CAPM and Jensen's Alpha

How the Capital Asset Pricing Model sets the bar, and how alpha measures skill against it.

1. Why CAPM exists

Every investor faces the same fundamental question: **how much return should I demand before I accept the risk of holding a security?** The Capital Asset Pricing Model (CAPM), developed in the 1960s by William Sharpe, John Lintner, and Jan Mossin, is the canonical answer. It provides a single, clean equation that tells you the return a rational investor should *expect* from an asset given the one risk that cannot be diversified away — its exposure to the overall market.

CAPM's great insight is the distinction between two kinds of risk: **systematic risk**, which affects every asset (recessions, interest-rate shifts, geopolitical shocks) and cannot be diversified away; and **idiosyncratic risk**, which is specific to a single asset (a product recall, a CEO scandal) and disappears when you hold a large enough portfolio. CAPM asserts that only systematic risk deserves a return premium, because idiosyncratic risk can be eliminated for free through diversification.

2. The CAPM equation

The expected return of an asset (or portfolio) *i* is:

$$E(R_i) = R_f + \beta_i \times (R_m - R_f)$$

Each term:

$E(R_i)$	The **expected return** of asset i. Read as: the return a rational investor should demand for holding this asset given the risk they are taking.
R_f	The **risk-free rate** — the return available from a zero-risk investment such as a short-term government bill. This is the floor: no one accepts risky returns below this.
R_m	The **market return** — the return of a broad benchmark portfolio (typically a diversified index such as the S&P 500 or a world equity index). It represents the return of the market as a whole.
$R_m - R_f$	The **market risk premium** — the extra return the market as a whole paid for bearing systematic risk. Historically around 5–7% per year for US equities, though highly variable.
β_i	The asset's **beta** — a measure of its systematic-risk exposure. $\beta=1$ means the asset moves with the market; $\beta>1$ amplifies market moves (more risk, more expected return); $\beta<1$ dampens them; $\beta<0$ moves against the market (rare; gold sometimes, short funds by design).

The equation says: start with what you could earn risk-free, then add a premium that is proportional to how much systematic risk you are accepting. The proportionality constant is beta. Double the beta, double the risk premium.

3. Understanding beta

Beta is estimated statistically, by regressing the asset's excess returns against the market's excess returns over a historical period (typically 2–5 years of weekly or monthly data). Formally:

$$\beta_i = \mathrm{Cov}(R_i, R_m) / \mathrm{Var}(R_m)$$

The numerator is how much asset *i* and the market move together; the denominator is the market's own variance. Beta is therefore the *slope* of the best-fit line in a scatter plot of the asset's returns against the market's returns.

- β = 1.0 — the asset moves **one-for-one** with the market. A diversified market-cap index fund has β ≈ 1 by construction.
- β = 1.5 — the asset moves **50% more** than the market. A 10% market move translates to a 15% move in the asset (on average). High-growth tech stocks often exhibit this.
- β = 0.5 — the asset moves **half as much** as the market. Defensive sectors like utilities and consumer staples typically sit here.
- β = 0 — the asset is **uncorrelated** with the market (e.g. certain hedge strategies, or a pure cash position).
- β < 0 — the asset moves **opposite** to the market. Gold, inverse ETFs, and deep-out-of-the-money index puts can fall here.

How this site computes β

The Beta panel in the Portfolio page does not ask you to guess each holding's β — it estimates one for every symbol you have ever traded by running an ordinary-least-squares (OLS) regression directly on our local end-of-day price database. Concretely, for each symbol *i* the backend pulls up to **10 years** of daily closing prices for both the symbol and the benchmark (SPY), converts them to **daily simple returns** ($r_t = p_t/p_{t-1} - 1$), aligns the two series on common trading days, drops missing observations, and computes:

$$\beta_i = \mathrm{Cov}(r_i^d, r_m^d) / \mathrm{Var}(r_m^d)$$

where r^d denotes a daily simple return and the covariance/variance are the sample (ddof = 1) estimators. The same arithmetic is the slope of an OLS line through the (r_m, r_i) scatter plot — exactly the textbook definition. Alongside β we also return:

- **R²** — the share of the symbol's daily-return variance that the benchmark explains. $R^2 = \mathrm{corr}(r_i, r_m)^2$. A name with high R² (say > 0.6) has a tight, trustworthy β; below ~0.2 the β estimate is noisy and should be treated as approximate.
- **n** — the number of overlapping daily return observations actually used. A 10-year window typically yields ~2,500 obs for liquid US equities; recently listed names will have far fewer. If **n < 30** we refuse to fit and tag the row *no fit*, falling back to β = 1.0 as a neutral default.
- **start / end** — the actual date range used (visible on hover). If the symbol's history is shorter than 10 years we silently shrink to the available overlap; the

symbol's listing date, not the benchmark's, is the binding constraint.

