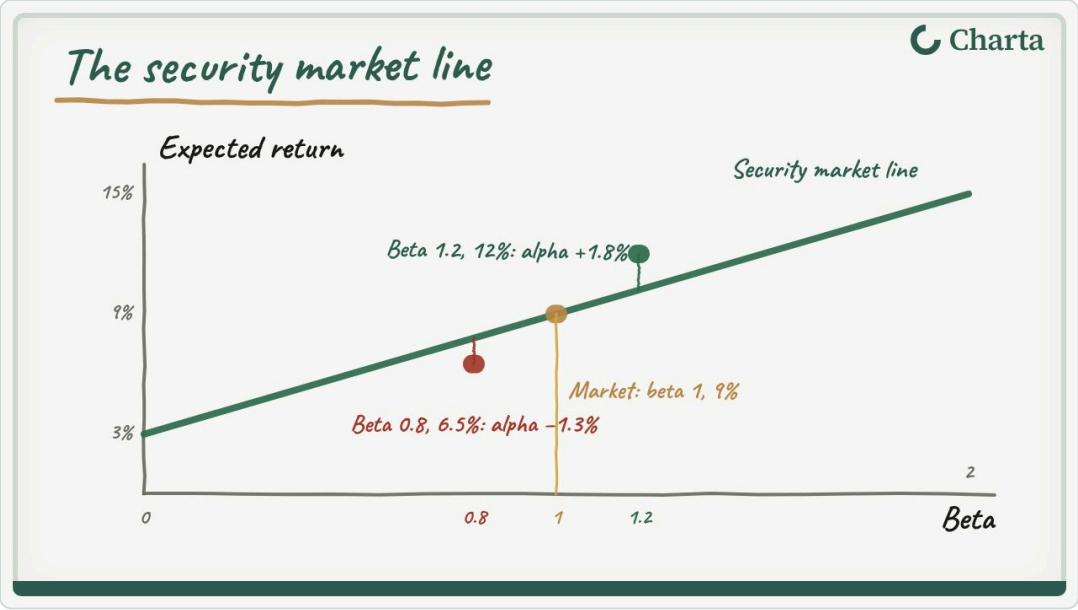
Illustration: stock with beta 1.2, so required return 10.2%

Expected payoff	105 price + 2 dividend	**107**
Discount factor	1 + 10.2%	**1.102**
Fair price today	107 ÷ 1.102	**$97.10**
If trading at $95	cheap: positive alpha	**Cheap**
If trading at $100	dear: negative alpha	**Dear**

IN THE LESSON

- The required return can also work as a discount rate, to find a fair price.
- Suppose our beta one point two stock is expected to be worth one hundred and five dollars in a year, and to pay a two dollar dividend.
- The expected payoff is price plus dividend: one hundred and seven.
- Divide by one plus the CAPM return, which is one point one zero two.
- The fair price is ninety seven dollars and ten cents.
- If the stock trades at ninety five, it is cheap: the buyer earns a return above the required ten point two.
- If it trades at one hundred, it is dear.
- Notice the logic: a riskier asset has a higher required return, which means a heavier discount, which means a lower price today.

19 The security market line



IN THE LESSON

- Here is the security market line.
- Along the bottom is beta, from zero to two.
- Up the side is expected return.
- The line starts at three per cent for a beta of zero, passes through nine per cent at a beta of one, the market, and reaches fifteen per cent at a beta of two.
- Its slope is the six per cent market premium.
- Our first stock, beta one point two, is forecast to return twelve per cent, above the required ten point two, so it plots above the line: a positive alpha, underpriced.
- Our second stock, beta point eight, forecast at six point five, plots below the line, where the required return is seven point eight: a negative alpha, overpriced.
- Fair pricing lies on the line itself.

20 What beta means

- Beta = sensitivity of an asset to the market
- Market = 1, risk-free = 0, above 1 amplifies the market
- Beta is the slope of the asset-on-market regression
- Beta is not total volatility, and not correlation

IN THE LESSON

- Let's look closely at beta.
- It measures how sensitive an asset is to the market.
- If the market rises ten per cent and a stock with a beta of one point five tends to rise fifteen, the stock amplifies market moves.
- A beta of one means it moves with the market.
- A beta of zero means no relationship with the market, like the risk free asset.
- If you plot the asset's returns against the market's and fit a line, beta is the slope of that regression.
- Beta is a statement about co movement with the market, not about total volatility.
- Two assets can have the same volatility and very different betas, and the CAPM pays only for beta.

