Capital Markets and the Pricing of Risk

From diversification to a defensible required return

Summer 2026

Computational Finance / FinTech Summer School

Day 4 | final-year undergraduate / introductory master's level

Day 3 used a 9.5% cost of equity; Day 4 must defend it

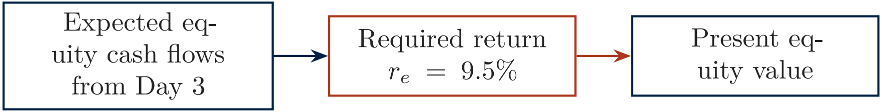

- Cash-flow forecasts came from operations and payout policy.
- The 9.5% rate was declared, not estimated.
- Today we ask what risk earns a premium and how that premium enters valuation.

Cost of equity r_e = expected return required by shareholders for bearing the equity claim's risk.

Decision brief: estimate a GBP required return for Aster

Valuation date	14 August 2026
Company	Aster Payments plc; fictional case
Cash-flow currency	Nominal pounds sterling
Observed decision input	Day 3 cost of equity: 9.5%
Evidence	Market returns, industry portfolios, and a disclosed synthetic Aster sample
Committee question	Is 9.5% a defensible base hurdle, and what range should surround it?

Required return and cash flow must share currency and inflation convention. A hurdle rate is the minimum expected return required for accepting an investment of comparable risk.

Day 4 roadmap: six questions turn risk into a required return

1. **Measure:** what do return and risk statistics describe?
2. **Diversify:** which risk survives inside a portfolio?
3. **Price:** why does beta determine required return?
4. **Estimate:** what is observed, assumed, and uncertain?
5. **Apply:** does market risk support Aster's 9.5%?
6. **Judge:** when is CAPM useful, and when should it be challenged?

The route separates three tasks often collapsed into one number: economic model, empirical estimate, and decision judgment.

By the end, you should be able to defend the estimate and its limits

- **Calculate** expected return, volatility, covariance, and portfolio risk.
- **Explain** why diversification removes firm-specific but not market risk.
- **Interpret** beta and derive a CAPM required return.
- **Estimate** beta while reporting benchmark, sample, interval, and residual risk.
- **Match** company or project risk to the appropriate hurdle rate.
- **Distinguish** CAPM evidence from market-efficiency evidence.
- **Recommend** a base rate, a range, and evidence that would change it.

CAPM = Capital Asset Pricing Model, an equilibrium benchmark linking expected return to exposure to market risk.

Days 1–4 complete the first valuation chain

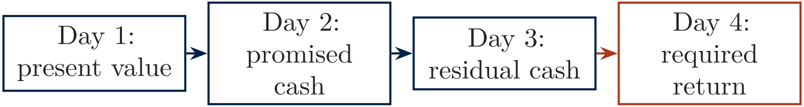

$$\text{value} = \sum_t \frac{\mathbb{E}[\text{cash flow}_t]}{(1 + \text{required return})^t}.$$

Days 1–3 developed the numerator and discounting mechanics. Day 4 studies the risk adjustment in the denominator.

Measure return and risk

What do return and risk statistics actually describe?

Begin with future states and observed samples; keep expectation separate from realised performance.

A holding-period return combines cash income and price change

$$R_{t+1} = \frac{D_{t+1} + P_{t+1} - P_t}{P_t}.$$

- P_t: purchase price at the start of the period.
- D_{t+1}: cash distribution during the period.
- P_{t+1}: ex-distribution price at the end.
- The realised return is known only after D_{t+1} and P_{t+1} occur.

Holding period = interval over which the investor owns the asset. Return is a rate; profit is a monetary amount.

Appendix A1 · Identity Python P2 · Returns

Before the period, return is a distribution rather than a fact

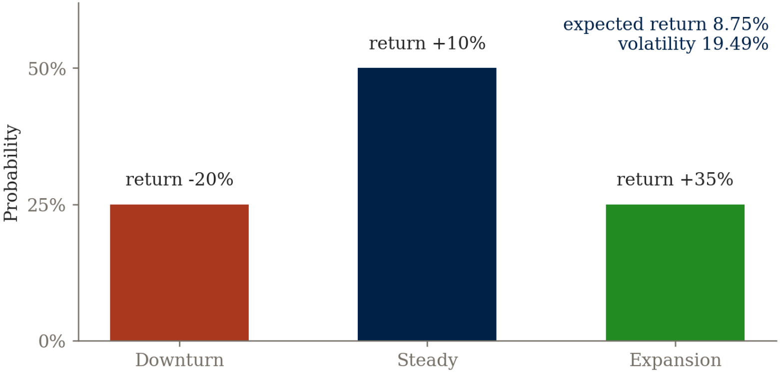

Constructed Aster state example. Returns and probabilities are illustrative inputs.

A state is one mutually exclusive future outcome. A probability distribution assigns a probability to every possible return.

Expected return is the probability-weighted centre

$$\mathbb{E}[R] = \sum_{s=1}^{S} p_s R_s, \qquad \sum_{s=1}^{S} p_s = 1.$$

- R_s is the return if state s occurs.
- p_s is the probability assigned to that state.
- The expected return is not the outcome that must occur; it is a probability-weighted average across outcomes.

$\mathbb{E}[\cdot]$ denotes an expectation. Probabilities are model inputs estimated from evidence or stated as scenarios.

Volatility measures dispersion around expected return

$$\mathrm{Var}(R) = \sum_{s=1}^{S} p_s \big(R_s - \mathbb{E}[R]\big)^2, \qquad \sigma_R = \sqrt{\mathrm{Var}(R)}.$$

- Squaring prevents positive and negative deviations from cancelling.
- Variance is in squared-return units; standard deviation σ_R returns to return units.
- Volatility treats upside and downside dispersion symmetrically.

Volatility is the standard deviation of returns. It measures dispersion, not the probability of a permanent loss.

Aster's three states imply 8.75% expected return and 19.49% volatility

State	Probability	Return	$p_s R_s$
Downturn	25%	-20%	-5.00%
Steady	50%	10%	5.00%
Expansion	25%	35%	8.75%
Total	100%		**8.75%**

$$\sigma_R = \sqrt{\sum_s p_s (R_s - 0.0875)^2} = 19.49\%.$$

Percentages enter the calculation as decimals. The table reports rounded components; the volatility uses full precision.

Appendix A2 · Full calculation

Observed returns replace probabilities with sample frequencies

$$\bar{R} = \frac{1}{T}\sum_{t=1}^{T} R_t, \qquad s_R = \sqrt{\frac{1}{T-1}\sum_{t=1}^{T}(R_t - \bar{R})^2}.$$

- T is the number of observed periods.
- $\bar{R}$ summarises the sample; it does not reveal the population mean exactly.
- $T-1$ corrects the usual sample variance for estimating $\bar{R}$ from the same data.

A sample is the observed subset of possible outcomes. A population parameter describes the unknown return-generating process.

Appendix A3 · Sample moments

Arithmetic mean and CAGR answer different questions

Measure	Calculation	Question answered
Arithmetic mean	$\bar{R} = T^{-1} \sum_t R_t$	What is the sample estimate of a one-period expected return?
CAGR	$\left[\prod_t (1 + R_t)\right]^{1/T} - 1$	What constant rate reproduces realised compound wealth?

CAGR = compound annual growth rate. It is a realised multi-period performance measure, not generally an unbiased forecast of next period's return.

Volatility lowers compound growth even when arithmetic means match

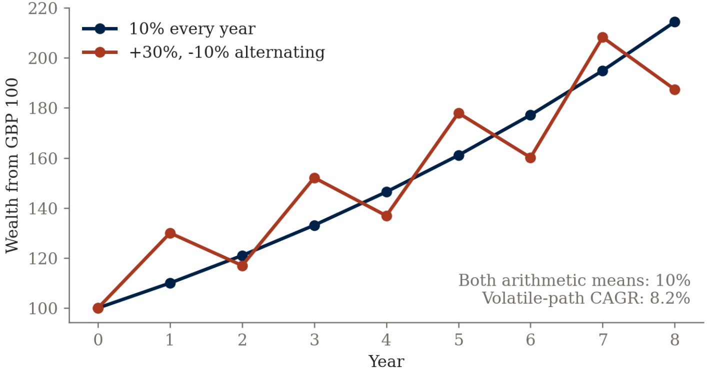

Author calculation: eight annual returns; no fees, tax, or cash flows.

Both paths average 10% arithmetically. The volatile path compounds at 8.17% because losses apply to a reduced wealth base.

The sample mean can be far less precise than the sample volatility

$$\mathrm{SE}(\bar{R}) \approx \frac{s_R}{\sqrt{T}}, \qquad 95\% \text{ interval} \approx \bar{R} \pm 1.96\ \mathrm{SE}(\bar{R}).$$

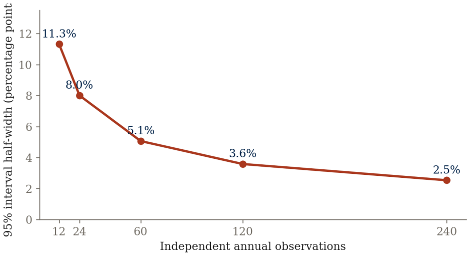

Illustration assumes independent annual returns with 20% volatility; normal approximation.

SE = standard error, the sampling uncertainty of an estimator. Independence and a stable distribution are strong assumptions for financial returns.

Appendix A5 · Confidence interval

Total volatility describes a standalone asset, not its role in a portfolio

- The same stock can be risky alone but useful beside assets that move differently.
- Expected portfolio return depends on weights; portfolio risk also depends on co-movement.
- Investors can choose combinations, so a pricing model must ask which risk remains after combination.

$$\text{next object:} \qquad \mathrm{Cov}(R_i, R_j).$$

Co-movement describes whether two returns tend to be above or below their means at the same time.

