This chapter examines how combining risky assets shapes a portfolio's risk and return. The central insight is that portfolio risk is not simply the average of its components' risks: because asset returns do not move in lockstep, combining them can push risk below what any single asset offers on its own. This effect, known as **diversification**, is one of the most fundamental ideas in modern finance.

We begin with two risky assets to build intuition, deriving the **investment opportunity set** they generate and identifying the portfolio within it that achieves the lowest possible risk, the **minimum-variance portfolio**. We then examine how diversification behaves as more assets are added to a portfolio, and generalize the two-asset formulas to any number of risky assets using matrix notation. This will allow us to characterize the investment opportunity set achievable with many risky assets.

A Portfolio of Two Risky Assets

Suppose we have two risky assets A and B. The return of a portfolio P in which we invest $1 - w$ in A and w in B is

$$r_P = (1 - w)r_A + wr_B. \tag{1}$$

The expected return of the portfolio is

$$\mu_P = (1 - w)\mu_A + w\mu_B, \tag{2}$$

whereas its variance can be computed as

$$\begin{aligned} \sigma_P^2 &= \mathrm{V}((1 - w)r_A + wr_B) \\ &= (1 - w)^2\sigma_A^2 + w^2\sigma_B^2 + 2w(1 - w)\sigma_{A,B}. \end{aligned} \tag{3}$$

Here $\sigma_{A,B}$ denotes the covariance of returns between A and B.

Example 1 (A Portfolio with Two Risky Assets)**.** Consider two risky assets A and B for which you have the following information.

Asset	Expected Return	Standard Deviation
A	10%	20%
B	15%	35%

The correlation between the assets returns is 0.4. If you invest 40% in A and 60% in B, your portfolio will have an expected return of

$$\mu = 0.4 \times 0.10 + 0.6 \times 0.15 = 13\%.$$

Since the covariance of returns is $\sigma_{AB} = 0.20 \times 0.35 \times 0.4 = 0.028$, the standard deviation of the portfolio returns is

$$\sigma = \sqrt{0.4^2 \times 0.20^2 + 0.6^2 \times 0.35^2 + 2 \times 0.4 \times 0.6 \times 0.028} = 25.29\%.$$

□

Suppose that now we want to vary w and see all the possible combinations of expected return (μ_P) and standard deviation (σ_P) that we can obtain. In the previous example, we split the wealth between the two assets, so that $0 \leq w \leq 1$ and $0 \leq 1 - w \leq 1$. What if we wanted to invest more than 100% in asset B?

In financial markets, it is possible to borrow an asset and sell it immediately, a transaction known as **short-selling** the asset. There is typically no fixed date by which the asset must be returned to the lender; rather, it must be returned whenever the lender demands it.[1]

We assume throughout this class that short-selling is allowed. This means the weight w can be greater than one, in which case we short A to overinvest in B, or less than zero, in which case we short B to overinvest in A.

By varying w from $-\infty$ to ∞, we obtain what is called the **investment opportunity set** generated by the two risky assets. If we plot μ_P on the y-axis as a function of σ_P on the x-axis, the resulting figure is a hyperbola.

[1] In practice, the lender (typically a broker) can recall the borrowed shares at their discretion, forcing the short-seller to buy them back on short notice. The short-seller must also compensate the lender for any dividends paid on the asset while the short position is open.

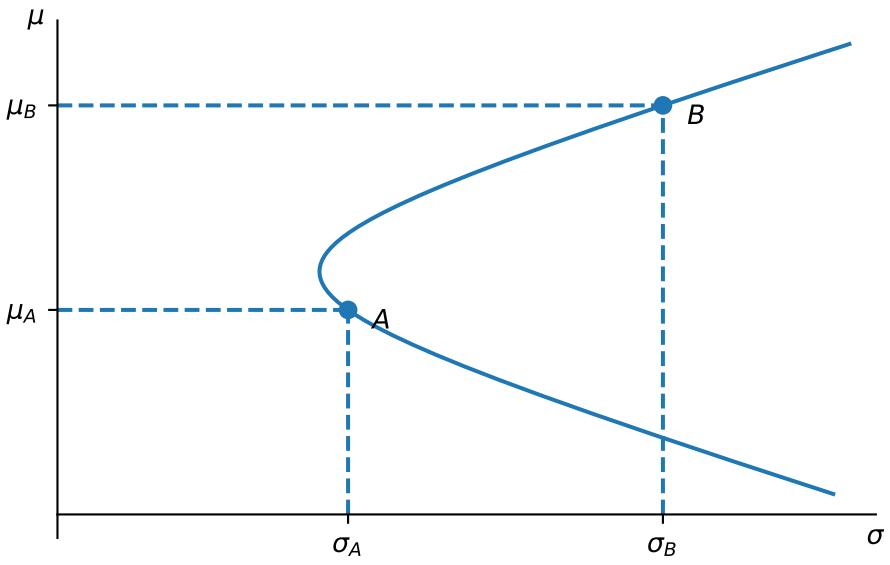


Figure 1: The figure shows the investment opportunity set generated by two risky assets A and B.

