Adding a Risk-Free Asset

We know from Portfolios of Risky Assets that, with no risk-free asset available, the investment opportunity set of several risky assets is a hyperbola-shaped region bounded on the left by the minimum-variance frontier.

Adding a risk-free asset changes what is available to investors. Recall from Optimal Capital Allocation that combining the risk-free asset with a single risky portfolio P produces a straight line, the CAL of P, with intercept r_f and slope equal to the Sharpe ratio of P. With many risky assets to choose from, a different CAL is available for every risky portfolio P in the hyperbola-shaped opportunity set, not just for A, B, or C but for every combination of the three.

The new investment opportunity set, once the risk-free asset is available, is the union of all these CALs. As we show below, this union is no longer the hyperbola: it is a triangle, with the best possible CAL, the one through the tangency portfolio Q, as its upper edge.

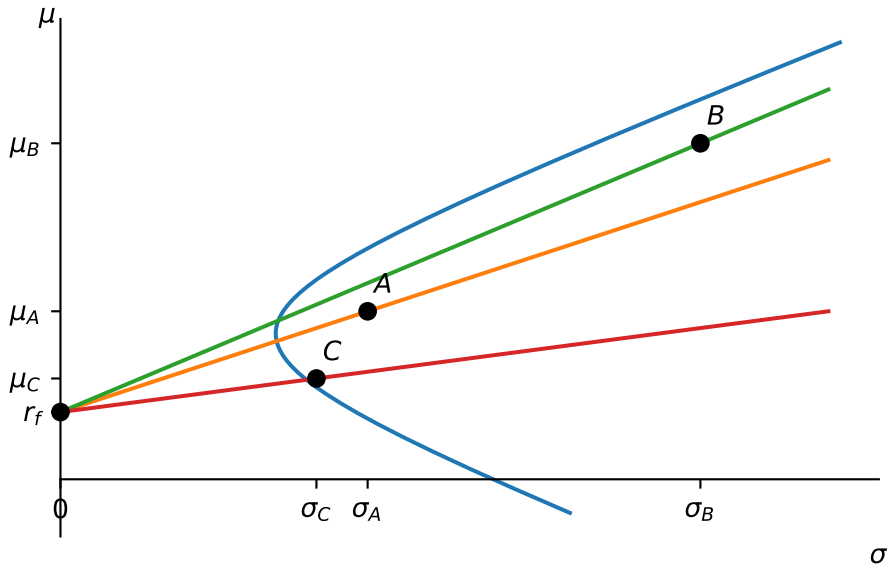


Figure 1: The figure shows the investment opportunity set generated by three risky assets A, B, and C together with their capital allocation lines. The different slopes of the CALs illustrate different Sharpe ratios for alternative risky portfolios.

By moving the CALs higher and higher, we reach a limit on how high the Sharpe ratio of the CALs can be. Indeed, there is one portfolio Q that achieves the highest Sharpe ratio. We call this

portfolio the **tangency portfolio** because it is, in effect, the portfolio that generates a CAL that is tangent to the original investment opportunity set generated by the three risky assets.