21 Two ways to compute beta

$$\beta = \mathrm{Cov}(R_i, Rm) \div \mathrm{Var}(Rm)$$
$$= \rho \times \sigma_i \div \sigma m$$

Beta	Sensitivity to the market
Covariance with market	How asset and market move together
Variance of market	Variance, not standard deviation
Correlation, rho	Correlation with the market

IN THE LESSON

- Beta has two equivalent formulas.
- First, beta equals the covariance of the asset's return with the market's return, divided by the variance of the market return.
- Second, beta equals the correlation between the two, times the asset's volatility, divided by the market's volatility.
- The second follows from the first, because covariance is correlation times both volatilities.
- Notice that the volatility of the asset enters beta, which is why a volatile but poorly correlated stock can have a modest beta.
- Watch out for a classic slip: dividing the covariance by the market's standard deviation, instead of its variance.
- That gives the wrong number.

22 Example 6: computing beta

Stock volatility 30%, market volatility 15%, correlation 0.6

Covariance	0.6 × 0.30 × 0.15	**0.0270**
Market variance	0.15^2	**0.0225**
Beta	0.0270 ÷ 0.0225	**1.2**
Check	0.6 × 30% ÷ 15%	**1.2**

IN THE LESSON

- Suppose a stock has a volatility of thirty per cent and a correlation with the market of point six.
- The market has a volatility of fifteen per cent, so its variance is point zero two two five.
- The covariance is the correlation times both volatilities: point six times point three times point one five, which is point zero two seven.
- Beta is covariance over market variance: point zero two seven divided by point zero two two five, which is one point two.
- Check with the second formula: point six times thirty over fifteen, point six times two, one point two again.
- Now the trap.
- If you had divided the covariance by the market's standard deviation, you would get point one eight, which is not a beta at all.

23 Example 7: splitting the variance

Stock: beta 1.2, volatility 30%. Market variance 0.0225. Model-based split

Systematic variance	$1.2^2 \times 0.0225$	**0.0324**
Idiosyncratic variance	0.09 − 0.0324	**0.0576**
Share explained (R^2)	0.0324 ÷ 0.09	**36%**
Idiosyncratic volatility	√0.0576	**24%**

IN THE LESSON

- Beta also splits an asset's variance in two.
- In the market model regression, where the leftover part is uncorrelated with the market, the systematic part is beta squared times the market variance.
- For our stock: one point four four times point zero two two five, equal to point zero three two four.
- The stock's total variance is thirty per cent squared, point zero nine.
- The rest, point zero five seven six, is idiosyncratic, with a volatility of twenty four per cent.
- The share of variance explained by the market is point zero three two four over point zero nine, which is thirty six per cent.
- That is the regression R squared, and it equals the correlation squared, point six squared.
- Only the thirty six per cent is rewarded.
- The rest can be diversified away.

24 Example 8: beta of a whole portfolio

Weights are on the whole portfolio, cash included at beta 0

40% in beta 1.2	0.4 × 1.2	**0.48**
40% in beta 0.8	0.4 × 0.8	**0.32**
20% in cash	0.2 × 0	**0.00**
Portfolio beta	0.48 + 0.32 + 0	**0.80**

IN THE LESSON

- Portfolio beta is the value weighted average of the asset betas, and cash enters at zero.
- Suppose you hold forty per cent of your portfolio in a stock with beta one point two, forty per cent in a stock with beta point eight, and twenty per cent in cash.
- The portfolio beta is point four times one point two, plus point four times point eight, plus point two times zero: point four eight, plus point three two, plus nothing, equals point eight.
- The weights must be on the whole portfolio, including cash or borrowing.
- If instead you borrowed, the weights on the stocks would add to more than one, and beta would be scaled up.
- The CAPM return for our portfolio is three plus point eight times six, which is seven point eight per cent.

25 Three ways to judge performance

Sharpe = (Rp − Rf) ÷ σp Treynor = (Rp − Rf) ÷ βp
Alpha = Rp − [Rf + βp × (Rm − Rf)]

Excess return	Return above the risk-free rate
Volatility, sigma	Total risk: Sharpe divides by it
Beta	Systematic risk: Treynor divides
Expected return	CAPM return: alpha subtracts it

IN THE LESSON

- Now we judge how well a fund did, and the key is which risk we penalise.
- The Sharpe ratio is the return above the risk free rate, divided by total risk, the standard deviation.
- It suits a whole portfolio.
- The Treynor ratio is the same excess return, divided by beta instead, so it penalises only systematic risk.
- Jensen's alpha is different: it is the fund's return minus what the CAPM says it should have earned for its beta.
- Sharpe and Treynor are ratios and rank funds.
- Alpha is in return units and says how far the fund sits above or below the security market line.
- Forgetting to subtract the risk free rate is the usual slip.