The Minimum Variance Portfolio (MVP)

As shown in the previous plot, there is a portfolio that has the minimum variance among all portfolios between A and B. To find its composition, we can use standard optimization techniques:

$$\begin{aligned}\frac{d}{dw}\sigma_P^2 &= \frac{d}{dw}(1-w)^2\sigma_A^2 + w^2\sigma_B^2 + 2w(1-w)\sigma_{A,B} \\ &= -2(1-w)\sigma_A^2 + 2w\sigma_B^2 + 2(1-2w)\sigma_{A,B} = 0.\end{aligned} \tag{4}$$

Thus,

$$-(1-w)\sigma_A^2 + w\sigma_B^2 + (1-2w)\sigma_{A,B} = 0.$$

or

$$w_{MV} = \frac{\sigma_A^2 - \sigma_{A,B}}{\sigma_A^2 + \sigma_B^2 - 2\sigma_{A,B}}. \tag{5}$$

Property 1 (The Minimum Variance Portfolio)**.** *Given two risky assets A and B, there is a portfolio that has the minimum variance among all possible portfolios that can be built with A and B. The*

weights of the minimum variance portfolio are given by

$$w_A = \frac{\sigma_B^2 - \sigma_{AB}}{\sigma_A^2 + \sigma_B^2 - 2\sigma_{AB}},$$
$$w_B = \frac{\sigma_A^2 - \sigma_{AB}}{\sigma_A^2 + \sigma_B^2 - 2\sigma_{AB}}.$$

In the previous expression, σ_A and σ_B denote the standard deviation or volatility of returns. The term σ_{AB} denotes the covariance of A and B and is equal to

$$\sigma_{AB} = \sigma_A \sigma_B \rho_{AB},$$

where ρ_{AB} is the correlation between the returns of A and B.

Example 2 (Minimum Variance Portfolio)**.** Using the data of Example 1, we find that the weights of the minimum variance portfolio are given by

$$w_A = \frac{0.35^2 - 0.028}{0.20^2 + 0.35^2 - 2 \times 0.028} = 88.73\%$$
$$w_B = \frac{0.20^2 - 0.028}{0.20^2 + 0.35^2 - 2 \times 0.028} = 11.27\%$$

Thus, the expected return and volatility of the minimum variance portfolio are

$$\mu = 0.8873 \times 0.10 + 0.1127 \times 0.15 = 10.56\%,$$
$$\sigma = \sqrt{0.8873^2 \times 0.20^2 + 0.1127^2 \times 0.35^2 + 2 \times 0.8873 \times 0.1127 \times 0.028}$$
$$= 19.66\%.$$

□

The MVP is the point farthest to the left of the investment opportunity set.

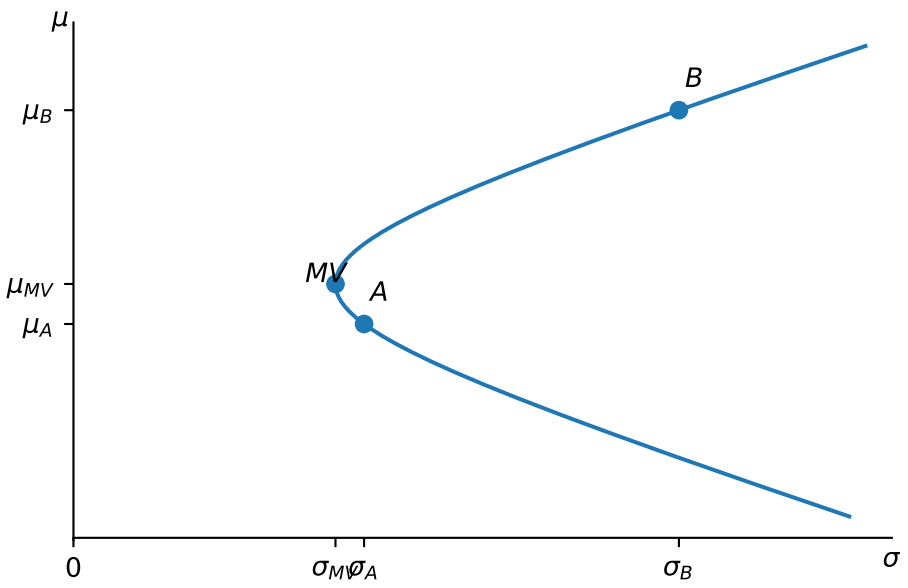


Figure 2: The figure shows the minimum variance portfolio obtained from combining two risky assets.

