



**Research Paper**

**March 2026**

**Portfolio Optimization**

**7 Steps to Best Practice**

**Redouane Raki**

**redouane.raki@saa-forum.com**

# Table of Contents

## Management Summary

### 1 Model Assumptions

- 1.1 Capital Market Model
- 1.2 Investor Profile

### 2 The Optimal Portfolio

- 2.1 Theorem
- 2.2 Proof
- 2.3 Corollary
- 2.4 Properties
- 2.5 Instability of the Covariance Matrix

### 3 7-Step Best Practice Code

- 3.1 Perform Empirical Optimization
- 3.2 Formulate Economically Grounded Return Views
- 3.3 Adjust the Covariance Matrix
- 3.4 Assess and, if Necessary, Stabilize the Covariance Matrix
- 3.5 Choose Constraints Deliberately
- 3.6 Conduct Sensitivity and Robustness Analyses
- 3.7 Interpret and Validate the Portfolio

### 4 References

## Management Summary

The mean-variance solution under linear constraints can be interpreted as a linear transformation of expected returns through the inverse covariance matrix. This structure implies that optimization not only processes information but also amplifies it. In particular, inversion of the covariance matrix acts as a sensitivity operator, so that even small estimation errors in covariances can have disproportionately large effects on the resulting portfolio weights.

While uncertainty in expected returns enters the optimization directly, errors in the covariance matrix affect the solution at a structural level and are further amplified by matrix inversion. Consequently, mean-variance optimization is inherently vulnerable to estimation errors and, if applied uncritically, can produce unstable portfolios that are difficult to interpret economically.

Optimization therefore amplifies the underlying model assumptions: high-quality, consistent, and robust inputs lead to stable and plausible portfolio allocations, whereas flawed or inconsistent inputs produce fragile and potentially misleading results. Against this background, a central objective of this paper is to develop a structured best practice code that translates the theoretical properties of the optimization problem into actionable recommendations.

The paper proceeds in three stages. First, capital market assumptions and the investor profile are formalized within a multivariate distribution model. Next, the optimal portfolio under budget and linear constraints is derived analytically using Lagrange methods and interpreted both algebraically and geometrically. Building on these insights, the theoretical findings are translated into a comprehensive seven-step best practice code addressing, in particular, covariance matrix conditioning, disciplined return assumptions, structural consistency, sensitivity analysis, and regularization.

Overall, this paper demonstrates that portfolio optimization should not be viewed as a purely technical procedure but rather as a structurally sensitive mapping problem whose validity depends critically on the numerical stability, statistical robustness, and economic consistency of the underlying assumptions.

# 1 Model Assumptions

## 1.1 Capital Market Model

The capital market model is represented by an investment universe comprising  $n$  assets or asset classes.

Each asset  $i$  has a time series of empirical return vectors  $R_i$ .

These  $n$  vectors yield means  $\mu_i$ , volatilities  $\sigma_i$ , correlations  $r(i,j)$ , and covariances  $c(i,j)$ . In aggregate, we consider an  $n$ -dimensional return distribution  $R(m,C)$ .<sup>1</sup>

Alternatively, a covariance matrix can be expressed as the product of volatilities and correlations:  $C = D p D$ , where  $D = \text{diag}(\sigma_i, i=1, \dots, n)$  is the diagonal matrix of volatilities and  $p$  is the correlation matrix.

This representation allows the individual components of the risk structure to be analyzed separately.

## 1.2 Investor Profile

In capital markets, the principle is that higher returns require greater risk, and vice versa. This means that the pursuit of higher returns, or the willingness to assume more risk, should be compensated by greater risk exposure or higher expected returns, respectively. From an investor's perspective, returns are desirable, while risk is undesirable. This gives rise to the investment dilemma: How much risk should an investor take in light of their return expectations? This dilemma can be expressed through a utility function:<sup>2</sup>

$$\text{Utility} = \text{Investor's Expected Return} - \text{Investor's Risk Aversion}$$

Every rational investor faces the challenge of maximizing their own utility function. While expected returns can be quantified, individual risk aversion is more subjective.

Investor's expected return: This is calculated from the vector of expected returns across all asset classes, weighted by the investor's portfolio allocation.