26 Example 9: Sharpe ratios of three funds

Rf 3%, market 9%. X: 11%, risk 20%, β1.2. Y: 10%, 16%, β0.9. Z: 11%, 25%, β1.0

Fund X	(11 − 3) ÷ 20	**0.40**
Fund Y	(10 − 3) ÷ 16	**0.44**
Fund Z	(11 − 3) ÷ 25	**0.32**
Ranking by Sharpe		**Y, X, Z**

IN THE LESSON

- Three funds, with made up numbers, in a market where the risk free rate is three per cent and the market return nine.
- Fund X returned eleven per cent with a volatility of twenty and a beta of one point two.
- Fund Y returned ten, volatility sixteen, beta point nine.
- Fund Z returned eleven, volatility twenty five, beta one.
- Sharpe ratio is excess return over volatility.
- Fund X: eight over twenty, point four.
- Fund Y: seven over sixteen, point four four.
- Fund Z: eight over twenty five, point three two.
- On total risk, fund Y is the best and fund Z the worst.
- Compare with the market's own Sharpe ratio of point four.
- Fund X only matches it, while fund Y beats it slightly.

27 Example 10: Treynor ratios of the same funds

Treynor = excess return ÷ beta. Market Treynor = 6% ÷ 1 = 6.0

Fund X	(11 − 3) ÷ 1.2	**6.67**
Fund Y	(10 − 3) ÷ 0.9	**7.78**
Fund Z	(11 − 3) ÷ 1.0	**8.00**
Ranking by Treynor		**Z, Y, X**

IN THE LESSON

- Now the Treynor ratio, which divides the same excess returns by beta instead.
- Fund X: eight divided by one point two, which is six point seven.
- Fund Y: seven divided by point nine, seven point eight.
- Fund Z: eight divided by one, eight point zero.
- The ranking has changed.
- Fund Z is now the best and fund X the worst.
- Compare each with the market's Treynor ratio, which is the premium, six per cent, because the market's beta is one.
- All three beat the market on this measure.
- Why did Z jump?
- Because it carries a lot of risk that is specific to the fund, which the Sharpe ratio punishes and the Treynor ratio ignores.
- Treynor is the right measure when the fund is one slice of a larger, well diversified portfolio.

28 Example 11: Jensen's alpha

Alpha = fund return − CAPM required return (Rf 3%, premium 6%)

Fund X	11% − (3% + 1.2×6%)	**+0.8%**
Fund Y	10% − (3% + 0.9×6%)	**+1.6%**
Fund Z	11% − (3% + 1.0×6%)	**+2.0%**
Ranking by alpha		**Z, Y, X**

IN THE LESSON

- Finally alpha.
- For each fund, work out what the CAPM requires, then subtract.
- Fund X, beta one point two: required three plus one point two times six, ten point two.
- Return eleven, so alpha is plus point eight.
- Fund Y, beta point nine: required three plus five point four, eight point four.
- Return ten, so alpha is plus one point six.
- Fund Z, beta one: required nine.
- Return eleven, so alpha is plus two.
- All three have positive alpha, and the ranking is Z, Y, X, the same as Treynor here.
- Both measure return against beta, so they often agree, but they need not.
- Notice the three rankings do not agree.
- A fund with high idiosyncratic risk looks worse on Sharpe than on Treynor or alpha, so always say which risk you are penalising.

29 Example 12: the Sortino ratio

Yearly returns 12, −4, 8, 2, 10 (%). Target 4%

Average return	28 ÷ 5	**5.6%**
Squared shortfalls	$8^2 + 2^2 = 68$, ÷ 5	**13.6**
Downside deviation	√13.6	**3.69%**
Sortino ratio	(5.6 − 4) ÷ 3.69	**0.43**

IN THE LESSON

- The Sortino ratio fixes a complaint about Sharpe: volatility penalises upside surprises just as much as downside ones.
- Sortino counts only returns below a target.
- It is the average return minus the target, divided by the downside deviation.
- Take five yearly returns: twelve, minus four, eight, two and ten per cent, with a target of four.
- The average is five point six.
- Only two years fall short: minus four misses by eight, and two misses by two.
- Square the shortfalls, sixty four and four, add to sixty eight, and divide by all five observations, not just the two, to get thirteen point six.
- The square root, the downside deviation, is three point six nine.
- Sortino is five point six minus four, over three point six nine: point four three.