An Equally-Weighted Portfolio

The minimum-variance frontier demonstrates that combining two risky assets can alter the risk-return profile of the original assets. This concept is known as **portfolio diversification**.

To understand this phenomenon better, say that we have N securities and we invest the same amount in each, so the return of this portfolio (that we call P) is

$$r_P = \frac{1}{N}\sum_{i=1}^{N} r_i.$$

The variance of P is given by

$$\sigma_P^2 = \frac{1}{N^2}\sum_{i=1}^{N}\sum_{j=1}^{N} \text{Cov}(r_i, r_j). \tag{6}$$

The previous expression looks complex, but it is possible to decompose it into two intuitive

parts. First, note that the average variance of the securities is given by

$$\text{Avg. Variance} = \frac{1}{N}\sum_{i=1}^{N} \mathrm{V}(r_i).$$

Furthermore, the average covariance of the securities can be computed as

$$\text{Avg. Covariance} = \frac{1}{N(N-1)}\sum_{i=1}^{N}\sum_{\substack{j=1 \\ i\neq j}}^{N} \mathrm{Cov}(r_i, r_j)$$

Using these two quantities, we can write (6) as

$$\sigma_P^2 = \frac{1}{N}(\text{Average Variance}) + \frac{N-1}{N}(\text{Average Covariance}). \tag{7}$$

As N increases in (7), we have that

$$\sigma_P^2 \xrightarrow[N\to\infty]{} \text{Average Covariance}.$$

The average variance disappears! This is the part of the risk that can be diversified away. What remains is the average covariance, which sets a floor on how much risk diversification can remove.

The 2008 financial crisis illustrated this limit. Correlations among mortgage-related assets that had appeared only loosely related rose sharply as the crisis unfolded, so portfolios that seemed well diversified suffered simultaneous losses.

A Portfolio of Many Risky Assets

The two-asset formulas above are a special case of a more general framework. Suppose there are N risky assets with returns $r_1, r_2, \ldots, r_N$. Collect the returns and the portfolio weights into

vectors

$$\mathbf{r} = \begin{pmatrix} r_1 \\ r_2 \\ \vdots \\ r_N \end{pmatrix}, \qquad \boldsymbol{w} = \begin{pmatrix} w_1 \\ w_2 \\ \vdots \\ w_N \end{pmatrix},$$

where w_i is the fraction of wealth invested in asset i. Since short-selling is allowed, each w_i can be any real number, but the weights must still add up to the investor's total wealth,

$$\boldsymbol{w}^{\top}\mathbf{1} = 1, \tag{8}$$

where **1** denotes the N-dimensional vector of ones. The portfolio return is the linear combination

$$r_P = \boldsymbol{w}^{\top}\mathbf{r} = \sum_{i=1}^{N} w_i r_i. \tag{9}$$

By linearity of expectation, the expected return is $\mu_P = \sum_{i=1}^{N} w_i \mu_i = \boldsymbol{w}^{\top}\boldsymbol{\mu}$, where $\boldsymbol{\mu} = (\mu_1, \ldots, \mu_N)^{\top}$ and $\mu_i = \mathrm{E}(r_i)$. The variance, however, involves every pairwise covariance among the assets: writing $r_P - \mu_P = \sum_i w_i (r_i - \mu_i)$ and squaring,

$$\sigma_P^2 = \mathrm{E}\left[\left(\sum_{i=1}^{N} w_i (r_i - \mu_i)\right)^2\right] = \sum_{i=1}^{N}\sum_{j=1}^{N} w_i w_j \,\mathrm{Cov}(r_i, r_j).$$

This double sum is the N-asset generalization of (3); keeping track of its N variances and $N(N-1)/2$ distinct covariances quickly becomes unwieldy as N grows.

Matrix notation packages this double sum compactly. Denote by $\boldsymbol{\Sigma}$ the **covariance matrix** of **r**, the $N \times N$ symmetric matrix whose (i, j) entry is $\sigma_{ij} = \mathrm{Cov}(r_i, r_j)$, with $\sigma_{ii} = \sigma_i^2$ on the diagonal. The double sum above is exactly the quadratic form $\boldsymbol{w}^{\top}\boldsymbol{\Sigma}\boldsymbol{w}$, so

$$\mu_P = \boldsymbol{w}^{\top}\boldsymbol{\mu}, \qquad \sigma_P^2 = \boldsymbol{w}^{\top}\boldsymbol{\Sigma}\boldsymbol{w}. \tag{10}$$

The same expand-and-take-expectations argument, applied to the covariance of two different

portfolios P and Q with weight vectors $\boldsymbol{w}_P$ and $\boldsymbol{w}_Q$, gives[2]

$$\text{Cov}(r_P, r_Q) = \boldsymbol{w}_P^{\top} \boldsymbol{\Sigma} \boldsymbol{w}_Q, \tag{11}$$

which reduces to (10) when $Q = P$.