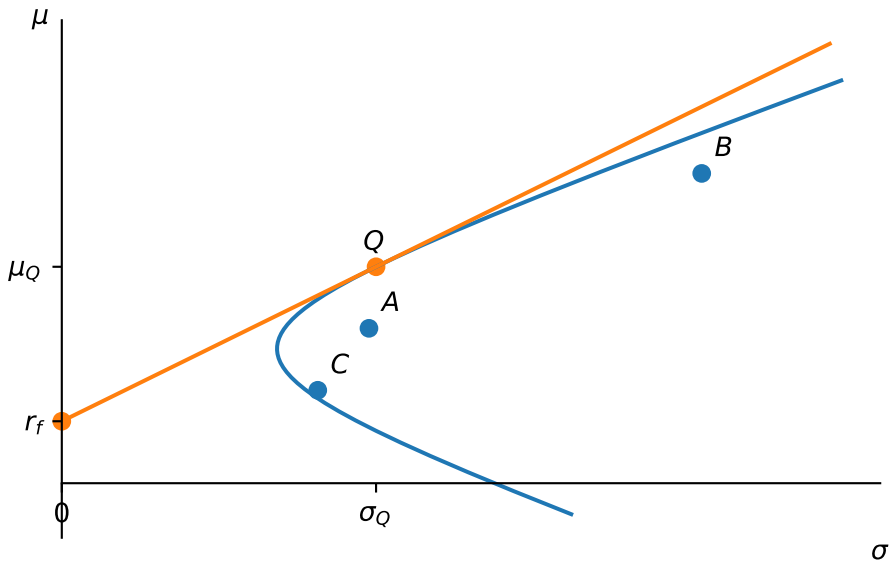


Figure 2: The figure shows the capital allocation line of the tangency portfolio Q obtained from the three risky assets A, B, and C.

Investors that prefer portfolios with the highest Sharpe ratio will choose a portfolio located in the CAL of the tangency portfolio. We call this CAL the **efficient frontier**. All portfolios in this CAL are efficient since they all have the highest Sharpe ratio. Any two portfolios in this line are enough to generate all the other efficient portfolios. Any other portfolios in this economy will have lower Sharpe ratios than the tangency portfolio. To determine if a portfolio is efficient or not, just compute its Sharpe ratio and compare it to the Sharpe ratio of an efficient portfolio.

No combination of the risk-free asset with a risky portfolio can ever beat the CAL of Q, since Q has the highest possible Sharpe ratio; this line is tangent to the hyperbola exactly at Q. Going the other way, shorting Q and investing more than 100% of wealth in the risk-free asset traces out the mirror-image line through $(0, r_f)$ with the same slope in absolute value but pointing down instead of up. Unlike the CAL of Q, this line is *not* tangent to the hyperbola: at every risk level it lies below even the worst mean-variance *inefficient* combination of A, B, and C with that same volatility. This makes sense: shorting the single best portfolio to fund extra risk-free holdings is a strategy built from an efficient portfolio, so it can push further in the bad direction than any

ordinary combination of A, B, and C could. Still, it is a true bound on what is achievable, so together with the CAL of Q it encloses the entire hyperbola in a triangular wedge. Once the risk-free asset is available, this triangle, not the hyperbola, is the new investment opportunity set.

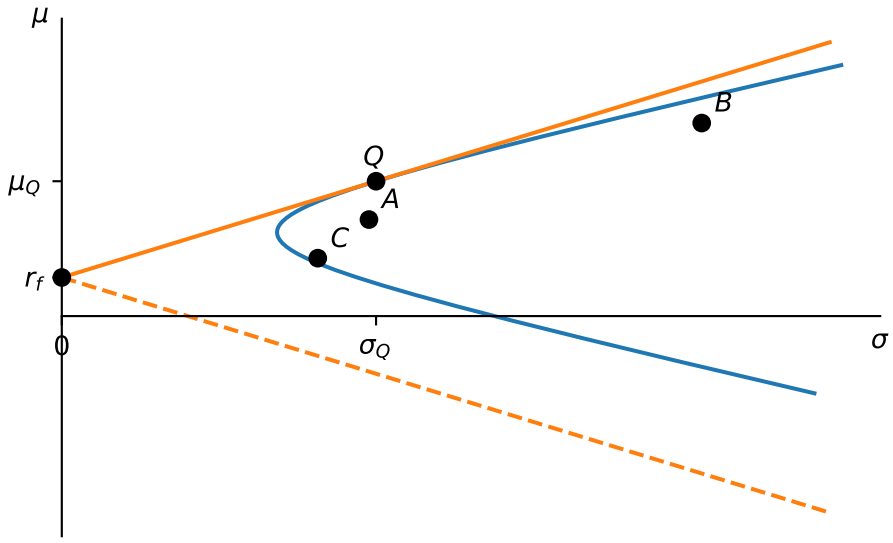


Figure 3: The figure shows how the risky-asset hyperbola of A, B, and C is enclosed by two straight lines through r_f. The upper line is the CAL of the tangency portfolio Q and is tangent to the hyperbola. The lower line, obtained by shorting Q instead, is not tangent to the hyperbola but still bounds it from below. Once the risk-free asset is available, this triangle, not the hyperbola, is the investment opportunity set.