The full *portfolio* beta — the number applied to the Alpha modal — is then the market-value-weighted average of the per-symbol fits:

$$\beta_p = \Sigma_i\, w_i \times \beta_i, \quad w_i = MV_i \,/\, \Sigma_j\, MV_j$$

MV_i is the holding's current market value (qty × latest close), or its book cost if no price is available (these rows are flagged with an asterisk). Closed positions (qty = 0) are listed for transparency but excluded from the weights; short positions (qty < 0) keep their sign, so they contribute *negative* weight — the same way they offset the portfolio's market exposure in real life. Any cell can be overridden by typing a different number; manual entries are highlighted in amber and the weighted total recomputes live.

Why these specific choices

- **Daily, not monthly, returns.** A 10-year window gives ~2,500 daily observations versus ~120 monthly ones. The standard error of a regression slope shrinks roughly with √n, so the daily estimate is ~4–5× tighter. The classical objection to daily β — non-synchronous trading bias for thinly-traded names — is negligible for SPY-listed equities, which essentially all trade in every session.

- **10-year lookback.** Long enough to span at least one full market cycle (post-GFC bull, COVID crash, 2022 drawdown) so β is not dominated by a single regime, yet short enough that pre-period structural shifts (a company changes business model, a sector restructures) don't pollute the estimate. We deliberately do *not* let users vary this in the modal — switching window lengths post-hoc to find a flattering β is exactly the kind of reverse-engineering the regression should not enable.

- **Raw returns, not excess returns.** Strict CAPM theory regresses ($R_i - R_f$) on ($R_m - R_f$). At daily frequency R_f/252 is on the order of 1–2 basis points per day; subtracting the same constant from both sides shifts the regression intercept but leaves the *slope* — β — virtually unchanged (typically within a few thousandths). Skipping the subtraction keeps the math transparent without materially altering any number.

- **SPY as the default benchmark.** SPY is the most heavily traded ETF tracking the S&P 500 and is the canonical proxy for the US equity market in most academic and industry work. Portfolios concentrated outside US large-caps will see weak fits (low R^2) — that is itself useful information: it tells you SPY is the wrong yardstick, and the resulting β should be taken with a grain of salt.

- **Spike filtering.** Before computing returns we drop implausible single-day prints (corporate-action artefacts, data-vendor errors). One bad close inserts two huge offsetting daily returns and can move β by ±0.2 — large enough to matter at the portfolio level. The filter is conservative; it removes only obvious outliers (multi-σ moves immediately reversed the next day).

- **Per-symbol fallback to 1.0.** When a symbol has fewer than 30 overlapping daily observations the regression has too few degrees of freedom to be meaningful; we report β = 1.0 (the market-neutral default) rather than a number with no statistical support. This is conservative — it tilts the portfolio β estimate toward 1 for new IPOs and recently added tickers, and you can override it manually if you have a stronger prior.

4. The Security Market Line (SML)

Plotting expected return against beta gives a straight line — the **Security Market Line**. It passes through (β=0, R_f) and (β=1, R_m), and its slope is the market risk premium. Every correctly priced asset sits exactly on the line. CAPM's verdict about any asset is a single question: *is it on the line, above it, or below it?*

- **On the line** — the asset earns exactly what CAPM predicts for its beta. Fair value.
- **Above the line** — the asset earns more than predicted. Either it is undervalued (buy it), or the manager has generated skill that CAPM does not account for. *This is alpha.*
- **Below the line** — the asset earns less than predicted. Overvalued or poorly managed.

5. Assumptions and limitations

CAPM is a *model*, not a law of nature. Its derivation relies on several strong assumptions, and each one is violated in practice to some degree:

- Investors are rational and hold mean–variance-optimal portfolios (often not the case — behavioural biases are well documented).
- All investors have the same information and time horizon (obviously false).
- There are no taxes or transaction costs (clearly false).
- Investors can borrow and lend freely at the risk-free rate (individual investors cannot).
- Returns are normally distributed (empirically they have fat tails — extreme moves happen more often than a normal distribution predicts).
- A single factor — the market — explains all systematic return (contradicted by the Fama-French 3- and 5-factor models, which find size, value, profitability, and investment also matter).

Despite these flaws, CAPM remains the starting point for virtually every risk/return discussion. Its elegance and the interpretive power of β make it the default benchmark even in contexts where everyone involved knows it is only approximately right.

6. Jensen's Alpha — the measure of skill

Jensen's alpha, introduced by Michael Jensen in 1968 to evaluate mutual fund managers, is the **difference between the portfolio's actual return and the return CAPM predicted** for its level of systematic risk. It directly measures whether a portfolio beat its CAPM-implied benchmark after accounting for beta.

$$\alpha = R_p - [\ R_f + \beta \times (R_m - R_f)\ ]$$

Components:

α	**Jensen's alpha** — the portfolio's excess return beyond what CAPM expected. Expressed in the same units as returns (percentage points).
R_p	The **portfolio's actual return** over the evaluation period.
$R_f + \beta(R_m - R_f)$	The **expected return** under CAPM — exactly the formula from section 2. This is the benchmark you must beat to earn a positive alpha.