Equations (9) and (10) replace (2) and (3) for any number of assets, and collapse back to the two-asset formulas once $N = 2$.[3]

Example 3 (A Portfolio of Three Risky Assets)**.** Consider three risky assets A, B, and C with the following expected returns and standard deviations.

Asset	Expected Return	Standard Deviation
A	8%	15%
B	12%	20%
C	15%	25%

The correlations between their returns are $\rho_{AB} = 0.3$, $\rho_{AC} = 0.1$, and $\rho_{BC} = 0.4$, so the

[2] For any weight vectors $\boldsymbol{a}$ and $\boldsymbol{b}$,

$$\text{Cov}(\boldsymbol{a}^{\top}\mathbf{r}, \boldsymbol{b}^{\top}\mathbf{r}) = \sum_{i=1}^{N}\sum_{j=1}^{N} a_i b_j \,\text{Cov}(r_i, r_j) = \boldsymbol{a}^{\top}\boldsymbol{\Sigma}\boldsymbol{b},$$

by the same argument used above for σ_P^2.

[3] With $N = 2$, set $\boldsymbol{w} = (1 - w, w)^{\top}$ and

$$\boldsymbol{\Sigma} = \begin{pmatrix} \sigma_A^2 & \sigma_{A,B} \\ \sigma_{A,B} & \sigma_B^2 \end{pmatrix}.$$

Expanding the quadratic form gives

$$\begin{aligned} \boldsymbol{w}^{\top}\boldsymbol{\Sigma}\boldsymbol{w} &= \begin{pmatrix} 1-w & w \end{pmatrix} \begin{pmatrix} \sigma_A^2 & \sigma_{A,B} \\ \sigma_{A,B} & \sigma_B^2 \end{pmatrix} \begin{pmatrix} 1-w \\ w \end{pmatrix} \\ &= (1-w)^2\sigma_A^2 + w^2\sigma_B^2 + 2w(1-w)\sigma_{A,B}, \end{aligned}$$

which is exactly (3).

covariance matrix is

$$\boldsymbol{\Sigma} = \begin{pmatrix} 0.0225 & 0.0090 & 0.00375 \\ 0.0090 & 0.0400 & 0.0200 \\ 0.00375 & 0.0200 & 0.0625 \end{pmatrix}.$$

If you invest $\boldsymbol{w} = (0.3, 0.3, 0.4)^\top$ in the three assets, (10) gives an expected return of

$$\mu_P = \boldsymbol{w}^\top \boldsymbol{\mu} = 12.0\%$$

and a variance of $\sigma_P^2 = \boldsymbol{w}^\top \boldsymbol{\Sigma} \boldsymbol{w} = 0.0229$, so that $\sigma_P = 15.15\%$. Computing μ_P and σ_P only requires matrix multiplication, regardless of how many assets are in the portfolio. □

The Investment Opportunity Set with Many Risky Assets

With two risky assets, tracing out (σ_P, μ_P) as w ranges over all real numbers produces a single curve: each value of w pins down one portfolio, and there is exactly one free parameter to vary. With N risky assets, the weights $\boldsymbol{w}$ still have to satisfy the single budget constraint (8), which leaves $N - 1$ degrees of freedom. For $N \geq 3$, there are therefore many different portfolios that share the same expected return μ_P but differ in variance σ_P^2. As a result, the set of achievable pairs (σ_P, μ_P), the **investment opportunity set**, is no longer a curve but a filled two-dimensional region.

That region still has a well-defined left boundary, however. For each attainable level of expected return m, exactly one portfolio achieves it with the smallest possible variance; every other portfolio with that same expected return lies strictly to its right. Collecting these minimum-variance points as m varies traces out a curve called the **minimum-variance frontier**, and every portfolio in the investment opportunity set lies on or to the right of it.

We find the frontier by adding a target-return constraint to the variance minimization. For a target expected return m, we minimize

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

Introducing a Lagrange multiplier for each constraint and following the same steps as for the minimum-variance portfolio[4] gives a frontier variance of

$$\sigma_P^2(m) = \frac{Cm^2 - 2Am + B}{D}, \tag{12}$$

where

$$A = \boldsymbol{\mu}^\top \boldsymbol{\Sigma}^{-1} \mathbf{1}, \qquad B = \boldsymbol{\mu}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}, \qquad C = \mathbf{1}^\top \boldsymbol{\Sigma}^{-1} \mathbf{1}, \qquad D = BC - A^2.$$

Equation (12) is a parabola in (σ_P^2, μ_P)-space, which is exactly why the frontier traces out a hyperbola in (σ_P, μ_P)-space, generalizing the hyperbola we found earlier with only two assets. Its minimum over m occurs at $m = A/C$, with $\sigma_P^2 = 1/C = 1/(\mathbf{1}^\top \boldsymbol{\Sigma}^{-1} \mathbf{1})$, which is the minimum-variance portfolio derived in the next section.