Diversify risk

Which risk survives inside a portfolio?

Move from weights and covariance to efficient portfolios and the systematic-risk floor.

Portfolio return is the value-weighted average of component returns

$$R_p = \sum_{i=1}^{N} w_i R_i, \qquad \sum_{i=1}^{N} w_i = 1.$$

- w_i is asset i's fraction of portfolio value at the start of the period.
- Expectations are linear: $\mathbb{E}[R_p] = \sum_i w_i \mathbb{E}[R_i]$.
- Risk is not generally a weighted average because returns can offset or reinforce one another.

A portfolio is a collection of claims. A long-only portfolio has $w_i \geq 0$; short positions permit negative weights.

Covariance records the direction and scale of co-movement

$$\mathrm{Cov}(R_i, R_j) = \mathbb{E}[(R_i - \mathbb{E}[R_i])(R_j - \mathbb{E}[R_j])] \,.$$

- Positive covariance: returns tend to be on the same side of their means.
- Negative covariance: one tends to be high when the other is low.
- Zero covariance: no linear co-movement; it does not imply complete independence.

Covariance is scale-dependent and measured in squared-return units, which makes direct comparisons difficult.

Correlation standardises covariance to a common scale

$$\rho_{ij} = \frac{\mathrm{Cov}(R_i, R_j)}{\sigma_i \sigma_j}, \qquad -1 \leq \rho_{ij} \leq 1.$$

- $\rho = 1$: perfect positive linear co-movement; no offsetting benefit.
- $\rho < 1$: some diversification benefit is possible.
- $\rho = -1$: one exact combination can eliminate volatility.

Correlation ρ is dimensionless. It describes linear dependence and can change across samples and market regimes.

Two-asset risk has two own-risk terms and one covariance term

$$\sigma_p^2 = w_A^2\sigma_A^2 + w_B^2\sigma_B^2 + 2w_A w_B \rho_{AB}\sigma_A\sigma_B.$$

- The first two terms scale each asset's own variance.
- The third term determines how much their movements reinforce or offset.
- Expected return needs means and weights; risk additionally needs covariance.

Portfolio variance σ_p^2 is quadratic in weights because both the size and interaction of positions matter.

Appendix A7 · Derivation Python P3 · Calculate

Lower correlation bends portfolio risk below a weighted average

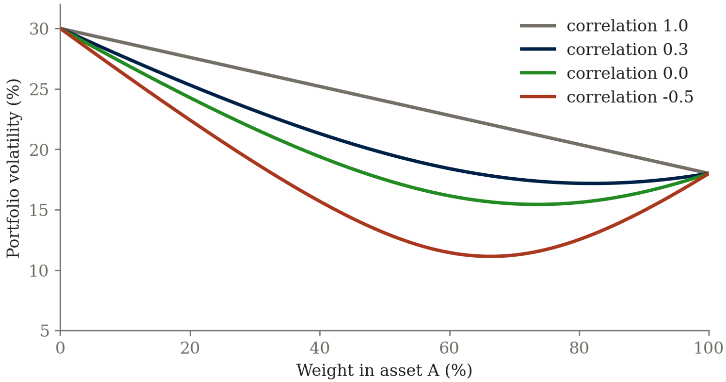

Author calculation: $\sigma_A = 18\%$, $\sigma_B = 30\%$; long-only weights.

The endpoints are individual-asset volatility. Interior curvature is created by correlation, not by changing either asset's standalone risk.

Appendix A8 · Boundary cases

A 60/40 portfolio can be less risky than the weighted-average shortcut

Input	Asset A	Asset B	Portfolio
Weight	60%	40%	100%
Expected return	8%	14%	10.4%
Volatility	18%	30%	**18.40%**

$$\rho_{AB} = 0.30, \qquad 0.6(18\%) + 0.4(30\%) = 22.8\% \quad \text{is not portfolio volatility.}$$

The correct 18.40% comes from portfolio variance. The 22.8% shortcut ignores offsetting movements.

Appendix A7 · Arithmetic

Independent firm-specific risk falls as the portfolio grows

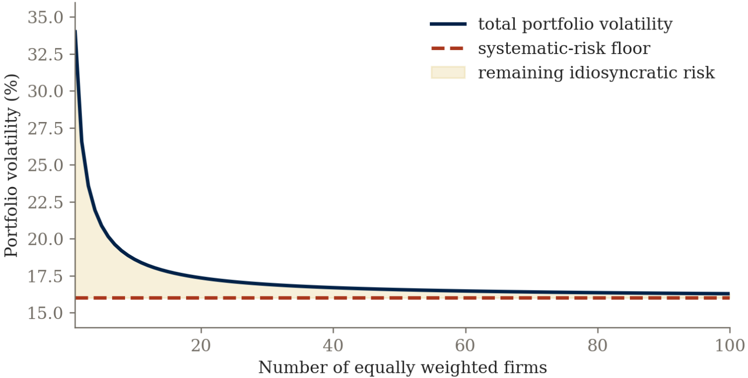

Constructed model: equal weights; 16% common volatility and 30% independent firm-specific volatility per stock.

With equal weights and independent firm-specific shocks, that variance component falls in proportion to $1/N$. Volatility falls in proportion to $1/\sqrt{N}$.

Appendix A9 · Many assets Python P4 · Simulate

A common shock creates the diversification floor

$$R_i = \underbrace{a_i + b_i F}_{\text{common component}} + \underbrace{\varepsilon_i}_{\text{firm-specific component}}.$$

- Adding firms averages out independent ε_i shocks.
- The common factor F affects many holdings together and remains.
- Diversification changes which risk is relevant; it does not make investment riskless.

A factor is a common source of return variation. The coefficient b_i measures asset i's exposure to that source.

One hundred independent loans have the same mean but one-tenth the risk

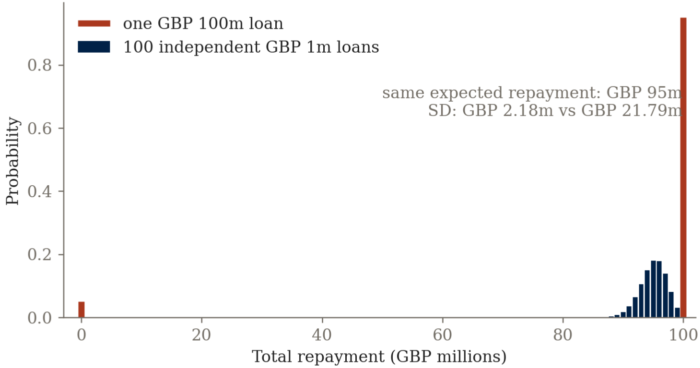

Adapted and recalculated from Berk–DeMarzo Ch. 10: 5% independent default probability; zero recovery.

Both banks expect GBP95m. Standard deviation is GBP2.18m for 100 independent GBP1m loans versus GBP21.79m for one GBP100m loan.

Appendix A10 · Binomial audit

Systematic risk moves many firms; diversifiable risk is local

Mostly systematic	Mostly diversifiable
Recession reduces broad demand	One plant is damaged
Unexpected inflation reprices discount rates	A product launch fails
Economy-wide credit conditions tighten	A key employee leaves
Broad tax or regulatory change	One supplier defaults

Systematic risk cannot be removed by holding more ordinary securities. Diversifiable, idiosyncratic, and firm-specific risk refer to the removable component.

The efficient frontier discards avoidable portfolio combinations

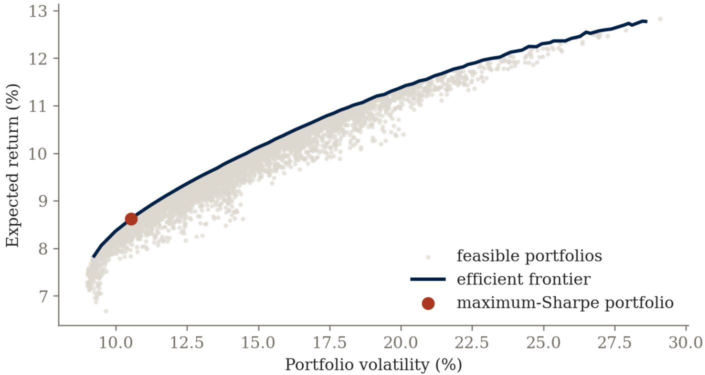

Author simulation of 30,000 long-only portfolios from four assets; declared means and covariance matrix.

An efficient portfolio offers the highest expected return for its volatility, or the lowest volatility for its expected return.

Appendix A11 · Construction

A risk-free asset creates a straight capital allocation line

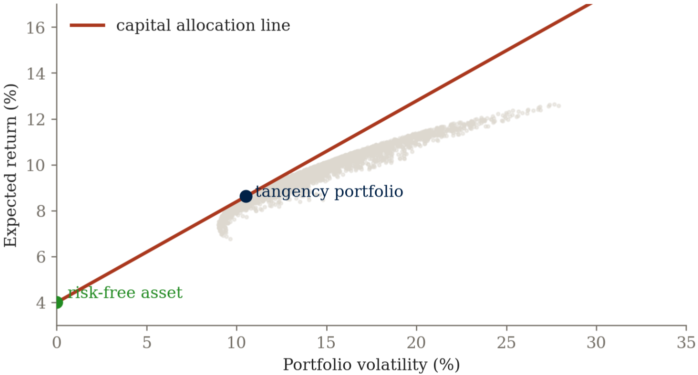

Same simulation; risk-free rate 4%; unrestricted lending or borrowing illustrated.

Capital allocation line = expected-return and volatility combinations from a risk-free asset and one risky portfolio. Borrowing extends the line beyond the risky portfolio.

