The Capital Asset Pricing Model

Lorenzo Naranjo

March 2026

Introduction

Beta Pricing is Mathematics, Not Economics

The preceding notebooks established two benchmark facts. First, in frictionless mean-variance geometry, beta pricing holds for any frontier portfolio. Second, with a risk-free asset, this pricing relation takes the single-beta form based on the tangency portfolio q, the risky portfolio with the highest Sharpe ratio. The short-sales notebook showed how constraints can break that global relation. In this notebook we return to the frictionless benchmark and ask when the tangency portfolio should be interpreted economically as the market portfolio. Applied to q, beta pricing gives

$$\mathrm{E}(r_i) = r_f + \beta_i(\mathrm{E}(r_q) - r_f),$$

where $\beta_i = \mathrm{Cov}(r_i, r_q)/\sigma^2(r_q)$. It is essential to understand what this result is and what it is not.

Beta pricing is a *mathematical identity*. It follows from the geometry of the minimum-variance frontier and holds for the tangency portfolio q by construction. Its only economic content is the no-arbitrage condition embedded in the economy's definition (see Definition 1 and the footnote in the beta pricing notebook): we need asset payoffs to be arbitrage-free so that the law of one price holds and expected returns are well-defined. Beyond that, beta pricing makes no claim about *which* portfolio investors actually hold, why they hold it, or how prices are determined in equilibrium. It is silent on preferences, on how beliefs are formed, and on market clearing. Any mean-variance efficient portfolio satisfies it.

Equally important, beta pricing imposes no assumptions on the *distribution* of returns. The entire derivation requires only that the covariance matrix $\mathbf{V}$ is invertible — i.e., that returns have finite means and variances. Returns need not be normally distributed, nor elliptically distributed, nor follow any other parametric family. This is a significant generality: the formula holds whether returns are fat-tailed, skewed, or exhibit any other non-Gaussian features, as long as the second moment exists. The CAPM, by contrast, typically requires either normally distributed returns or quadratic utility to justify mean-variance preferences as the basis for investor choice, both of which are strong additional assumptions.

The CAPM as an Economic Theory

The Capital Asset Pricing Model (CAPM), developed independently by Sharpe (1964) and Lintner (1965), is an *economic* statement. It takes the frictionless mathematical framework of the previous notebooks and adds assumptions about investor behavior and market equilibrium to answer a question that beta pricing cannot: *which* portfolio is the relevant pricing factor?

The CAPM's answer is the *market portfolio* — the value-weighted portfolio of all risky assets in the economy. This identification does not follow from mathematics alone; it requires an argument about how individual investor demands aggregate. The assumptions are economic in nature: investors share the same beliefs, face the same opportunity set, and markets clear. Under these conditions, every mean-variance investor optimally holds a combination of the risk-free asset and the same tangency portfolio q — the risky portfolio with the highest Sharpe ratio. Since all investors agree on q, and in equilibrium aggregate risky-asset holdings must equal total supply, the tangency portfolio must be the market portfolio. Beta pricing then applies with $q = m$, giving the CAPM formula.

The distinction matters for interpretation and testing. A finding that expected returns are linear in beta with respect to some portfolio is always true if that portfolio is efficient — it proves nothing about the CAPM. Verifying the CAPM requires the additional economic step of showing that the efficient portfolio in question is the market portfolio, which is precisely the content of Roll's critique discussed below.

The CAPM

The Equilibrium Argument

We maintain the frictionless setup of the risk-free asset notebook and add the equilibrium assumptions needed to identify the tangency portfolio with the market portfolio.

Assumption 1 (CAPM Assumptions)**.** *In addition to the economy described in Definition 1, we assume:*

i. All investors have mean-variance preferences and share identical beliefs about the joint distribution of asset returns.
ii. All investors face the same risk-free rate r_f and can borrow or lend at this rate without restriction.
iii. Asset markets clear, so aggregate demand for each risky asset equals its exogenous supply.

Under these assumptions[1] every investor faces the same mean-variance problem. With a risk-free asset, the optimal risky portfolio is the one that maximizes the Sharpe ratio — which is the tangency portfolio q. Investors differ only in how much they allocate between r_f and q, depending on their risk aversion, but they all hold the same risky portfolio q. Since q is the only risky portfolio held by any investor, in aggregate the demand for each risky asset is proportional to its weight in q. Market clearing then requires these weights to equal the market-capitalization weights, so q coincides with the market portfolio m.