The figure below illustrates this for the three assets from Example 3. Each gray dot is a portfolio obtained by perturbing the equally-weighted portfolio with a random shock that is symmetric across the three assets and averages to zero, so the budget constraint continues to hold. Points near the frontier are denser than points far from it: moving away from the frontier at a fixed expected return raises the variance quadratically, so interior points require much larger short and levered positions.

The Minimum-Variance Portfolio with Many Risky Assets

With many risky assets, finding the minimum-variance portfolio means solving

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

[4] With two constraints, the Lagrangian is $L(\boldsymbol{w}, \lambda, \gamma) = \boldsymbol{w}^\top \boldsymbol{\Sigma} \boldsymbol{w} - 2\lambda(\boldsymbol{w}^\top \boldsymbol{\mu} - m) - 2\gamma(\boldsymbol{w}^\top \mathbf{1} - 1)$. The first-order condition gives $\boldsymbol{w} = \boldsymbol{\Sigma}^{-1}(\lambda \boldsymbol{\mu} + \gamma \mathbf{1})$, and substituting into the two constraints yields the linear system

$$\begin{pmatrix} B & A \\ A & C \end{pmatrix} \begin{pmatrix} \lambda \\ \gamma \end{pmatrix} = \begin{pmatrix} m \\ 1 \end{pmatrix},$$

so that $\lambda = (Cm - A)/D$ and $\gamma = (B - Am)/D$. Plugging $\boldsymbol{w}$ back into $\boldsymbol{w}^\top \boldsymbol{\Sigma} \boldsymbol{w} = \lambda m + \gamma$, using the two constraints, gives (12).

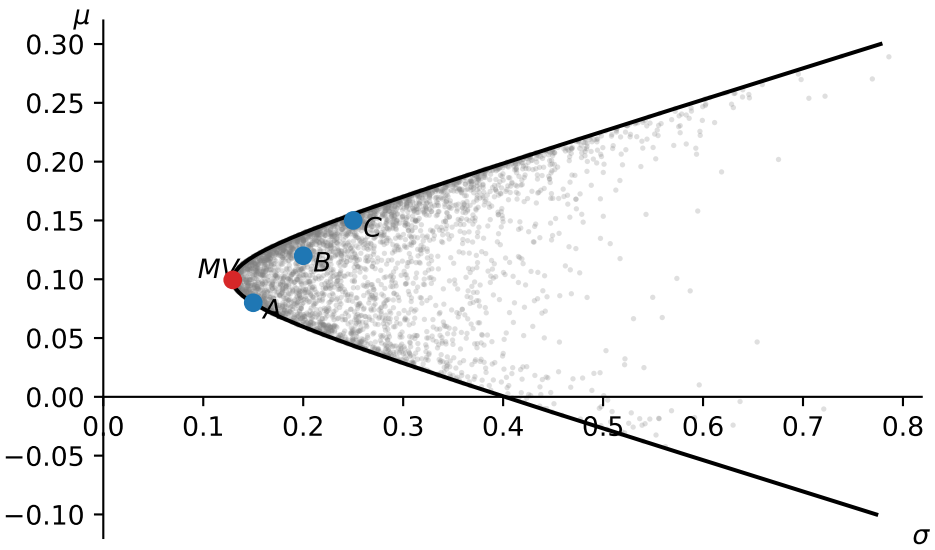


Figure 3: The figure shows simulated portfolios of the three risky assets from Example 3, obtained by perturbing the equally-weighted portfolio with a random shock, together with the minimum-variance frontier and the minimum-variance portfolio.