Appendix A12 · Line equation

The Sharpe ratio measures excess return per unit of total portfolio risk

$$\text{Sharpe}(p) = \frac{\mathbb{E}[R_p] - r_f}{\sigma_p}.$$

- It is the slope of the capital allocation line through portfolio p.
- Among feasible risky portfolios, the steepest line offers the best expected excess return per unit of volatility.
- For a single stock added to a diversified portfolio, total volatility is no longer the relevant marginal risk.

Excess return = return above the risk-free rate r_f. The Sharpe ratio is a portfolio-performance statistic, not CAPM beta.

Price systematic risk

Why does beta, rather than total volatility, determine required return?

View each security through the diversified investor's portfolio and derive the Security Market Line.

The marginal investor can diversify before choosing one more stock

- A diversified investor does not evaluate a security as an isolated gamble.
- Firm-specific shocks can be offset across many holdings without sacrificing expected return.
- Competitive prices therefore need not reward risk that investors can remove cheaply.
- The relevant question is how the security changes risk in the investor's broad portfolio.

> priced risk $\neq$ standalone volatility.

Marginal investor = investor whose trade helps determine the market price at the margin.

The market portfolio is the benchmark diversified risky portfolio

- In CAPM, all investors combine the risk-free asset with the same market portfolio.
- The theoretical market contains every risky asset, weighted by market value.
- In practice, an investable equity index is a proxy, not the complete theoretical market.
- Benchmark choice becomes part of the estimate rather than an invisible convention.

Market portfolio = value-weighted portfolio of all risky assets in the CAPM. Market proxy = observed index used to approximate it.

Appendix A21 · Market proxy

Risk is costly when a payoff is weak in bad market states

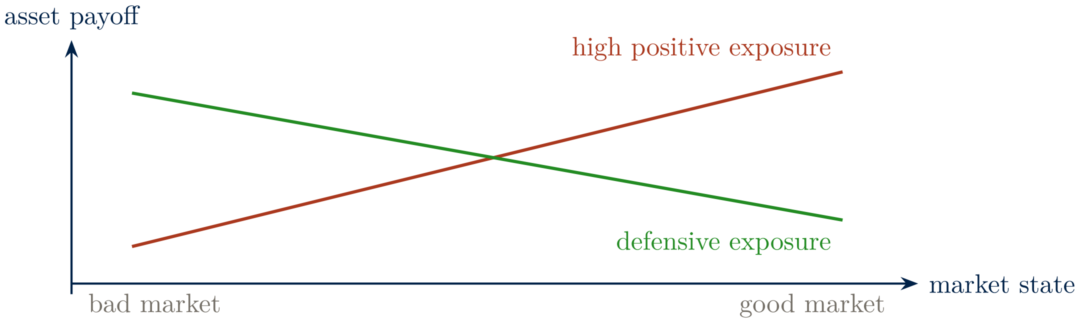

- A claim that pays poorly when the market is poor worsens already-bad portfolio outcomes.
- Investors require compensation for that covariance with market returns.

Defensive exposure tends to preserve payoff in weak market states; it need not have low standalone volatility.

Beta scales covariance with the market to market variance

$$\beta_i = \frac{\mathrm{Cov}(R_i, R_M)}{\mathrm{Var}(R_M)}.$$

- The numerator measures co-movement with market return R_M.
- The denominator expresses that co-movement per unit of market variance.
- Beta is dimensionless and measures linear market exposure, not total risk.

Equity beta β_i measures the sensitivity of a stock's return to the market return. The market's beta with itself equals one.

Appendix A13 · Covariance form

Read beta as a sensitivity, not as a promised movement

Beta	Conditional interpretation
0	No linear market exposure; return can still be volatile for other reasons
1	Same market sensitivity as the market portfolio
1.5	A 1 percentage-point market excess-return movement is associated with about 1.5 points
< 0	Tends to move against the market; rare for ordinary operating companies

The interpretation is conditional and average. A regression beta does not predict each realised return because residual shocks remain.

A two-state example exposes the slope behind beta

State	Market return	Cyclical stock	Defensive stock
Strong economy	30%	43%	-22%
Weak economy	-10%	-17%	18%
Return spread	40 points	60 points	-40 points
Beta from spread		**1.5**	**-1.0**

With two equally likely states, beta equals the stock's state-return spread divided by the market's spread. Adapted from Berk–DeMarzo Ch. 10.

Portfolio beta is the value-weighted average of component betas

$$\beta_p = \sum_{i=1}^{N} w_i \beta_i.$$

Holding	Weight	Beta	Contribution
Defensive	30%	0.50	0.15
Market-like	40%	1.00	0.40
Cyclical	30%	1.50	0.45
Portfolio	100%		**1.00**

Beta contributions add because covariance is linear in portfolio weights.

Appendix A14 · Proof

CAPM prices one unit of beta at the market risk premium

$$\mathbb{E}[R_i] = r_f + \beta_i(\mathbb{E}[R_M] - r_f).$$

- r_f: return for waiting without market risk.
- $\mathbb{E}[R_M] - r_f$: market risk premium, paid per unit of beta.
- β_i: quantity of market exposure in the security.
- Expected return rises linearly with beta in equilibrium.

Market risk premium = expected market return above the risk-free rate. CAPM required return is an expectation, not the next realised return.

Appendix A15 · Equilibrium logic

The Security Market Line maps beta into required return

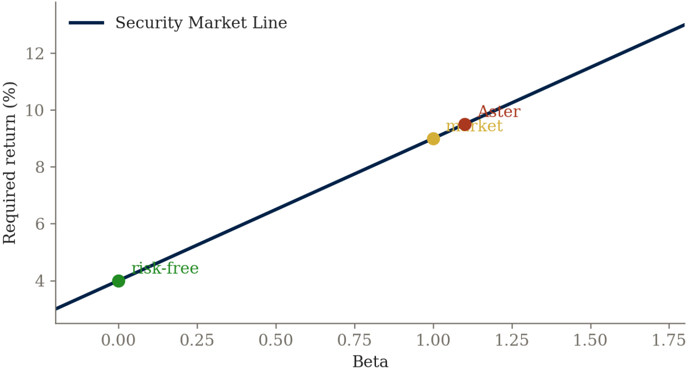

Author calculation: $r_f = 4.0\%$, market risk premium 5.0%; Aster beta 1.10.

Security Market Line (SML) = CAPM relation between beta and expected return. Its intercept is r_f; its slope is the market risk premium.

Alpha measures return relative to the model benchmark

$$\alpha_i = \mathbb{E}[R_i] - \left[r_f + \beta_i\big(\mathbb{E}[R_M] - r_f\big)\right].$$

- Positive alpha means expected return above the CAPM benchmark for estimated beta.
- Negative alpha means expected return below that benchmark.
- An estimated historical alpha can reflect skill, mispricing, omitted factors, or noise.

Alpha is model-relative abnormal return. It is not automatically evidence of arbitrage or managerial skill.

Appendix A16 · SML geometry

Equal total volatility need not imply equal required return

Security	Total volatility	Beta	CAPM return
Firm-specific biotech	35%	0.60	7.0%
Cyclical manufacturer	35%	1.40	11.0%

- The biotech's large residual risk can be diversified across holdings.
- The manufacturer's market exposure worsens broad portfolio losses.

Illustration uses $r_f = 4\%$ and a 5% market risk premium. The security labels are conceptual, not observed companies.

CAPM gives a benchmark; estimation determines what number enters the model

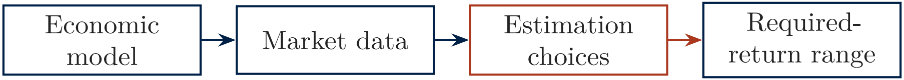

- Beta requires a sample, frequency, and market proxy.
- The risk-free rate requires currency and horizon consistency.
- The market premium is expected and therefore cannot be read directly from one price.

A model output inherits uncertainty and judgment from every input.

Estimate required return

What is observed, assumed, and uncertain?

Estimate beta transparently, then align benchmark, currency, horizon, and market premium.

The market model turns beta into a regression slope

$$R_{i,t} - R_{f,t} = \alpha_i + \beta_i\big(R_{M,t} - R_{f,t}\big) + \varepsilon_{i,t}.$$

- Left side: security excess return in period t.
- Market excess return supplies the systematic explanatory variable.
- $\varepsilon_{i,t}$ is the residual return not explained by this linear market relation.
- Ordinary least squares chooses the line that minimises squared residuals.

Ordinary least squares (OLS) is a regression method. A residual is observed return minus fitted return. Historical market-model alpha is not expected CAPM alpha automatically.

Appendix A17 · OLS

Aster's synthetic sample makes the estimated slope visible

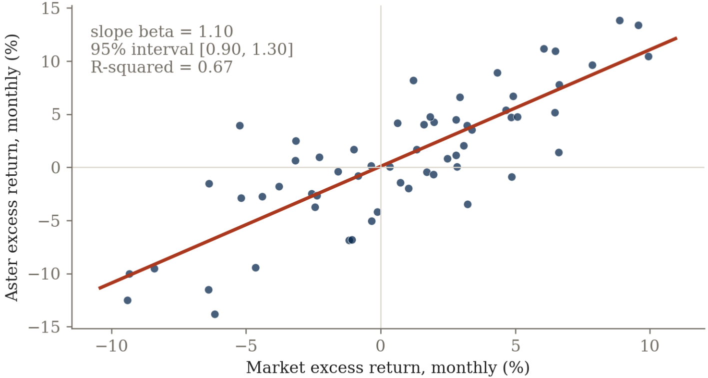

Synthetic Aster monthly returns and Kenneth French market factor, 2021:07–2026:06; 60 observations; seed 20260814.

The line summarises average linear co-movement. The vertical distance from each point to the line is that month's residual.