This equilibrium argument also explains why the no-short-sales constraint is naturally non-binding in the CAPM benchmark. If every risky asset is in positive net supply, market clearing requires the market portfolio to assign it a strictly positive weight. Since the tangency portfolio coincides with the market portfolio in equilibrium, it must also assign strictly positive weight to

[1] Mean-variance preferences can be justified in two standard ways. If utility is quadratic, $u(W) = W - \frac{b}{2}W^2$, then expected utility depends only on the mean and variance of wealth for any return distribution (Mossin 1966). Alternatively, if returns are jointly normally distributed, any expected-utility maximizer with increasing concave utility acts as a mean-variance optimizer, since the normal is fully characterized by its first two moments (Tobin 1958). A common special case combines constant absolute risk aversion (CARA) utility with normality, which delivers closed-form portfolio weights. Both approaches yield the CAPM under the remaining equilibrium assumptions, but each imposes a restriction — on preferences or on the return distribution — that the direct mean-variance assumption avoids.

every asset. Thus prices adjust so that the unconstrained tangency portfolio already satisfies the non-negativity constraints. At q, the constrained and unconstrained minimum-variance frontiers therefore coincide.

This observation is also the starting point of the Black and Litterman (1992) approach. If market-capitalization weights are taken as an equilibrium tangency portfolio, then the CAPM can be inverted to recover the vector of equilibrium expected returns consistent with those weights. Investor views are then introduced as deviations from that benchmark rather than as a replacement for it.

Property 1 (Capital Asset Pricing Model)**.** *Under Assumption 1, the tangency portfolio is the market portfolio m. For any asset or portfolio i,*

$$\mathrm{E}(r_i) = r_f + \beta_i \left(\mathrm{E}(r_m) - r_f\right),$$

where $\beta_i = \mathrm{Cov}(r_i, r_m)/\sigma^2(r_m)$.

The proof follows immediately from the beta-pricing result in the risk-free asset notebook applied to $q = m$. The quantity $\mathrm{E}(r_m) - r_f$ is the *equity risk premium*, and β_i measures the systematic exposure of asset i to market risk. Idiosyncratic risk carries no premium: two assets with the same β_i must have the same expected return regardless of their total variance.

Black's zero-beta CAPM (Black 1972) is the main theoretical extension: if investors cannot all borrow and lend freely at the same risk-free rate, the linear pricing relation survives but the intercept is no longer r_f and is instead pinned down by the expected return of a zero-beta portfolio.

The Security Market Line

The CAPM implies a linear relationship between an asset's expected return and its beta. This line — defined by the two points $(0, r_f)$ and $(1, \mathrm{E}(r_m))$ in $(\beta, \mathrm{E}(r))$ space — is called the *security market line* (SML). Every asset and portfolio must lie on the SML in equilibrium.