i.e., minimizing (10) subject to the budget constraint (8). Introducing a Lagrange multiplier λ for the constraint, define the Lagrangian

$$L(\boldsymbol{w}, \lambda) = \boldsymbol{w}^\top \boldsymbol{\Sigma} \boldsymbol{w} - 2\lambda \left(\boldsymbol{w}^\top \mathbf{1} - 1 \right),$$

where the factor of 2 is just a convenient normalization for λ. The first-order condition with respect to $\boldsymbol{w}$ is[5]

$$2\boldsymbol{\Sigma}\boldsymbol{w} - 2\lambda \mathbf{1} = \mathbf{0} \quad \Longrightarrow \quad \boldsymbol{w} = \lambda \boldsymbol{\Sigma}^{-1} \mathbf{1},$$

where $\boldsymbol{\Sigma}^{-1}$ exists because $\boldsymbol{\Sigma}$ is positive definite: $\boldsymbol{w}^\top \boldsymbol{\Sigma} \boldsymbol{w}$ is the variance of the portfolio $\boldsymbol{w}$, which is positive as long as no combination of the risky assets is riskless. Substituting into the budget constraint pins down the multiplier,

$$\lambda \mathbf{1}^\top \boldsymbol{\Sigma}^{-1} \mathbf{1} = 1 \quad \Longrightarrow \quad \lambda = \frac{1}{\mathbf{1}^\top \boldsymbol{\Sigma}^{-1} \mathbf{1}},$$

which, substituted back, gives the following result.

[5] Differentiating $\boldsymbol{w}^\top \boldsymbol{\Sigma} \boldsymbol{w}$ with respect to $\boldsymbol{w}$ gives $2\boldsymbol{\Sigma}\boldsymbol{w}$ because $\boldsymbol{\Sigma}$ is symmetric, and differentiating $\boldsymbol{w}^\top \mathbf{1}$ gives $\mathbf{1}$.

Property 2 (The Minimum-Variance Portfolio with Many Assets)**.** *Given N risky assets with covariance matrix $\boldsymbol{\Sigma}$, the minimum-variance portfolio has weights*

$$\boldsymbol{w}_{MV} = \frac{\boldsymbol{\Sigma}^{-1}\mathbf{1}}{\mathbf{1}^{\top}\boldsymbol{\Sigma}^{-1}\mathbf{1}} \tag{13}$$

and variance

$$\sigma^2_{MV} = \frac{1}{\mathbf{1}^{\top}\boldsymbol{\Sigma}^{-1}\mathbf{1}}. \tag{14}$$

To see why the variance takes this simple form, plug $\boldsymbol{w}_{MV} = \lambda\boldsymbol{\Sigma}^{-1}\mathbf{1}$ into (10). Since $\boldsymbol{\Sigma}\boldsymbol{w}_{MV} = \lambda\mathbf{1}$,

$$\sigma^2_{MV} = \boldsymbol{w}^{\top}_{MV}\boldsymbol{\Sigma}\boldsymbol{w}_{MV} = \boldsymbol{w}^{\top}_{MV}(\lambda\mathbf{1}) = \lambda\left(\boldsymbol{w}^{\top}_{MV}\mathbf{1}\right) = \lambda,$$

using the budget constraint $\boldsymbol{w}^{\top}_{MV}\mathbf{1} = 1$ in the last step, which is exactly (14).

Equation (13) is the general version of (5): with $N = 2$ it reduces exactly to the weights w_A and w_B from Property 1.[6]

Example 4 (The Minimum-Variance Portfolio with Three Assets)**.** Using the covariance matrix $\boldsymbol{\Sigma}$ from Example 3, (13) gives minimum-variance weights of 64.07% in A, 19.05% in B, and 16.88% in C. This portfolio has an expected return of $\mu_{MV} = 9.94\%$ and, by (14), a variance of $\sigma^2_{MV} = 0.0168$, i.e., a standard deviation of $\sigma_{MV} = 12.95\%$. Notice that σ_{MV} is lower than the volatility of any of the three individual assets, illustrating diversification with more than two risky assets. □

The minimum-variance portfolio also satisfies a remarkable property, valid for any number of assets. Let $\boldsymbol{w}_P$ be the weights of an arbitrary portfolio P, so that $\boldsymbol{w}^{\top}_P\mathbf{1} = 1$. By (11) and

---

[6] With $N = 2$ and $\boldsymbol{\Sigma} = \begin{pmatrix} \sigma^2_A & \sigma_{AB} \\ \sigma_{AB} & \sigma^2_B \end{pmatrix}$, the formula for the inverse of a 2×2 matrix gives

$$\boldsymbol{\Sigma}^{-1}\mathbf{1} = \frac{1}{\sigma^2_A\sigma^2_B - \sigma^2_{AB}}\begin{pmatrix} \sigma^2_B - \sigma_{AB} \\ \sigma^2_A - \sigma_{AB} \end{pmatrix},$$

so that $\mathbf{1}^{\top}\boldsymbol{\Sigma}^{-1}\mathbf{1} = \dfrac{\sigma^2_A + \sigma^2_B - 2\sigma_{AB}}{\sigma^2_A\sigma^2_B - \sigma^2_{AB}}$, and dividing each entry of $\boldsymbol{\Sigma}^{-1}\mathbf{1}$ by this quantity recovers exactly w_A and w_B from Property 1.