Python P5 · Estimate beta

The OLS slope reproduces the covariance definition of beta

$$\widehat{\beta}_i = \frac{\sum_t (x_t - \bar{x})(y_t - \bar{y})}{\sum_t (x_t - \bar{x})^2} = \frac{\widehat{\mathrm{Cov}}(x, y)}{\widehat{\mathrm{Var}}(x)},$$

- $x_t = R_{M,t} - R_{f,t}$: observed market excess return.
- $y_t = R_{i,t} - R_{f,t}$: observed security excess return.
- A steeper fitted line means greater estimated sensitivity to market movements.

A hat marks an estimate from a sample. Centring by sample means allows the regression to estimate an intercept.

A useful beta report includes uncertainty and fit

Statistic	Estimate	Interpretation
Monthly alpha	0.088%	Average intercept in the synthetic sample
Beta	1.099	Estimated market sensitivity
Standard error of beta	0.101	Sampling uncertainty under OLS assumptions
Approx. 95% interval	[0.90, 1.30]	Plausible slope range under the model
R^2	0.670	Share of sample excess-return variance explained
Observations	60	Monthly data, 2021:07–2026:06

R^2, read "R-squared," measures in-sample fit. It does not prove or disprove CAPM.

Appendix A18 · Standard error

Regression separates market-related variance from residual variance

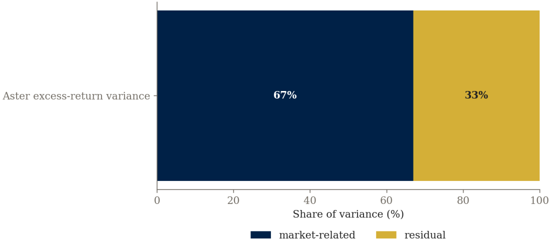

Synthetic Aster regression, 60 monthly observations; variance shares calculated from fitted and residual components.

- Market-related variation drives beta and the CAPM premium.
- Residual variation still affects a concentrated owner but can be diversified in a broad portfolio.

The decomposition is sample- and model-dependent; omitted common factors may appear in the residual.

Appendix A19 · Variance audit

The beta estimate changes when the sample window changes

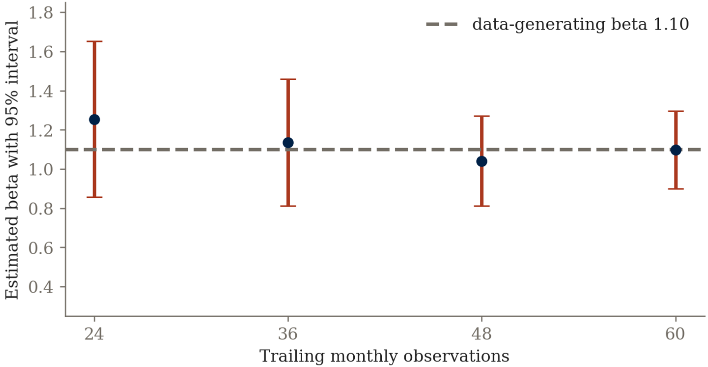

Same synthetic sample; trailing 24, 36, 48, and 60 months; approximate 95% OLS intervals.

A shorter window adapts faster but is noisier. A longer window is more precise only if the company's exposure is sufficiently stable.

Appendix A20 · Design choices

Data frequency trades more observations against more measurement problems

Frequency	Potential benefit	Main risk
Daily	Many observations; responsive	Non-synchronous trading, thin liquidity, market microstructure noise
Weekly	Compromise between count and noise	Day-of-week and alignment choices
Monthly	Cleaner alignment; common in corporate finance	Fewer observations and slower recognition of structural change

Market microstructure describes how trading rules, liquidity, bid–ask spreads, and timing affect observed prices. Frequency should be reported, not treated as a hidden default.

Appendix A20 · Frequency audit

Beta is conditional on the chosen market benchmark

- A UK equity index matches Aster's primary listing and domestic investor base.
- A global index may better represent internationally diversified shareholders.
- A sector index measures peer co-movement but is not the CAPM market portfolio.
- Thin or concentrated indices can distort both covariance and market variance.

$$\widehat{\beta}_i(\text{UK market}) \neq \widehat{\beta}_i(\text{global market}) \quad \text{in general.}$$

Benchmark risk is the sensitivity of an estimate or decision to the selected comparison portfolio.

Appendix A21 · Proxy problem

The risk-free input must match currency, horizon, and convention

- Use a GBP nominal rate for GBP nominal cash flows.
- Match duration to the valuation problem rather than copying a one-month cash rate blindly.
- Government yields may include liquidity or credit effects; an OIS curve is another market proxy.
- Record observation date and whether the rate is spot, par, or a term-structure approximation.

OIS = overnight indexed swap, a contract exchanging a fixed rate for compounded overnight rates.
"Risk-free" is a modelling approximation, not an observable label.

Historical and implied market premiums answer different questions

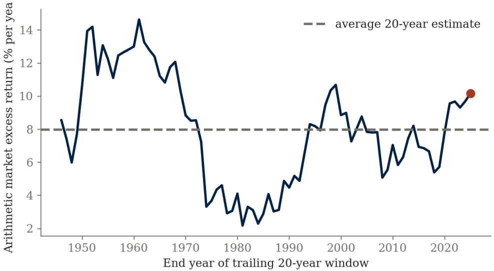

Kenneth French US market factor, July 1926–June 2026; trailing 20-year arithmetic annualised means.

- Historical premium summarises realised excess returns over a chosen sample.
- Implied premium reverses current market prices and forecasts into a forward-looking discount-rate estimate.

Equity risk premium (ERP) is expected market return above r_f. The latest 20-year historical estimate is 10.15%; it is not our 5% forward input.

Appendix A22 · ERP methods

A current implied premium is a dated benchmark, not a universal constant

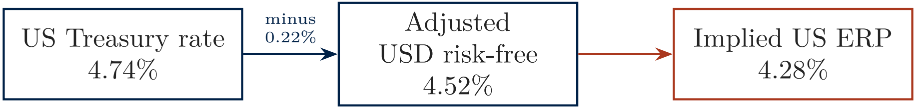

- The observation is USD, market-wide, and dated 1 August 2026.
- Aster's 5.0% GBP market premium is a declared decision input, not copied from this estimate.
- A professional report explains any difference in market, currency, date, or method.

Aswath Damodaran, NYU Stern implied equity risk premium update, observed 1 August 2026.

Implied ERP solves for the discount-rate premium consistent with current index level, expected cash distributions, and growth assumptions.

Industries with similar total risk can carry different market exposure

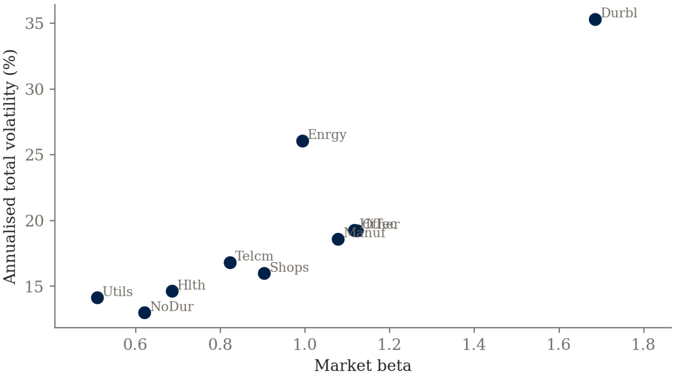

Kenneth French ten value-weighted US industry portfolios and market factor, 2006:07–2026:06; monthly returns. Annualised volatility equals monthly standard deviation times $\sqrt{12}$. Industry abbreviations and construction are documented in Appendix A23.

Appendix A23 · Data source

Report an estimate package rather than a naked beta

Report field	Aster case estimate
Security and benchmark	Synthetic Aster excess return versus Kenneth French US market factor
Sample	Monthly, 2021:07–2026:06, 60 observations
Point estimate	Beta 1.099; monthly alpha 0.088%
Uncertainty	Beta SE 0.101; approximate 95% interval [0.90, 1.30]
Fit	$R^2 = 0.670$; residual monthly SD 3.61%
Judgment	UK/global proxy and post-listing business mix require separate validation

The synthetic sample demonstrates estimation mechanics. It is not evidence about a real company and cannot substitute for a real Aster data audit.

Apply CAPM

Does market risk support Aster's 9.5% cost of equity?

Translate the business risk story into inputs, sensitivity, valuation, and a committee recommendation.

Begin Aster's estimate with its business exposure, not a database beta

- Merchant transaction volumes rise and fall with consumer spending.
- Payment acceptance has operating leverage through fixed technology and compliance costs.
- Settlement and fraud losses can add firm-specific volatility.
- Recurring software fees and diversified merchants can moderate cyclicality.
- Leverage transfers more operating risk to each pound of equity.

Operating leverage means fixed operating costs amplify the effect of revenue changes on operating profit. Financial leverage means debt amplifies the risk borne by equity.

Aster's base inputs are declared in one currency and on one date

Input	Base case	Status
GBP nominal risk-free rate	4.0%	Illustrative input for valuation date
Market risk premium	5.0%	Forward-looking illustrative input
Equity beta	1.10	Rounded estimate supported by synthetic regression
Cost of equity	9.5%	CAPM output

These inputs are internally coherent but are not live investment recommendations. A real report would replace every assumption with a documented market and company estimate.

CAPM reproduces Day 3's 9.5% cost of equity

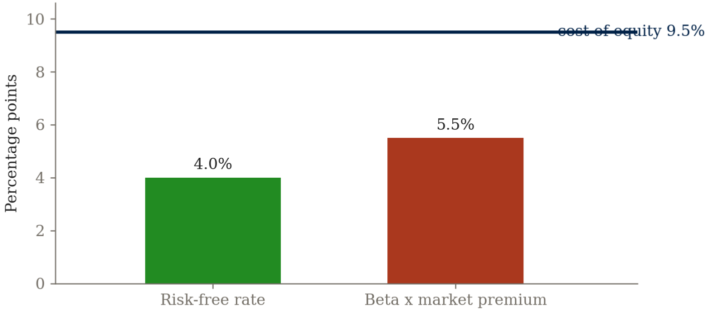

$$r_e = 4.0\% + 1.10(5.0\%) = 9.5\%.$$

The beta-scaled premium contributes 5.5 percentage points. "9.5%" is more precise than the inputs justify; the result is reported to one decimal place.