This is the real payoff of Figure 3: once the risk-free asset is available, the hyperbola itself stops mattering. Every efficient portfolio is just a combination of the risk-free asset and Q, so *if* an investor has mean-variance utility, their problem collapses back to the simple capital allocation problem of Optimal Capital Allocation, how much to invest in Q and how much in the risk-free asset, except that now Q is not just some risky asset, but specifically the tangency portfolio built from A, B, and C. Such an investor never needs to think about A, B, or C individually again: comparing any portfolio's Sharpe ratio to that of Q is all that is needed.

Not every investor has mean-variance utility, so this is not a claim that everyone should hold some combination of Q and the risk-free asset. An investor with different preferences may rationally choose an inefficient portfolio. What survives regardless of preferences is *pricing*: as the next section shows, once Q is known to be mean-variance efficient, the expected return

of *every* asset, whether or not it is itself efficient, is pinned down by its beta with Q and the risk-free rate. Comparing Sharpe ratios to Q's is still the right way to check whether a *portfolio* is efficient, but the beta pricing relation is not something an asset can fail: it holds by construction for every asset in the economy.

Example 1 (Testing for Efficiency)**.** You know that the tangency portfolio has an expected return of 15% with a standard deviation of returns of 20%. The risk-free rate is 5% per year. You have the following information of two risky assets.

Asset	Expected Return	Standard Deviation
A	10%	20%
B	25%	40%

Are these portfolios efficient? The Sharpe ratio of the tangency portfolio is $(15-5)/20 = 0.5$. The Sharpe ratio of A is $(10-5)/20 = 0.25 < 0.5$, whereas the Sharpe ratio of B is $(25-5)/40 = 0.5$. Thus only B is an efficient portfolio. □

Example 2 (Constructing an Efficient Portfolio)**.** In Example 1, note that by investing 50% in the risk-free asset and 50% in the tangency portfolio, we obtain a portfolio with the same expected return as A but with a lower standard deviation of $0.5 \times 0.2 = 10\%$.

A mean-variance investor targeting a 10% expected return would prefer this efficient portfolio to A: same expected return, less risk. □

Beta Pricing

The triangle of Figure 3 describes what is attainable by combining the risk-free asset with A, B, and C. But its real usefulness goes beyond this: the tangency portfolio Q is mean-variance efficient by construction (no other portfolio has a higher Sharpe ratio), and this alone, with no reference to market equilibrium, is enough to derive a beta pricing formula for *every* asset or portfolio in the economy, whether or not it is one of the three risky assets used to construct Q, and whether or not it even lies inside the triangle.

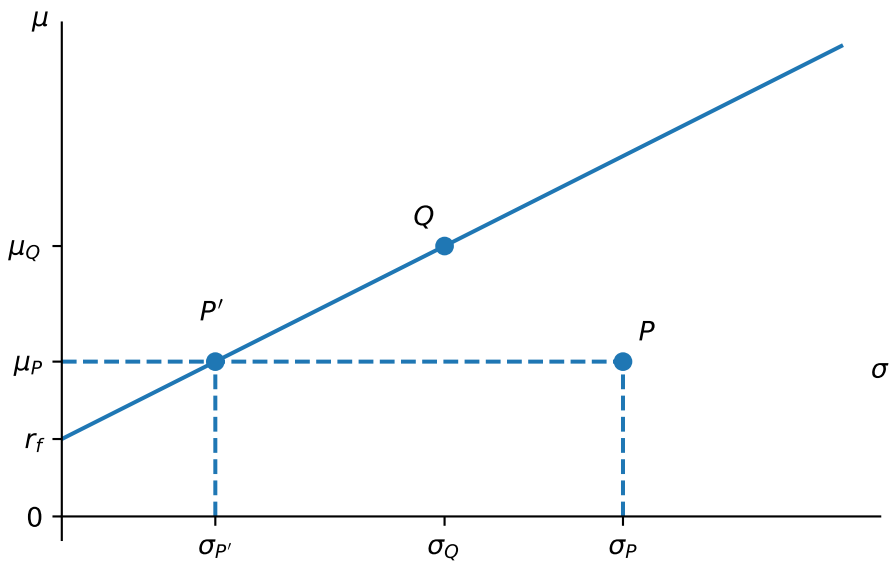


Figure 4: The figure shows the tangency portfolio Q obtained from the risky assets A, B, and C. Portfolio P' is efficient and has the same expected return as P, but a lower standard deviation.