Investor's risk aversion: Risk aversion refers to the tendency to prefer less risky outcomes over riskier ones, even when the latter offer potentially higher returns. In portfolio theory, risk aversion is represented by a fraction or multiple ( $\lambda$ ) of portfolio variance.

Mathematical representation:

Let  $C$  denote the covariance matrix of the asset classes in the investment universe,  $r$  the corresponding vector of expected returns, and  $x$  the vector of portfolio weights of our  $\lambda$ -investor. Then:

$$\text{Investor's expected return} = r'x$$

$$\text{Investor's risk aversion} = \lambda * x' C x$$

$$\text{Utility function} = r'x - \lambda * x' C x$$

<sup>1</sup>  $C_{(i,j)} = \rho_{(i,j)} * \sigma_i * \sigma_j$

<sup>2</sup> Always in relative terms, i.e., compared with a risk-free or low-risk asset class.

## 2 The Optimal Portfolio

Consider an investment universe with  $n$  assets and a  $\lambda$ -investor who maximizes the utility function:

$U(x) = r^T x - \lambda x^T C x$  subject to the following constraints:

$1^T x = 1$ : Budget constraint

$A^T x = b$ : Linear investment constraint<sup>3</sup>

Here,  $x \in \mathbb{R}^n$  denotes portfolio weights;  $r \in \mathbb{R}^n$  the investor's expected returns;  $\lambda > 0$  risk aversion;  $C \in \mathbb{R}^{n \times n}$  the covariance matrix;  $A \in \mathbb{R}^{n \times m}$  the matrix of linear constraints; and  $b \in \mathbb{R}^m$  the vector of constraint targets.

### 2.1 Theorem

The investor's optimal portfolio  $x$  subject to  $1^T x = 1$  and  $A^T x = b$  is given by

$$x = \frac{1}{2\lambda} C^{-1} (r + \gamma 1 + A\eta)$$

where the multipliers  $\gamma$  and  $\eta$  are determined by the linear Lagrange system:

$$\begin{pmatrix} 1^T C^{-1} 1 & 1^T C^{-1} A \\ A^T C^{-1} 1 & A^T C^{-1} A \end{pmatrix} \begin{pmatrix} \gamma \\ \eta \end{pmatrix} = \begin{pmatrix} 2\lambda - 1^T C^{-1} r \\ 2\lambda b - A^T C^{-1} r \end{pmatrix}$$

### 2.2 Proof

To solve the optimization problem, we formulate the Lagrangian:<sup>4</sup>

$$L(x, \gamma, \eta) = r^T x - \lambda x^T C x + \gamma(1^T x - 1) + \eta^T(A^T x - b)$$

The first-order optimality condition follows by taking the gradient with respect to  $x$ :

$$\nabla_x L = r - 2\lambda Cx + \gamma 1 + A\eta = 0$$

It follows that:  $2\lambda Cx = r + \gamma 1 + A\eta$ . Hence:  $x = \frac{1}{2\lambda} C^{-1} (r + \gamma 1 + A\eta)$

Substituting the expression for  $x$  into the constraints (budget constraint  $1^T x = 1$  and investment constraint  $A^T x = b$ ) yields two equations for the Lagrange multipliers  $\gamma$  and  $\eta$ .

$$\frac{1}{2\lambda} 1^T C^{-1} (r + \gamma 1 + A\eta) = 1$$

$$\frac{1}{2\lambda} A^T C^{-1} (r + \gamma 1 + A\eta) = b$$

A suitable algebraic rearrangement of these two equations yields precisely the linear Lagrange system.

<sup>3</sup>Constraints of the form  $B^T x \leq u$  can be rewritten as  $B^T x + e = u$ , or  $(B, I)^T(x, e) = u$ , and the optimal vector  $(x, e)$  can then be determined.

<sup>4</sup>The Lagrangian transforms a constrained optimization problem into an unconstrained one.

### 2.3 Corollary

With  $M = [1^T; A]$ ,  $d = [1; b]$ , and  $\xi = (\gamma, \eta)$ , we obtain the compact formula for the optimal portfolio:

$$x^* = \frac{1}{2\lambda} C^{-1} (r + M^T \xi^*)$$

where  $M^T x = d$  and  $\xi^* = (M^T C^{-1} M)^{-1} (2\lambda d - M^T C^{-1} r)$

Geometrically, the optimal portfolio  $x^*$  can be viewed as the projection of the unconstrained optimal portfolio  $C^{-1}r$  onto the feasible space  $M^T x = d$ . In other words, the constraints redirect the unconstrained solution into the feasible space.