Python P6 · CAPM

A beta revision changes the hurdle rate one-for-one with the premium

Beta	Cost of equity	Change from base
0.90	8.5%	-1.0 point
1.10	**9.5%**	Base
1.30	10.5%	+1.0 point

$$\frac{\partial r_e}{\partial \beta} = \mathbb{E}[R_M] - r_f = 5.0\%.$$

A 0.20 beta change moves the required return by 1.0 percentage point when the market premium is 5.0%.

The market-premium choice matters more for higher-beta equity

Market premium	Aster beta	Cost of equity
4.0%	1.10	8.4%
5.0%	1.10	**9.5%**
6.0%	1.10	10.6%

$$\frac{\partial r_e}{\partial \mathrm{MRP}} = \beta = 1.10.$$

MRP = market risk premium. A one-point premium revision moves Aster's required return by 1.10 points at beta 1.10.

Joint sensitivity shows the input combinations hidden by one rate

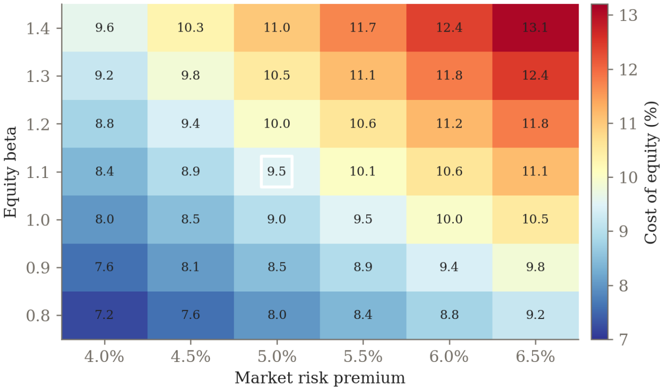

Author calculation; GBP risk-free rate fixed at 4.0%. White square marks beta 1.10 and 5.0% premium.

The table is a decision surface, not a probability distribution. A professional range should be tied to evidence for each input.

Appendix A24 · Sensitivity table

Required-return uncertainty moves the Day 3 valuation materially

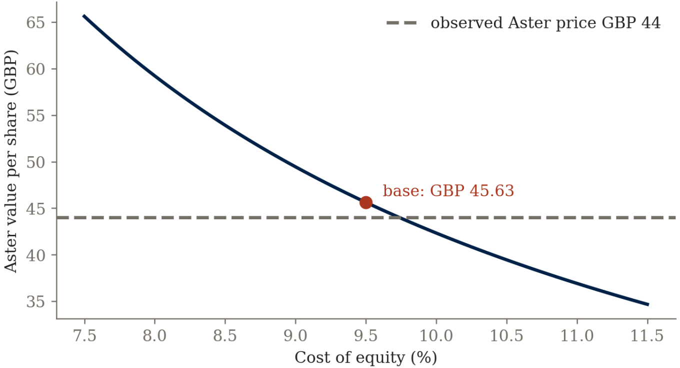

Rebuilt Day 3 base FCFF forecast; 80% equity, 20% debt, 4.5% after-tax debt cost, 3.0% terminal growth.

FCFF = free cash flow to the firm. Cost of equity enters WACC through the capital weights; it is not used to discount FCFF directly. Base value is GBP45.63 per share.

Appendix A26 · WACC audit

The company cost of capital is not every project's hurdle rate

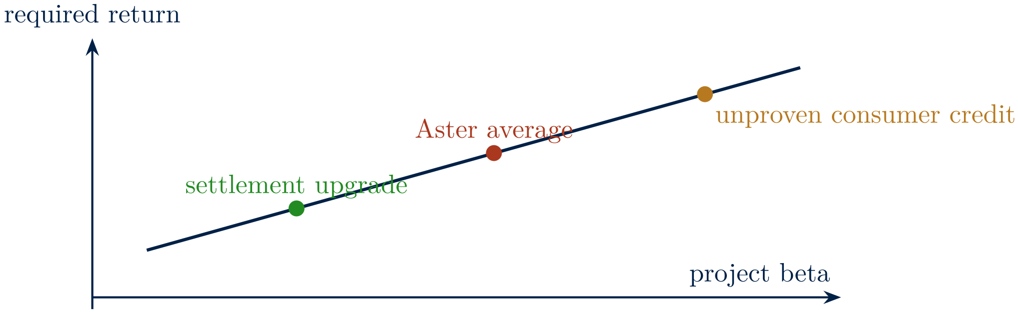

- Discount a project's cash flow at a rate for that project's systematic risk.
- Using one company-wide hurdle can reject safe projects and accept risky ones.

A project beta measures the market exposure of project cash flows, conceptually independent of the financing used to fund it.

Financial leverage raises equity beta above asset beta

$$\beta_E = \beta_A + \frac{D}{E}(\beta_A - \beta_D) \quad \approx \quad \beta_A\left(1 + \frac{D}{E}\right) \text{ if } \beta_D \approx 0.$$

- Asset beta β_A reflects operating risk borne by all capital providers.
- Equity beta β_E concentrates that risk after debt has a prior claim.
- Comparable-company betas must be unlevered and relevered when capital structures differ materially.

D/E is the market-value debt-to-equity ratio; β_D is debt beta. Tax-adjusted conventions are common in practice and must be stated.

Appendix A25 · Unlevering

Cost of equity enters WACC beside the after-tax cost of debt

$$\text{WACC} = \frac{E}{D+E} r_e + \frac{D}{D+E} r_d (1 - T_c).$$

Component	Weight	Contribution
Equity at 9.5%	80%	7.6%
Debt after tax at 4.5%	20%	0.9%
WACC	100%	**8.5%**

WACC = weighted average cost of capital; T_c is the corporate tax rate. Use market-value weights and apply WACC only to operating cash flows of comparable risk.

Appendix A26 · Full bridge

Recommendation: retain 9.5% as a base, but manage an 8.5–10.5% range

- **Base:** 9.5% follows transparently from $4.0\% + 1.10 \times 5.0\%$.
- **Range:** 8.5–10.5% represents beta 0.90–1.30 at the base premium; it is a decision range, not a statistical interval for all inputs.
- **Validation:** compare UK and global benchmarks, inspect listed payment peers, and test structural breaks.
- **Project use:** adjust away from the company rate when project systematic risk differs.
- **Update trigger:** revise after material leverage, business-mix, benchmark, or market-premium changes.

A defendable recommendation states model, inputs, evidence, uncertainty, scope, and update rule.

Judge the model

When is CAPM useful, and when should it be challenged?

Separate assumptions, empirical evidence, and decision fitness before choosing a richer model.

CAPM's clean result rests on demanding assumptions

- Investors care only about expected return and variance over one common horizon.
- They share beliefs about returns, variances, and covariances.
- They can trade without taxes, transaction costs, or binding short-sale constraints.
- They can borrow or lend at the same risk-free rate.
- All relevant assets are tradable and the market portfolio is observable.

A model can remain useful when assumptions are approximate if its output is stable, interpretable, and appropriate for the decision.

An empirical CAPM test is a joint test of model and market proxy

- CAPM predicts expected returns relative to the true market portfolio.
- Researchers observe realised returns and substitute an equity-index proxy.
- Rejection may reflect the model, the proxy, estimated betas, changing premiums, or sampling noise.
- Failure to reject does not prove the assumptions are literally true.

$$\text{test result} = f(\text{pricing model}, \text{market proxy}, \text{estimation design}, \text{sample}).$$

This is the Roll critique: the theoretical market portfolio is broader than any observed stock index, so CAPM is difficult to test in isolation.

Appendix A27 · Test design

One realised return cannot validate or reject an expected-return model

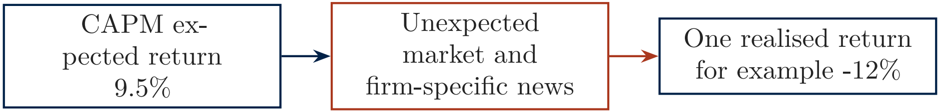

- Expectations are conditional means; realised returns contain large shocks.
- Testing requires many assets or periods and an explicit statistical design.

A model can assign a 9.5% expectation while allowing a wide distribution of realised outcomes around it.

Realised industry returns do not sit exactly on the CAPM benchmark

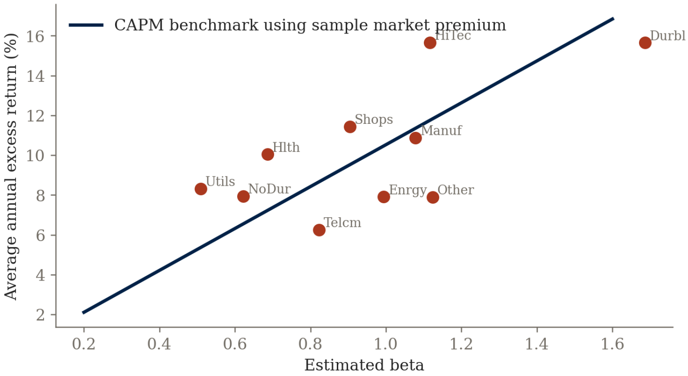

Kenneth French ten industries and market factor, 2006:07–2026:06; monthly estimates annualised arithmetically.

Points above or below the line are realised sample deviations, not automatically persistent alpha. Betas and the market premium are estimated from the same 240-month sample.