In Figure 4, portfolio P' is efficient, representing a combination of the tangency portfolio and the risk-free asset. Its return can be expressed as:

$$r_{P'} = (1 - \beta)r_f + \beta r_Q.$$

The residual $\varepsilon = r_P - r_{P'}$ is constructed to have a mean of zero, leading to:

$$\mathrm{E}(r_P) = (1 - \beta)r_f + \beta\,\mathrm{E}(r_Q).$$

Additionally, the residual is orthogonal to its projection $r_{P'}$,[1] which implies:

$$0 = \mathrm{Cov}(\varepsilon, r_{P'}) = \mathrm{Cov}(r_P - r_{P'}, r_{P'}) = \beta\,\mathrm{Cov}(r_P, r_Q) - \beta^2\,\mathrm{V}(r_Q),$$

or equivalently:

$$\beta = \frac{\mathrm{Cov}(r_P, r_Q)}{\mathrm{V}(r_Q)}.$$

Nothing in this argument used market clearing, equilibrium, or any assumption about the *shape* of the return distribution: no normality, no symmetry, nothing beyond Q being mean-variance efficient. All that is required is that returns have finite means, variances, and covariances. The same orthogonal projection argument works for *any* portfolio P against *any* efficient portfolio Q.

Property 1 (Beta Pricing)**.** *Let Q be any mean-variance efficient portfolio. The return of any asset P can be decomposed as*

$$r_P = (1 - \beta_P) r_f + \beta_P r_Q + \varepsilon_P,$$

where

$$\beta_P = \frac{\mathrm{Cov}(r_P, r_Q)}{\mathrm{V}(r_Q)}$$

[1] With the risk-free asset available, let $\boldsymbol{w}$ hold the weights in the risky assets only, with the remainder $1 - \boldsymbol{w}^\top \mathbf{1}$ held in the risk-free asset; the portfolio's mean and variance are then $r_f + \boldsymbol{w}^\top(\boldsymbol{\mu} - r_f\mathbf{1})$ and $\boldsymbol{w}^\top\boldsymbol{\Sigma}\boldsymbol{w}$, with the budget constraint already built in. Since Q is mean-variance efficient, its weights $\boldsymbol{w}_Q$ minimize variance among all portfolios with mean μ_Q; the Lagrangian $L(\boldsymbol{w}, \lambda) = \boldsymbol{w}^\top\boldsymbol{\Sigma}\boldsymbol{w} - 2\lambda\left[\boldsymbol{w}^\top(\boldsymbol{\mu} - r_f\mathbf{1}) - (\mu_Q - r_f)\right]$ for this problem has first-order condition

$$\boldsymbol{\Sigma}\boldsymbol{w}_Q = \lambda_Q(\boldsymbol{\mu} - r_f\mathbf{1})$$

for some scalar λ_Q. The same condition holds for any efficient portfolio on the CAL of Q, with λ rescaled accordingly, except at the risk-free intercept itself, where $\boldsymbol{w} = \mathbf{0}$ and $\lambda = 0$ give a degenerate, uninformative relation. So for *any* portfolio i (with weights $\boldsymbol{w}_i$ in the risky assets and the remainder in the risk-free asset)

$$\begin{aligned}\mathrm{Cov}(r_i, r_Q) &= \boldsymbol{w}_i^\top\boldsymbol{\Sigma}\boldsymbol{w}_Q \\ &= \lambda_Q\boldsymbol{w}_i^\top(\boldsymbol{\mu} - r_f\mathbf{1}) = \lambda_Q(\mu_i - r_f),\end{aligned}$$