The key intuition: simply beating the market ($R_p > R_m$) is *not* enough to claim skill. A high-beta portfolio should beat the market in up years just from its beta exposure; we have not yet controlled for that. Only after subtracting the CAPM-implied return — which bakes in the beta adjustment — can we say whether any return was truly alpha (skill) rather than beta (risk-taking).

7. Interpreting alpha

Sign of α	Meaning	Verdict
α > 0	Portfolio earned more than CAPM predicted for its level of systematic risk.	Manager added value.
α = 0	Portfolio earned exactly what CAPM predicted. All of its return is compensation for systematic risk.	Market-neutral skill; no alpha.
α < 0	Portfolio earned less than CAPM predicted for its risk level.	Underperformance after adjusting for risk.

A positive α is what every active manager is hired to produce. The unpleasant empirical finding from decades of fund-performance studies is that, after fees, the *average* active manager delivers α close to zero or mildly negative. A minority do deliver sustained positive alpha, but identifying them in advance is notoriously difficult.

8. Worked example

Suppose over the last year you observe the following for your portfolio and the market:

Variable	Value
Portfolio return Rp	14.0%
Benchmark return Rm (S&P 500)	10.0%

Risk-free rate Rf (3M T-Bill)	5.0%
Portfolio beta β	1.30

Step 1 — market risk premium:

$$R_m - R_f = 10.0\% - 5.0\% = 5.0\%$$

Step 2 — the premium, scaled by beta:

$$\beta \times (R_m - R_f) = 1.30 \times 5.0\% = 6.5\%$$

Step 3 — CAPM expected return:

$$E(R_p) = R_f + 6.5\% = 5.0\% + 6.5\% = 11.5\%$$

Step 4 — Jensen's alpha:

$$\alpha = R_p - E(R_p) = 14.0\% - 11.5\% = \mathbf{+2.5\%}$$

Conclusion: the portfolio outperformed its CAPM-implied benchmark by **2.5 percentage points**. Note that the naïve comparison $R_p - R_m = 14\% - 10\% = 4\%$ overstates the skill, because it does not account for the fact that a β=1.30 portfolio is structurally expected to beat the market by 1.5% in a year when the market beats the risk-free rate by 5%. Jensen's alpha isolates the part of that 4% outperformance that *cannot* be explained by extra risk-taking.

9. Limitations of Jensen's alpha

- Inherits **all of CAPM's assumptions**. If CAPM is wrong (and it is, partially), α is contaminated with whatever the model failed to capture. A high α may simply reflect a missing risk factor rather than true skill.
- **Beta itself must be estimated** from historical data. A mis-estimated β directly distorts α. A portfolio concentrated in a single sector, or one with changing leverage, will have unstable β.
- α is **backward-looking**. Past alpha does not guarantee future alpha — and for most managers, it does not predict it either.
- The choice of **benchmark** matters. Measuring a small-cap portfolio's alpha against the S&P 500 is a category error; the correct benchmark is a small-cap index.
- The evaluation **period length** matters. Short windows are dominated by noise; long windows may blend multiple market regimes.

10. Close cousins of Jensen's alpha

- **Treynor ratio** = $(R_p - R_f) / \beta$. Excess return per unit of systematic risk. Same numerator as a simplified alpha, but normalised by β so portfolios of different risk can be ranked.
- **Sharpe ratio** = $(R_p - R_f) / \sigma_p$. Uses total volatility instead of β, so it penalises idiosyncratic risk as well — useful when the portfolio isn't well diversified.

- **Information ratio** = α / tracking-error. Measures alpha per unit of benchmark-deviation risk — standard for evaluating active managers benchmarked to an index.
- **Fama–French alpha** — α computed against a 3- or 5-factor model (market + size + value [+ profitability + investment]) rather than CAPM alone. A portfolio with positive CAPM α but zero Fama–French α was not skilled; it was just loaded on cheap and small stocks.

11. Practical checklist

- Choose a benchmark that genuinely reflects the portfolio's opportunity set (asset class, geography, size/style).
- Use **the same period** for R_p, R_m and R_f. Mismatched periods are the most common source of bogus alpha.
- Estimate β from at least two years of data, using the same return frequency (daily, weekly, or monthly) throughout.
- Check that α is **statistically significant**. A single-point estimate of +2.5% is not evidence of skill if its standard error is ±6%.
- Remember that positive alpha is **zero-sum**: the sum of all investors' alphas must equal the costs of trading. Most investors cannot have positive alpha simultaneously.