Appendix A23 · Data construction

Size and value patterns motivate richer descriptions of average returns

- Small-cap and high book-to-market portfolios have historically shown return patterns not fully captured by market beta.
- These patterns may reflect additional risks, behavioural mispricing, changing samples, or data-mining.
- A factor that describes returns does not automatically identify a causal risk premium.
- Added factors increase flexibility and estimation burden.

Book-to-market = accounting book equity divided by market equity. Small minus big (SMB) and high minus low (HML) are Fama–French factor returns.

A multi-factor model extends the benchmark, but every factor needs a reason

$$R_i - r_f = \alpha_i + \beta_{iM}(R_M - r_f) + \beta_{iS}\, SMB + \beta_{iH}\, HML + \varepsilon_i.$$

- Additional betas describe sensitivity to additional factor returns.
- The model can improve historical fit without supplying a unique required return.
- For corporate hurdle rates, CAPM often remains the transparent starting benchmark.

A factor loading is a regression sensitivity. A priced factor additionally requires a defensible expected premium.

Efficient markets and CAPM are different claims

Claim	Core question	Possible failure
Market efficiency	Do prices incorporate available information?	Predictable mispricing after costs
CAPM	Does market beta alone explain expected-return differences?	Other priced risks or model misspecification

A market can be informationally efficient under a pricing model richer than CAPM. A surprising realised return is consistent with both.

Choose the simplest model that is fit for the decision

Approach	Use when	Main discipline
CAPM	Transparent company or project hurdle; limited data	Audit beta, risk-free rate, and premium
Comparable beta	Company is unlisted or history is unstable	Match business risk; unlever and relever
Multi-factor model	Portfolio attribution or strong factor mandate	Justify factors and expected premiums
Build-up judgment	Market data are weak or special risks dominate	Avoid double-counting diversifiable risk

Model complexity is useful only when it changes the decision more reliably than it adds estimation error.

Six failure checks catch most required-return mistakes

1. **Date:** are all market inputs observed or estimated for the same decision date?
2. **Currency:** does the risk-free rate match cash-flow currency and inflation?
3. **Benchmark:** is beta tied to a defensible market proxy?
4. **Leverage:** is equity beta adjusted for financing differences?
5. **Project:** does the hurdle reflect project rather than company-average risk?
6. **Uncertainty:** are intervals, sensitivities, and update triggers disclosed?

A spreadsheet cell containing one rate is not an audit trail.

Day 4 synthesis: price only the risk that survives diversification

- Return and volatility describe outcomes; covariance describes portfolio interaction.
- Diversification removes firm-specific risk, leaving market exposure to be priced.
- CAPM maps beta to expected return, but every input is estimated or assumed.
- Aster's 9.5% is a transparent base case, not an immutable truth.
- Day 5 applies pricing and risk management to forwards and futures.

Technical appendix Python practical

References and data sources

- Brealey, Myers, Allen and Edmans, *Principles of Corporate Finance*, Ch. 8; user-provided lecture slides.
- Berk and DeMarzo, *Corporate Finance*, Ch. 10; accompanying slides and solutions.
- Markowitz (1952), "Portfolio Selection"; Sharpe (1964), "Capital Asset Prices"; Lintner (1965), valuation of risky assets.
- Jensen (1968), mutual-fund performance; Roll (1977), critique of CAPM tests.
- Fama and French (1993), common risk factors in stock and bond returns.
- MIT OpenCourseWare, 15.401 Finance Theory I, CAPM lecture.
- Kenneth R. French Data Library, US research factors and ten value-weighted industry portfolios, CRSP 202606 vintage.
- Aswath Damodaran, NYU Stern, current data and implied equity risk premium pages; observed 1 August 2026.
- Aster Payments returns and case assumptions are constructed data; calculations frozen 14 August 2026.

Direct URLs, units, samples, transformations, and file names are documented in Appendix A28 and the local metadata file.

Using these slides

- Course materials for the Computational Finance / FinTech Summer School.
- Course materials and updates: fatih.ai.
- Personal study by enrolled students is permitted.
- Do not copy, distribute, publish, or adapt these slides without asking the author first.
- Third-party names, data, and extracts remain subject to their original rights and terms.

Technical appendix map

← **Return to Day 4 synthesis**

Python Practical map

A1. The holding-return identity follows from wealth change

$$W_{t+1} = P_{t+1} + D_{t+1}, \qquad R_{t+1} = \frac{W_{t+1} - P_t}{P_t}.$$

$$R_{t+1} = \frac{P_{t+1} - P_t}{P_t} + \frac{D_{t+1}}{P_t} = \text{capital-gain yield} + \text{income yield}.$$

- Reinvesting an interim distribution requires tracking the reinvestment date and price.
- Index total-return series perform this reinvestment convention; price indices do not.

The formula uses simple return. A log return is $\ln(1 + R)$ and adds across time, but portfolio log returns do not add across assets.

← **Return to holding return**

Appendix map

A2. Aster's state moments can be audited row by row

State	p_s	R_s	$R_s - 0.0875$	$p_s(R_s - 0.0875)^2$
Downturn	0.25	-0.20	-0.2875	0.020664
Steady	0.50	0.10	0.0125	0.000078
Expansion	0.25	0.35	0.2625	0.017227
Total	1.00			**0.037969**

$$\mathbb{E}[R] = 0.0875, \qquad \sigma_R = \sqrt{0.03796875} = 0.19486.$$

Variance uses squared decimal returns. Multiplying a decimal variance by 10,000 expresses it in squared percentage points.

← **Return to state example**

Appendix map

A3. Sample variance loses one degree of freedom

$$\sum_{t=1}^{T}(R_t - \bar{R}) = 0.$$

- Once $T-1$ deviations and $\bar{R}$ are known, the final deviation is fixed.
- Dividing the squared deviations by $T-1$ makes the usual sample variance unbiased under independent, identically distributed sampling.
- Financial returns often violate identical-distribution and independence assumptions; the formula remains descriptive, but inference needs care.

Degrees of freedom count independent pieces of information remaining after parameters are estimated.

← **Return to sample statistics**

Appendix map

A4. CAGR is the geometric mean of gross returns

$$\frac{W_T}{W_0} = \prod_{t=1}^{T}(1 + R_t) \quad \Longrightarrow \quad \mathrm{CAGR} = \left(\frac{W_T}{W_0}\right)^{1/T} - 1.$$

$$\frac{1}{T}\sum_{t=1}^{T} \ln(1 + R_t) = \ln(1 + \mathrm{CAGR}).$$

- With constant 10% returns, arithmetic mean and CAGR both equal 10%.
- With alternating +30% and -10%, the arithmetic mean is 10% but CAGR is 8.17%.

Gross return is $1 + R_t$. The geometric mean is defined only when terminal wealth remains positive.

← **Return to arithmetic versus CAGR**

Appendix map

A5. A mean confidence interval depends on a sampling model

$$\frac{\bar{R} - \mu}{s_R/\sqrt{T}} \sim t_{T-1} \quad \Longrightarrow \quad \mu \in \bar{R} \pm t_{0.975,T-1} \frac{s_R}{\sqrt{T}}.$$

- With $s_R = 20\%$ and $T = 100$, the normal-approximation half-width is 3.92 percentage points.
- Serial correlation changes the standard error; overlapping multi-period returns create serial dependence mechanically.
- A confidence interval concerns repeated-sample behaviour under assumptions, not a 95% probability that a fixed parameter lies inside this one interval.

$t_{0.975,T-1}$ is the 97.5th percentile of a Student t distribution with $T - 1$ degrees of freedom.

← **Return to mean uncertainty**

Appendix map

A6. Covariance is a probability-weighted product of deviations

$$\operatorname{Cov}(R_A, R_B) = \sum_s p_s (R_{A,s} - \mathbb{E}[R_A])(R_{B,s} - \mathbb{E}[R_B]).$$

- Same-sign deviations contribute positively.
- Opposite-sign deviations contribute negatively.
- Multiplying by p_s weights each state by its likelihood.
- $\operatorname{Cov}(R_A, R_A) = \operatorname{Var}(R_A)$.

Covariance is bilinear: $\operatorname{Cov}(aX + bY, Z) = a \operatorname{Cov}(X, Z) + b \operatorname{Cov}(Y, Z)$. This property drives portfolio variance and beta aggregation.

← **Return to covariance**

Appendix map

A7. Expanding the squared portfolio return yields two-asset variance

$$R_p - \mathbb{E}[R_p] = w_A(R_A - \mathbb{E}[R_A]) + w_B(R_B - \mathbb{E}[R_B]).$$

$$\operatorname{Var}(R_p) = w_A^2\sigma_A^2 + w_B^2\sigma_B^2 + 2w_A w_B \operatorname{Cov}(R_A, R_B).$$

$$\operatorname{Cov}(R_A, R_B) = \rho_{AB}\sigma_A\sigma_B.$$

The factor 2 appears because the expansion contains both $w_A w_B$ cross-products. For the 60/40 example, $\sigma_p = 18.3957\%$.

← **Return to two-asset example**

Appendix map

A8. Correlation boundary cases explain the curve's shape

- $\rho = 1$: $\sigma_p = |w_A\sigma_A + w_B\sigma_B|$; risk is linear for long-only weights.
- $\rho = -1$: $\sigma_p = |w_A\sigma_A - w_B\sigma_B|$; choose $w_A = \sigma_B/(\sigma_A + \sigma_B)$ for zero volatility.
- $-1 < \rho < 1$: risk is curved and positive, with an interior minimum possible.
- Correlation can rise during stress, reducing diversification precisely when it is most valuable.

The zero-risk combination at $\rho = -1$ is a mathematical boundary case, not a stable empirical promise.