an affine function of μ_i alone, using $\boldsymbol{w}_i^\top(\boldsymbol{\mu} - r_f\mathbf{1}) = \mu_i - r_f$ from the mean formula above (which holds regardless of how much of i sits in the risk-free asset). Since $\varepsilon = r_P - r_{P'}$ was built so that $\mathrm{E}(r_{P'}) = \mathrm{E}(r_P)$, this gives $\mathrm{Cov}(r_P, r_Q) = \mathrm{Cov}(r_{P'}, r_Q)$, i.e. $\mathrm{Cov}(\varepsilon, r_Q) = 0$. Since also $\mathrm{Cov}(\varepsilon, r_f) = 0$ (as r_f is constant), bilinearity applied to $r_{P'} = (1 - \beta)r_f + \beta r_Q$ gives $\mathrm{Cov}(\varepsilon, r_{P'}) = 0$.

and $\mathrm{E}(\varepsilon_P) = \mathrm{Cov}(\varepsilon_P, r_Q) = 0$. *Taking expectations,*

$$\mathrm{E}(r_P) = (1 - \beta_P) r_f + \beta_P \, \mathrm{E}(r_Q).$$

Property 1 holds regardless of investor preferences and regardless of any assumption on the shape of the return distribution. All it needs is that returns have well-defined means, variances, and covariances. The reason is diversification, not optimization: for any asset P, the portfolio $\beta_P r_Q + (1 - \beta_P) r_f$ matches P's expected return while shedding ε_P, the piece of P's return that is firm-specific to P and uncorrelated with Q. That piece adds variance without adding expected return, so it cannot earn a premium of its own; $\mathrm{E}(r_P)$ is pinned down by β_P alone. This is a statement about *pricing*, not about what investors should hold. It says nothing about whether Q is anyone's optimal portfolio, only that its efficiency is enough to price every asset relative to it.[2]

The CAPM, covered in the next note, adds a single extra claim on top of Property 1: that the tangency portfolio Q coincides with the market portfolio M. Many textbooks blur this distinction and present the beta pricing equation as if it were a consequence of the CAPM's equilibrium assumptions, when in fact it holds for *any* efficient portfolio, whether or not it is the market.

Practice Problems

These problems give you a chance to practice the concepts introduced in this chapter. Try to solve each one on your own before expanding the solution.

Solutions to all problems can be found at https://lorenzonaranjo.com/class-materials/investments/.

[2] In more advanced treatments, this result is restated as the pricing equation $p = \mathrm{E}(mx)$ for a stochastic discount factor m: the mean-variance efficiency of Q is exactly what makes a linear function of r_Q a valid m for every asset in the economy. Existence of such an m only requires the law of one price (LOOP). Payoffs replicable at the same cost must earn the same price, which is strictly weaker than absence of arbitrage. Absence of arbitrage further requires $m > 0$ almost surely, and nothing here guarantees that: a mean-variance efficient Q prices every asset correctly under LOOP even if m happens to be negative in some states, in which case arbitrage opportunities can still exist.

Problem 1 (Screening Mutual Funds for Efficiency)**.** Elena Rossi, an advisor at Fairhaven Wealth Partners, uses a tangency portfolio Q with expected return $\mu_Q = 12\%$ and volatility $\sigma_Q = 16\%$. T-bills yield 4%. A client asks her about three funds:

Fund	Expected Return	Standard Deviation
Larch Value Fund (LVF)	9%	10%
Monarch Technology Fund (MTF)	16%	30%
Oriel Balanced Fund (OBF)	8%	12%

a. Which of the three funds are efficient?
b. For MTF, find the efficient portfolio with the same volatility and the efficient portfolio with the same expected return. Describe how each is built from Q and T-bills.