$\boldsymbol{w}_{MV} = \lambda \boldsymbol{\Sigma}^{-1}\mathbf{1}$,

$$\text{Cov}(r_P, r_{MV}) = \boldsymbol{w}_P^\top \boldsymbol{\Sigma} \boldsymbol{w}_{MV} = \boldsymbol{w}_P^\top \boldsymbol{\Sigma} \left(\lambda \boldsymbol{\Sigma}^{-1}\mathbf{1}\right) = \lambda \left(\boldsymbol{w}_P^\top \mathbf{1}\right) = \lambda = \sigma_{MV}^2,$$

using the budget constraint $\boldsymbol{w}_P^\top \mathbf{1} = 1$ and (14) in the last two steps. The covariance of the MVP with any other portfolio is always the same and equal to its own variance! This also implies that the correlation of the MVP with any other portfolio is always positive, regardless of how many assets are in the investment universe.

Example 5 (Computing a Correlation with the MVP)**.** Consider an investment opportunity set where the MVP has a standard deviation of returns of 20%. An asset A has a standard deviation of returns equal to 30%. Thus,

$$\sigma_A \sigma_{MV} \rho_{A,MV} = \sigma_{A,MV} = \sigma_{MV}^2,$$

or

$$\rho_{A,MV} = \frac{\sigma_{MV}}{\sigma_A} = \frac{2}{3}.$$

□

Equations (5) and (13) demonstrate that the weights of the MVP depend solely on variances and covariances, which can typically be estimated with high precision from the data. In contrast, constructing any other portfolio within the investment opportunity set requires estimating expected returns, which are generally more challenging to estimate accurately. This is the reason why some practitioners like to use the MVP for their asset allocation. Moreover, forbidding short sales makes the estimated minimum-variance portfolio more robust. With fewer months of returns than stocks, the sample covariance matrix is singular and the textbook MVP does not exist; ruling out short sales removes the long-short positions that depend on the directions the data cannot pin down, so the problem has a solution. When the sample matrix is invertible, Jagannathan and Ma (2003) show that the constraint is equivalent to shrinking the largest covariance estimates, and that once short sales are ruled out, the sample covariance matrix performs as well as covariance matrix estimates based on factor models, shrinkage estimators, and daily data.

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/.

Problem 1 (Covariance and Correlation of Two Securities)**.** Oliver Grant believes the economy has two equally likely states next month: boom (the good state) and recession (the bad state). At the beginning of the month, he can buy the following securities:

- Security 1 is a share of Sunvale Resorts (SVR), which trades at \$4. Its price increases by \$10 in the good state and remains unchanged in the bad state.
- Security 2 is a recession-protection contract sold by Bradwell Insurance, which trades at \$5. It pays nothing extra in the good state and is worth \$10 more in the bad state.

a. Compute the expected returns of securities 1 and 2.
b. Compute the standard deviations of returns for securities 1 and 2.
c. Compute the covariance and the correlation between the returns of the two securities.

Problem 2 (Expected Return and Volatility of a Two-Asset Portfolio)**.** Tomás Rivera is choosing between two mutual funds, the Dynamic Growth Fund (D) and the Emerging Opportunities Fund (E). Fund D is the steadier of the pair, with an expected return of $\mu_D = 12\%$ and a volatility of $\sigma_D = 18\%$, while Fund E offers a higher expected return of $\mu_E = 18\%$ but comes with substantially more risk, $\sigma_E = 30\%$. The two funds' returns are only mildly related, with a correlation of $\rho_{DE} = 0.25$. He decides to build a portfolio that puts 60% of his wealth in D and the remaining 40% in E. What expected return and volatility should he expect from this portfolio?

Problem 3 (The Minimum-Variance Portfolio with Negatively Correlated Assets)**.** Beatrice Hall holds shares of Anchor Steel Co. (STL), a stock with a volatility of $\sigma_{STL} = 25\%$ whose returns have a correlation of 0.20 with the S&P 500. To offset some of that market exposure, she considers adding a leveraged inverse ETF (INV) designed to deliver twice the opposite of the

index's daily return. The ETF has a volatility of $\sigma_{INV} = 40\%$, and since it moves against the index, its correlation with STL is $\rho_{STL,INV} = -0.20$. Find the weights of the minimum-variance portfolio of STL and INV, and compute its volatility. Compare this volatility to σ_{STL} and σ_{INV}.

Problem 4 (Diversification with Many Equally-Weighted Assets)**.** Kenji Watanabe, a portfolio manager at Ravenhill Partners, follows a large universe of small biotechnology stocks, each with a variance of returns of 0.36 (i.e., a volatility of 60%) and an average covariance between any two stocks of 0.09. If Kenji forms an equally-weighted portfolio of $N = 25$ of these stocks, what is the portfolio's volatility? What volatility does the portfolio approach as $N \to \infty$, and what does this limit represent?