← **Return to correlation curves**

Appendix map

A9. Equal weighting isolates the systematic-risk floor

$$\sigma_p^2 = \sigma^2 \left[\rho + \frac{1-\rho}{N}\right] \quad \text{for equal volatility and common pairwise correlation.}$$

$$\lim_{N\to\infty} \sigma_p^2 = \rho\sigma^2.$$

- The $(1-\rho)/N$ component vanishes as holdings increase.
- The common $\rho\sigma^2$ component remains.

The common-correlation model is a simplifying assumption; a covariance matrix allows heterogeneous volatilities and correlations.

← **Return to diversification curve**

Appendix map

A10. Independent loans reduce standard deviation through the square root of count

- One GBP100m loan pays GBP100m with probability 0.95 and zero with probability 0.05.
- $\mathbb{E}[X] = 95$, $\mathrm{SD}(X) = 100\sqrt{0.95(0.05)} = 21.794$ million.
- For 100 independent GBP1m loans, repayment count $Y \sim \mathrm{Binomial}(100, 0.95)$.

$$\mathbb{E}[Y] = 100(0.95) = 95, \qquad \mathrm{SD}(Y) = \sqrt{100(0.95)(0.05)} = 2.179.$$

Binomial distribution = number of successes in independent trials with a common success probability. Correlated defaults would reduce this diversification benefit.

← **Return to bank-loan case**

Appendix map

A11. The efficient frontier solves a constrained variance problem

$$\min_{\boldsymbol{w}} \ \boldsymbol{w}^\top \boldsymbol{\Sigma} \boldsymbol{w} \quad \text{subject to} \quad \boldsymbol{w}^\top \boldsymbol{\mu} = \bar{\mu}, \quad \boldsymbol{w}^\top \mathbf{1} = 1.$$

- $\boldsymbol{w}$: portfolio-weight vector.
- $\boldsymbol{\mu}$: expected-return vector.
- $\boldsymbol{\Sigma}$: covariance matrix.
- Vary target return $\bar{\mu}$ to trace minimum-variance portfolios; retain the upper efficient branch.

A covariance matrix stores every variance on its diagonal and every pairwise covariance off the diagonal. The lecture figure imposes non-negative weights.

← **Return to efficient frontier**

Appendix map

A12. The capital allocation line inherits its slope from one risky portfolio

- Put fraction y in risky portfolio p and $1 - y$ in the risk-free asset.

$$\mathbb{E}[R_c] = r_f + y\big(\mathbb{E}[R_p] - r_f\big), \qquad \sigma_c = |y|\sigma_p.$$

$$\mathbb{E}[R_c] = r_f + \frac{\mathbb{E}[R_p] - r_f}{\sigma_p}\sigma_c \quad \text{for } y \geq 0.$$

The coefficient on σ_c is portfolio p's Sharpe ratio. Borrowing corresponds to $y > 1$; holding some risk-free asset corresponds to $0 < y < 1$.

← **Return to capital allocation line**

Appendix map

A13. Beta converts market covariance into a unit-free exposure

$$\mathrm{Cov}(R_i, R_M) = \beta_i \,\mathrm{Var}(R_M).$$

- $\beta_M = \mathrm{Var}(R_M) / \mathrm{Var}(R_M) = 1.$
- A risk-free payoff has zero return covariance and therefore beta zero.
- If $R_i = a_i + b_i R_M + \varepsilon_i$ and $\mathrm{Cov}(\varepsilon_i, R_M) = 0$, then $\beta_i = b_i$.

Beta depends on the return horizon, market proxy, and sample used to estimate covariance and variance.

← **Return to beta definition**

Appendix map

A14. Portfolio beta follows from covariance linearity

$$\beta_p = \frac{\mathrm{Cov}(\sum_i w_i R_i, R_M)}{\mathrm{Var}(R_M)} = \frac{\sum_i w_i \,\mathrm{Cov}(R_i, R_M)}{\mathrm{Var}(R_M)} = \sum_i w_i \beta_i.$$

- No correlation terms between holdings are needed for portfolio beta.
- Those covariance terms remain essential for total portfolio variance.

Linearity makes beta easy to aggregate across positions, but an estimated beta still inherits each component's estimation error.

← **Return to portfolio beta**

Appendix map

A15. The market calibrates the price of covariance

- Mean–variance equilibrium makes expected excess return proportional to market covariance.

$$
\begin{aligned}
\mathbb{E}[R_i] - r_f &= \lambda \operatorname{Cov}(R_i, R_M), \\
\lambda &= \frac{\mathbb{E}[R_M] - r_f}{\operatorname{Var}(R_M)}, \\
\therefore \quad \mathbb{E}[R_i] - r_f &= \frac{\operatorname{Cov}(R_i, R_M)}{\operatorname{Var}(R_M)} (\mathbb{E}[R_M] - r_f) = \beta_i (\mathbb{E}[R_M] - r_f).
\end{aligned}
$$

- The second line calibrates λ by applying the first relation to the market portfolio.

This is an equilibrium argument, not an identity; it rests on CAPM's opportunity-set and investor assumptions.

← Return to CAPM

Appendix map

A16. Alpha is the vertical distance from the Security Market Line

$$\alpha_i = \mathbb{E}[R_i] - r_f - \beta_i(\mathbb{E}[R_M] - r_f).$$

- On the SML: $\alpha_i = 0$.
- Above the SML: expected return is high relative to CAPM beta.
- Below the SML: expected return is low relative to CAPM beta.
- In equilibrium CAPM predicts zero expected alpha, but realised sample alpha is noisy.

Do not confuse the SML, which uses beta, with the capital allocation line, which uses total portfolio volatility.

← **Return to alpha**

Appendix map

A17. OLS chooses the intercept and beta that minimise squared residuals

$$\min_{a,b} \sum_{t=1}^{T} (y_t - a - bx_t)^2.$$

$$\widehat{b} = \frac{\sum_t (x_t - \bar{x})(y_t - \bar{y})}{\sum_t (x_t - \bar{x})^2}, \qquad \widehat{a} = \bar{y} - \widehat{b}\bar{x}.$$

- At the optimum, residuals sum to zero and are sample-uncorrelated with x_t.
- For beta estimation, x_t and y_t are market and security excess returns.

OLS estimates a conditional linear relation. It does not establish that market movements cause every fitted security movement.

← **Return to market model**

Appendix map

A18. Beta precision rises with market variation and falls with residual noise

$$\mathrm{SE}(\widehat{\beta}) = \sqrt{\frac{\widehat{\sigma}_{\varepsilon}^{2}}{\sum_t (x_t - \bar{x})^2}}, \qquad \widehat{\sigma}_{\varepsilon}^{2} = \frac{\sum_t \widehat{\varepsilon}_t^{2}}{T-2}.$$

- More observations usually increase the denominator.
- Greater market-return dispersion identifies the slope more clearly.
- Greater residual variation widens the interval.
- Serial correlation or changing variance requires robust inference.

Aster's synthetic 60-month estimate is 1.099 with SE 0.101 and approximate 95% interval [0.90, 1.30].

← **Return to regression report**

Appendix map

A19. The market model decomposes variance under zero residual covariance

$$y_t = \alpha + \beta x_t + \varepsilon_t, \qquad \operatorname{Cov}(x_t, \varepsilon_t) = 0.$$

$$\operatorname{Var}(y) = \beta^2 \operatorname{Var}(x) + \operatorname{Var}(\varepsilon).$$

- $R^2 = \beta^2 \operatorname{Var}(x) / \operatorname{Var}(y)$ in a one-factor regression with an intercept.
- Aster's sample attributes approximately 67% of excess-return variance to the fitted market component and 33% to residual variation.

This is a statistical decomposition under the fitted model, not a claim that every residual shock is diversifiable in every portfolio.

← **Return to risk decomposition**

Appendix map

A20. Window and frequency choices change the estimand and its noise

Choice	What it gains	What it risks
Short window	Current business mix	Wide interval; event sensitivity
Long window	More observations	Stale leverage or operations
Daily data	Statistical count	Thin trading; non-synchronous prices
Monthly data	Cleaner alignment	Few observations; slow adaptation
Raw beta	Direct sample estimate	Mean reversion ignored
Adjusted beta	Shrinkage toward benchmark	Adjustment may hide real structure

An estimand is the population quantity the procedure aims to estimate. Changing the sample design can change both the estimand and estimator precision.

← **Return to window sensitivity**

Appendix map

A21. A market-proxy audit records economic and statistical fit

- **Coverage:** domestic equity, global equity, or broader risky wealth?
- **Currency:** are security, index, and risk-free returns expressed consistently?
- **Weighting:** value-weighted or equal-weighted benchmark?
- **Investability:** can the marginal investor hold the proxy at reasonable cost?
- **Concentration:** do a few firms or sectors dominate market variance?
- **Robustness:** does the conclusion survive a plausible alternative proxy?

CAPM's theoretical market includes human capital and other non-traded wealth; observed equity indices are practical proxies.

← **Return to benchmark choice**

Appendix map

A22. Historical and implied equity premiums use different information

- Historical arithmetic premium:

$$\widehat{\mathrm{ERP}}_{\mathrm{hist}} = \frac{1}{T}\sum_{t=1}^{T}(R_{M,t} - R_{f,t}).$$

- Implied premium solves a market valuation equation for the discount rate k:

$$P_M = \sum_{t=1}^{H}\frac{\mathbb{E}[CF_t]}{(1+k)^t} + \frac{\mathbb{E}[TV_H]}{(1+k)^H}, \qquad \mathrm{ERP}_{\mathrm{impl}} = k - r_f.$$

Historical estimates are sample-sensitive; implied estimates are forecast- and model-sensitive. Neither is an observed constant.