30 Example 13: tracking error and information ratio

Active return = fund − benchmark, in % per year: +2, −1, +3, 0, +1

Mean active return	5 ÷ 5	**1.0%**
Variance (sample)	10 ÷ 4	**2.5**
Tracking error	√2.5	**1.58%**
Information ratio	1.0 ÷ 1.58	**0.63**

IN THE LESSON

- Last, the measures for a fund that is judged against a benchmark.
- Active return is the fund's return minus the benchmark's return.
- Tracking error is the standard deviation of that active return.
- The information ratio is the mean active return divided by the tracking error.
- It shows how much extra return you earned for each unit of active risk taken.
- Suppose five years of active returns are plus two, minus one, plus three, zero and plus one per cent.
- The mean is one.
- The deviations from the mean are one, minus two, two, minus one and zero.
- Squares add to ten.
- Using the sample divisor, four, that gives two point five, so tracking error is one point five eight per cent.
- The information ratio is one over one point five eight, which is point six three.

31 Four mistakes to avoid

- Diversification does not remove all risk; only beta is paid
- CAPM: use the premium, add Rf, keep units consistent
- Beta = covariance ÷ variance, not correlation
- Match ratio to risk; subtract Rf; CML is not SML

IN THE LESSON

- Mistake one: thinking that diversification removes all risk, or that you are paid for total volatility.
- The CAPM rewards only beta, and real portfolios keep some risk.
- Mistake two: sloppy algebra in the CAPM.
- Use the market premium and not the market return, remember to add the risk free rate, and do not mix percentages with decimals.
- Mistake three: confusing beta with correlation, or dividing covariance by a standard deviation.
- Beta is covariance over variance.
- Mistake four: dividing by the wrong risk, or forgetting to subtract the risk free rate.
- Sharpe uses total risk, Treynor uses beta, alpha compares with the CAPM line, and Sortino uses downside only.
- And do not mix up the two market lines.

32 Recap: the main ideas

- Risk depends on correlation; frontier = efficient edge
- CAPM: E(R) = Rf + β × premium; strong assumptions
- CML: total risk, efficient only. SML: beta, every asset
- Beta = Cov ÷ Var; portfolio beta is the weighted average
- Sharpe: total risk. Treynor, alpha: beta. Sortino: downside

IN THE LESSON

- Let's pull it together.
- Portfolio return is a weighted average, but portfolio risk depends on correlation, so combining imperfectly correlated assets cuts risk.
- The efficient frontier is the upper boundary from the minimum variance portfolio.
- If everyone follows this approach, they hold the market portfolio plus the risk free asset, and the CAPM says expected return equals the risk free rate plus beta times the premium.
- It rests on strong assumptions.
- The capital market line has total risk on its axis and applies to efficient portfolios.
- The security market line has beta on its axis and applies to everything.
- Beta is covariance over market variance.
- Sharpe, Treynor, alpha, Sortino and the information ratio each judge a different kind of risk.

33 Try three quick questions

1. Rf 4%, market 10%, beta 1.5. CAPM required return?
   Answer: 4% + 1.5 × (10% − 4%) = 13%.
2. Return 12%, risk 15%, beta 1.2, Rf 4%. Sharpe and Treynor?
   Answer: Sharpe = 8 ÷ 15 = 0.53. Treynor = 8 ÷ 1.2 = 6.67.
3. Cov with market 0.03, market variance 0.02. Beta?
   Answer: 0.03 ÷ 0.02 = 1.5. Use the variance.

IN THE LESSON

- Now try three quick questions yourself.
- Pause after each one.
- Question one.
- The risk free rate is four per cent, the market return is ten, and a stock has a beta of one point five.
- What return does the CAPM require?
- Four plus one point five times six is thirteen per cent.
- Question two.
- A fund returns twelve per cent with a volatility of fifteen and a beta of one point two, when the risk free rate is four.
- What are its Sharpe and Treynor ratios?
- Eight over fifteen, point five three, and eight over one point two, six point seven.
- Question three.
- A stock has a covariance with the market of point zero three and the market variance is point zero two.
- What is its beta?
- One point five.