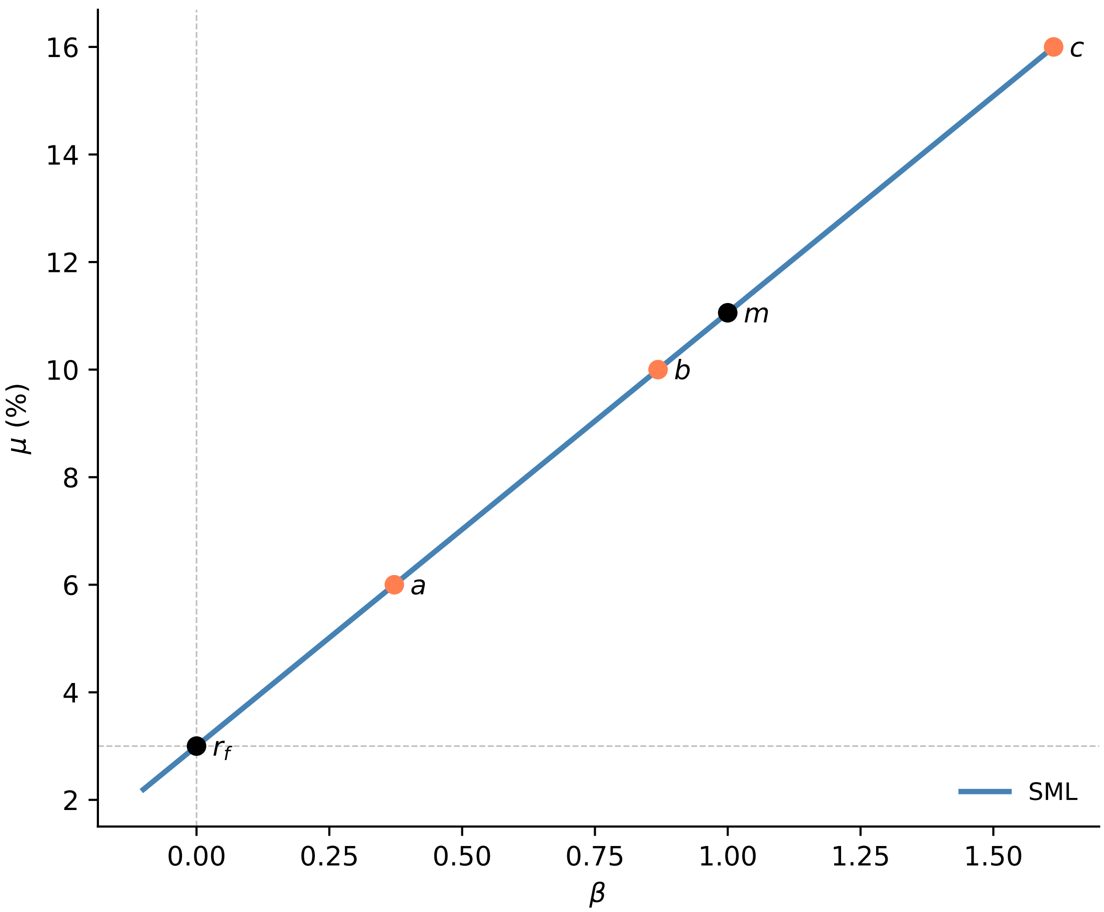

Figure 1: Security market line (SML). All assets must lie on the line in CAPM equilibrium. The market portfolio has $\beta = 1$ and expected return $E(r_m)$; the risk-free asset has $\beta = 0$ and expected return r_f.

Note that the SML is defined in (β, μ) space, which is distinct from the mean-standard deviation space used to display the minimum-variance frontier. The SML applies to all assets and portfolios, including those that are not on the frontier. The CAL, by contrast, only describes the set of efficient portfolios formed by combining the risk-free asset with the market portfolio.

Roll's Critique

Roll (1977) raised a fundamental objection to empirical tests of the CAPM. His argument has two parts.

First, the CAPM is a statement about the *true* market portfolio, which in principle includes every risky asset in the economy: publicly traded stocks and bonds, but also private equity, real estate, human capital, and all other stores of value. This portfolio is unobservable in practice. Empirical tests must therefore substitute a measurable proxy — typically a broad stock index.

Second, and more fundamentally, the beta-pricing relation $\mathrm{E}(r_i) = r_f + \beta_i(\mathrm{E}(r_p) - r_f)$ holds exactly for *any* mean-variance efficient portfolio p, as a mathematical consequence of the frontier analysis in the previous notebooks. It does not require equilibrium. This means that if a researcher finds that a given index is mean-variance efficient, the SML will hold by construction — but this says nothing about whether the index approximates the true market portfolio or whether the CAPM equilibrium condition is satisfied. Conversely, if the index is not efficient, the SML will be rejected even if the CAPM is true and the true market portfolio is efficient.

The implication is stark: any test of the CAPM using a proxy for the market is simultaneously a test of two distinct hypotheses — (i) that the CAPM equilibrium holds and (ii) that the chosen proxy is a good approximation to the true market portfolio. These hypotheses cannot be separated without observing the true market portfolio, making the CAPM fundamentally untestable.

This joint-hypothesis problem shaped the empirical literature that followed, including the classic cross-sectional tests of Fama and MacBeth (1973), which operationalized the CAPM with observable proxies while inheriting exactly the identification problem Roll emphasized.

References

Black, Fischer. 1972. "Capital Market Equilibrium with Restricted Borrowing." *Journal of Business* 45 (3): 444–55.

Black, Fischer, and Robert Litterman. 1992. "Global Portfolio Optimization." *Financial Analysts Journal* 48 (5): 28–43.

Fama, Eugene F., and James D. MacBeth. 1973. "Risk, Return, and Equilibrium: Empirical Tests." *Journal of Political Economy* 81 (3): 607–36.

Lintner, John. 1965. "The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets." *Review of Economics and Statistics* 47 (1): 13–37.

Mossin, Jan. 1966. "Equilibrium in a Capital Asset Market." *Econometrica* 34 (4): 768–83.

Roll, Richard. 1977. "A Critique of the Asset Pricing Theory's Tests Part I: On Past and Potential Testability of the Theory." *Journal of Financial Economics* 4 (2): 129–76.

Sharpe, William F. 1964. "Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk." *Journal of Finance* 19 (3): 425–42.

Tobin, James. 1958. "Liquidity Preference as Behavior Towards Risk." *Review of Economic Studies* 25 (2): 65–86.