← **Return to market premium**

Appendix map

A23. The French sample is frozen, labelled, and reproducible

Field	Construction
Provider	Kenneth R. French Data Library
Database vintage	CRSP 202606; files created August 2026
Market series	`Mkt-RF`; value-weighted US market excess return
Risk-free series	One-month T-bill return; source convention documented by provider
Industries	Ten value-weighted industry portfolios; portfolios formed at June end
Lecture sample	2006:07–2026:06; 240 aligned monthly observations
Units	Provider percent returns converted to decimal simple returns in local CSV
Missing codes	Provider flags -99.99 and -999; none retained in aligned sample

CRSP = Center for Research in Security Prices. The provider changed its underlying monthly-return format from FIZ to CIZ beginning with the January 2025 release; the current files use the provider's documented current construction.

← **Return to industry evidence**

Appendix map

A24. CAPM sensitivity has simple partial derivatives

$$r_e = r_f + \beta \, \mathrm{MRP}\,.$$

$$\frac{\partial r_e}{\partial r_f} = 1, \qquad \frac{\partial r_e}{\partial \beta} = \mathrm{MRP}, \qquad \frac{\partial r_e}{\partial \, \mathrm{MRP}} = \beta.$$

- At beta 1.10 and MRP 5.0%, a 0.10 beta change moves r_e by 0.50 points.
- A 0.50-point MRP change moves r_e by 0.55 points.
- Simultaneous changes interact through $\Delta\beta \, \Delta \, \mathrm{MRP}$.

MRP = market risk premium. Partial derivatives vary one input while holding the others fixed.

← **Return to joint sensitivity**

Appendix map

A25. Unlevering separates operating exposure from financing exposure

$$\beta_A = \frac{E}{D+E}\beta_E + \frac{D}{D+E}\beta_D \quad \Longrightarrow \quad \beta_E = \beta_A + \frac{D}{E}(\beta_A - \beta_D).$$

- Estimate peer equity betas.
- Unlever each peer using market values and an explicit debt-beta convention.
- Take a robust central asset beta for comparable operations.
- Relever to the target capital structure.

The Hamada approximation $\beta_E \approx \beta_A[1 + (1 - T_c)D/E]$ incorporates a particular tax-shield convention and near-zero debt beta; state the convention used.

← **Return to leverage and beta**

Appendix map

A26. Aster's WACC and valuation link can be audited directly

$$\text{WACC} = 0.80(9.5\%) + 0.20(4.5\%) = 8.5\%.$$

- Rebuild the Day 3 base FCFF forecast and discount at this WACC.
- With 3.0% terminal growth, the audited value is GBP45.63 per share.
- Holding debt assumptions fixed, cost of equity 8.5% implies GBP53.92; 10.5% implies GBP39.46.
- The sensitivity reflects both explicit cash-flow discounting and terminal value.

The 4.5% debt input is after tax. A pre-tax debt cost must be multiplied by $1 - T_c$ in the WACC formula.

← **Return to value sensitivity**

Appendix map

A27. Cross-sectional CAPM tests face a joint-hypothesis problem

$$\bar{R}_i - r_f = \gamma_0 + \gamma_1 \widehat{\beta}_i + u_i.$$

- CAPM predicts $\gamma_0 = 0$ and $\gamma_1 = \mathbb{E}[R_M] - r_f$.
- Betas are generated regressors estimated with error in a first stage.
- Realised sample means are noisy proxies for expected returns.
- The chosen market index may be mean–variance inefficient even if the unobserved true market is efficient.

A generated regressor is an explanatory variable estimated rather than observed. Inference should account for its estimation error.

← **Return to joint test**

Appendix map

A28. Provenance distinguishes observed, synthetic, and assumed inputs

Object	Status	Local file	Source / construction
Market factors	Observed	`factors.csv`	Kenneth French; CRSP 202606
Industry returns	Observed	`industries.csv`	Kenneth French; ten value-weighted portfolios
Aster returns	Synthetic	`aster.csv`	Seed 20260814; target beta 1.10
Aster CAPM inputs	Assumed	`metadata.json`	GBP $r_f = 4\%$, MRP 5%, beta 1.10
Aster valuation	Calculated	Same metadata JSON	Rebuilt Day 3 FCFF base case
Implied US ERP	Observed estimate	Slide source note	Damodaran; 1 August 2026

Table labels are shortened; full local filenames are in `day04_metadata_and_results.json`. Retrieval date: 14 August 2026.

← **Return to estimate package**

Appendix map

Python practical map

1. Load frozen data and verify units and dates
2. Compute simple returns, means, CAGR, and volatility
3. Calculate covariance and two-asset portfolio risk
4. Simulate common and firm-specific shocks
5. Estimate beta with covariance and OLS
6. Compute Aster's CAPM return and sensitivity grid
7. Assemble and audit the required-return decision table

The local CSV files are frozen. Returns and rates enter Python as decimals; multiply by 100 only for percentage display.

← **Return to Day 4 synthesis**

Appendix map

P1. Load the frozen data and validate its contract

```python
import pandas as pd

# A DataFrame is a labelled table. parse_dates reads date text as dates.
factors = pd.read_csv(
    "../data/cases/ff_factors_monthly_2006_07_2026_06.csv",
    parse_dates=["date"]
)
aster = pd.read_csv(
    "../data/cases/aster_monthly_returns_synthetic_2021_07_2026_06.csv",
    parse_dates=["date"]
)

# assert stops if a declared condition is false.
assert len(factors) == 240
assert len(aster) == 60
assert aster["date"].max() == pd.Timestamp("2026-06-01")
print(aster[["Mkt-RF", "aster_excess"]].head(2))
```

`pd.read_csv` reads comma-separated data. Returns are decimal simple monthly returns; Aster is synthetic and the market factor is observed.

← **Return to Aster brief**

Practical map

P2. Compute arithmetic return, CAGR, and sample volatility

```python
import numpy as np

# A NumPy array stores aligned numeric values.
r = np.array([0.30, -0.10, 0.30, -0.10])
arithmetic = r.mean()
ending_wealth = np.prod(1 + r)
cagr = ending_wealth ** (1 / len(r)) - 1
volatility = r.std(ddof=1)  # ddof=1 uses T-1

print(round(arithmetic, 4), round(cagr, 4),
      round(volatility, 4))
```

`np.prod` multiplies entries; `**` exponentiates. Expected output: `0.10 0.0817 0.2309`.

← **Return to arithmetic versus CAGR**

Practical map

P3. Calculate two-asset return and risk from covariance

```python
import numpy as np

weights = np.array([0.60, 0.40])
expected = np.array([0.08, 0.14])
sigma = np.array([0.18, 0.30])
rho = 0.30

covariance = rho * sigma[0] * sigma[1]
cov = np.array([[sigma[0] ** 2, covariance],
                [covariance, sigma[1] ** 2]])
portfolio_return = weights @ expected
portfolio_volatility = np.sqrt(weights @ cov @ weights)

print(round(portfolio_return, 4),
      round(portfolio_volatility, 4))
```

`@` performs matrix multiplication. Expected output: `0.1040 0.1840`; the weighted-average volatility shortcut gives 0.2280 and is wrong.

← **Return to two-asset example**

Practical map

P4. Simulate the risk that diversification can and cannot remove

```
import numpy as np

rng = np.random.default_rng(41)  # fixed seed: reproducible draws
common = rng.normal(0, 0.16, size=(10000, 1))
firm = rng.normal(0, 0.30, size=(10000, 100))
returns = common + firm

# mean(axis=1) forms one equal-weight portfolio return per simulation.
for n in [1, 10, 100]:
    portfolio = returns[:, :n].mean(axis=1)
    print(n, round(portfolio.std(ddof=1), 4))
```

`default_rng` creates reproducible random draws. Expected volatility declines from about 34% toward the 16% common-risk floor.

← **Return to diversification floor**

Practical map

P5. Estimate beta with covariance and verify it by OLS

```python
import numpy as np
import pandas as pd

data = pd.read_csv(
    "../data/cases/aster_monthly_returns_synthetic_2021_07_2026_06.csv"
)
x = data["Mkt-RF"].to_numpy()
y = data["aster_excess"].to_numpy()

beta_cov = np.cov(y, x, ddof=1)[0, 1] / np.var(x, ddof=1)
design = np.column_stack([np.ones(len(x)), x])
alpha_ols, beta_ols = np.linalg.lstsq(design, y, rcond=None)[0]

print(round(beta_cov, 4), round(beta_ols, 4),
      round(alpha_ols, 6))
```

`lstsq` fits the least-squares line. Expected output: beta `1.0987` by both methods and monthly intercept `0.000878`.

← **Return to Aster regression**

Practical map

P6. Compute Aster's CAPM return and a sensitivity grid

```python
import numpy as np

rf, beta, market_premium = 0.04, 1.10, 0.05
cost_of_equity = rf + beta * market_premium

betas = np.arange(0.8, 1.41, 0.10)
premiums = np.arange(0.04, 0.0651, 0.005)
grid = rf + np.outer(betas, premiums)

print(round(cost_of_equity, 4))
print(round(grid.min(), 4), round(grid.max(), 4))
```

`np.outer` evaluates every beta–premium pair. Expected base output: `0.0950`; grid range: `0.0720 0.1310`.

← **Return to CAPM sensitivity**

Practical map

P7. Assemble the decision table and assert consistency

```
import pandas as pd

cases = pd.DataFrame({
    "case": ["Low", "Base", "High"],
    "beta": [0.90, 1.10, 1.30],
    "market_premium": [0.05, 0.05, 0.05]
})
cases["cost_of_equity"] = 0.04 + cases["beta"] * cases["market_premium"]
cases["wacc"] = 0.80 * cases["cost_of_equity"] + 0.20 * 0.045

assert cases.loc[1, "cost_of_equity"] == 0.095
assert cases.loc[1, "wacc"] == 0.085
print(cases.round(4))
```

A new DataFrame column stores a calculation for every row. The base case must reproduce 9.5% cost of equity and 8.5% WACC.

← **Return to Aster recommendation**

Practical map