Problem 2 (Efficiency, Leverage, and a Constrained Investor)**.** An economy has many risky assets and a risk-free asset that yields 5%. The tangency portfolio Q has an expected return of 20% and a standard deviation of 30%. Two clients of Ashby & Cole Advisors are looking at the following funds:

Fund	Expected Return	Standard Deviation
Bellwether Growth Fund (BGF)	15%	25%
Cedar Income Fund (CIF)	8%	15%

The correlation between BGF and CIF is 0.24.

a. In a (σ, μ) diagram, draw the capital allocation line (CAL) of Q, and plot BGF and CIF. Which portfolios are efficient? Why?
b. Peter Novak wants an efficient portfolio with an expected return of 30%. What should he do? Specify the positions he should take and the standard deviation of his portfolio.
c. Christine Albright manages a parish endowment whose investment policy allows only BGF and CIF, with no T-bills or borrowing. She targets an expected return of 10%. What would you recommend? Give the composition and standard deviation of her portfolio, and say whether it is efficient.

Problem 3 (Inferring the Risk-Free Rate from an Efficient Portfolio)**.** An economy has many risky assets and a risk-free asset. The tangency portfolio Q has an expected return of 20% and a standard deviation of 30%. Peter Novak, who invests optimally, holds a portfolio P_1 with an expected return of 12% and a standard deviation of 15%.

a. Compute the risk-free rate r_f of this economy.
b. What is the composition of Peter's portfolio?
c. His colleague Michael Grant also wants to invest optimally, in a portfolio P_2 with a standard deviation of 40%. What would you recommend? Give the composition of his portfolio and the expected return he can achieve.

Problem 4 (Computing a Beta and an Expected Return)**.** Kingsbridge Rail (KBR) has a return volatility of 30% and a correlation of 0.6 with the tangency portfolio Q. Portfolio Q has expected return $\mu_Q = 12\%$ and volatility $\sigma_Q = 20\%$, and the risk-free rate is 4%. Compute the beta of KBR with respect to Q and its expected return.

Problem 5 (The Beta of a Portfolio)**.** The Tamsin County Retirement System holds three stocks: 50% in Hollis Utilities (HOL), with beta 0.6 relative to the tangency portfolio Q; 30% in Marlowe Retail (MRL), with beta 1.2; and 20% in Vantor Software (VNT), with beta 1.8. The risk-free rate is 3% and $\mu_Q = 10\%$.

a. Compute the expected return of each stock.
b. Compute the beta of the portfolio, and use it to find the portfolio's expected return. Check that it equals the weighted average of the three expected returns.

Problem 6 (Backing Out the Risk-Free Rate and the Risk Premium)**.** Jonah Whitley, an analyst at Calloway Research, has estimated betas relative to an efficient portfolio Q for two stocks. Alcott Foods (ALF) has beta 0.5 and expected return 7%, and Brennan Mining (BRM) has beta 1.5 and expected return 13%. Use beta pricing to find the risk-free rate and the expected return of Q.

Problem 7 (Residual Risk Earns No Premium)**.** Thornfield Pharmaceuticals (TFP) has a return volatility of 40% and a beta of 0.8 relative to the tangency portfolio Q, which has $\mu_Q = 12\%$ and $\sigma_Q = 20\%$. The risk-free rate is 4%.

a. Compute the expected return of TFP.
b. Split TFP's variance into the part explained by Q and the residual part. What is the volatility of the residual ε_{TFP}?
c. Find the efficient portfolio with the same expected return as TFP. What is its volatility, and what does the comparison say about the reward for bearing residual risk?

Problem 8 (An Asset with a Negative Beta)**.** Bastion Capital runs a tail-hedge fund that buys protection against market crashes, so its returns rise when the tangency portfolio Q falls. Its beta relative to Q is -0.4. The risk-free rate is 4% and $\mu_Q = 12\%$. Compute the fund's expected return. Why would anyone hold an asset that is expected to earn less than T-bills?

Problem 9 (A Forecast That Contradicts Beta Pricing)**.** Dana Kimura, a portfolio manager at Ellery Investments, treats a broad stock index as the efficient portfolio Q, with $\mu_Q = 11\%$, and the risk-free rate is 4%. Her analysts estimate that Wexford Logistics (WXL) has a beta of 1.2 with the index, and they forecast an expected return of 14% for the stock.

a. What expected return does beta pricing assign to WXL? By how much does the analysts' forecast exceed it?
b. Property 1 says that every asset satisfies the beta pricing relation when Q is efficient. If the forecast is right, what does it imply about the index? Should Dana hold more or less WXL than the index does?