Problem 5 (A Three-Asset Portfolio Using the Covariance Matrix)**.** Priya Desai is building a portfolio out of three stocks that span a range of risk and return. Pembroke Utilities (PBU) is the most conservative of the three, with an expected return of $\mu_{PBU} = 9\%$ and a volatility of $\sigma_{PBU} = 12\%$. Quarry Materials (QRM) sits in the middle, offering $\mu_{QRM} = 14\%$ at a volatility of $\sigma_{QRM} = 22\%$. Zephyr Biotech (ZPB) is the most aggressive, with $\mu_{ZPB} = 20\%$ and $\sigma_{ZPB} = 30\%$. The three stocks are not equally related to one another: PBU and QRM have a modest correlation of 0.2, PBU and ZPB are uncorrelated, and QRM and ZPB move together more closely, with a correlation of 0.5. Priya settles on the weights $\boldsymbol{w} = (0.5, 0.3, 0.2)^{\top}$ in PBU, QRM, and ZPB. Compute the expected return and volatility of the resulting portfolio.

Problem 6 (Correlation with the Minimum-Variance Portfolio)**.** The Steadfast Minimum Volatility Fund holds the minimum-variance portfolio of a large stock universe, and its volatility is 15%. Orrin Therapeutics (ORT), a biotech stock in that universe, is far more volatile, with a volatility of 45%. What is the correlation between the returns of ORT and the returns of the fund?

Problem 7 (Comparative Statics: Correlation and Minimum-Variance Risk)**.** Helena Brooks, the risk manager of the Marlow Foundation, wants to hold the minimum-variance mix of two funds, the Aldgate Fund (A) with a volatility of $\sigma_A = 20\%$ and the Brookfield Fund (B) with $\sigma_B = 30\%$. Correlations between asset classes shift over time, and Helena's estimate of the correlation between the two funds depends on the sample: data from one decade give $\rho = -0.3$, while data from another give $\rho = 0.8$.

a. Compute the minimum-variance weights and volatility under each correlation estimate.
b. Under which estimate is the minimum-variance portfolio less risky? Explain how each portfolio reduces risk, paying attention to the signs of the weights.
c. Recompute the minimum-variance weights and volatility with $\rho = 0.97$. Does a lower correlation always lead to a less risky minimum-variance portfolio?

Problem 8 (A Pension Fund's Minimum-Variance Portfolio)**.** Gloria Nash, who manages the Tri-County Firefighters' Pension Fund, is considering two mutual funds from Crestline Investments: a stock fund and a long-term government and corporate bond fund. The probability distribution of the funds is as follows:

Fund	Expected Return	Standard Deviation
Stock fund (S)	20%	30%
Bond fund (B)	12%	15%

The correlation between the fund returns is 0.10.

a. What are the investment proportions in the minimum-variance portfolio of the two funds?

b. What is the expected value and standard deviation of the minimum-variance portfolio's rate of return?

c. Compute the expected return and standard deviation of the following portfolios.

Proportion in S	Proportion in B	Expected Return	Standard Deviation
25%	75%		
50%	50%		
75%	25%		

Problem 9 (Adding a Concentrated Stock Position)**.** Daniel Reyes manages a $4,000,000 fully diversified portfolio for a client. The client unexpectedly receives a bonus of $1,000,000 in vested shares of her employer, Meridian Corp. Daniel has produced the following forecasts:

Asset	Expected Return	Standard Deviation
Original Portfolio	8.5%	10%
Meridian Corp	18%	22%

The correlation coefficient between Meridian Corp and the original portfolio is 0.30.

a. The bonus changes the client's overall portfolio, and she is deciding whether to keep the Meridian shares. Assuming she keeps them, calculate the expected return and standard deviation of her new $5,000,000 portfolio, which includes the Meridian shares.
b. If she sells the Meridian shares, she will invest the proceeds in risk-free Treasury bills yielding 4.5%. Assuming she sells the shares and replaces them with the Treasury bills, calculate the expected return and standard deviation of her new portfolio.
c. One of Daniel's colleagues suggests that instead of buying Treasury bills, she sell the Meridian shares and use the $1,000,000 to buy shares of Atlas Corp, which has the same expected return and standard deviation as Meridian Corp. The colleague argues, "It doesn't matter whether she keeps the Meridian shares or switches to Atlas instead." State whether the colleague's comment is correct or incorrect, and justify your answer briefly.

Jagannathan, Ravi, and Tongshu Ma. 2003. "Risk Reduction in Large Portfolios: Why Imposing the Wrong Constraints Helps." *Journal of Finance* 58 (4): 1651–83.