### 2.4 Properties

The optimal portfolio  $x^*$  represents a linear transformation of returns  $(r + M^T x^*)$  combined with an inversion of the risk structure ( $C^{-1}$ ). Accordingly, even small changes in  $r$  and  $C$ , as well as unstable covariance matrices, can lead to substantial portfolio changes. In particular, inversion of the covariance matrix can cause numerical instability and considerable portfolio turnover, as it depends strongly on both correlations and volatilities. The equation for  $x^*$  also makes clear that the optimal portfolio structure depends on the constraint matrix  $A$  and constraint vector  $b$ . These affect the structure of the Lagrange solution  $x^*$  through the Lagrange multipliers  $\gamma$  and  $\eta$ . Small changes in  $A$  and  $b$  can shift the entire portfolio and create discontinuities in the solution. Because the optimal portfolio is proportional to  $C^{-1} (r + M^T x^*)$ , mean-variance optimization exploits linear dependencies in the investment universe very aggressively. When assets are highly correlated, the covariance matrix becomes nearly singular, potentially producing unstable portfolio weights.<sup>5</sup>

### 2.5 Instability of the Covariance Matrix

**Decomposition of the covariance matrix:** The algebraic decomposition of a covariance matrix  $C$  into volatilities and correlations is expressed by  $C = D \cdot \rho \cdot D$ , where  $D$  is the diagonal volatility matrix and  $\rho$  is the correlation matrix.

**Inversion of the covariance matrix:** By the matrix inversion rule  $(ABC)^{-1} = C^{-1}B^{-1}A^{-1}$ , we obtain  $C^{-1} = D^{-1}\rho^{-1}D^{-1}$ .

The inverse separates into two components:

- (i)  $D^{-1}$ : a stable diagonal scaling component<sup>6</sup>
- (ii)  $\rho^{-1}$ : a structural and potentially explosive component<sup>7</sup>

**Source of instability:** Numerical stability is determined almost entirely by  $\rho^{-1}$ . At very high correlations (close to 1), the correlation matrix  $\rho$  becomes nearly singular. This leads to very small eigenvalues and a high condition number. Consequently, inversion of  $\rho$  greatly amplifies small perturbations in the input data, and this amplification is directly transmitted to  $C^{-1}$ .

<sup>5</sup>When the covariance matrix is poorly conditioned.

<sup>6</sup> $D$  contains no interactions between assets, and its condition number  $K(D) = \max \text{volatility} / \min \text{volatility}$  is generally not extremely large.

<sup>7</sup>High correlations lead to very small eigenvalues; consequently,  $K(\rho) = \text{maximum eigenvalue} / \text{minimum eigenvalue}$  may become extremely large.

## 3 7-Step Best Practice Code

### 3.1 Perform Empirical Optimization

An initial analysis should be based on empirical data to understand the fundamental structure of the return distribution. The covariance matrix describes the risk structure, and decomposing it into volatilities and correlations allows these components to be analyzed and interpreted separately. In this context, examining multicollinearity and the stability of correlations and volatilities is essential. A useful first exercise is to perform portfolio optimization using empirical data, because the two quantities  $\mu$  (long-term returns) and  $C$  (long-term covariance matrix) belong together and originate from the same economic distribution.

### 3.2 Formulate Economically Grounded Return Views

An investor may formulate return expectations that differ from long-term expectations. These views should be economically grounded, meaningful, and, above all, mutually consistent. We denote these expectations by  $r$ , where  $r \neq \mu$ . In classical mean-variance optimization, it is often assumed that only the return vector is adjusted, while the covariance matrix remains unchanged. However, if  $r \neq \mu$ , the investor implicitly describes a different distribution of future returns. This implies potential changes not only in the first moment (the mean) but also in higher moments and thus in the resulting correlations with other asset classes in the portfolio.

### 3.3 Adjust the Covariance Matrix

Assuming an empirical normal distribution,  $R \sim N(\mu, C)$ , the mean  $\mu$  lies at the center of the distribution. If an investor instead believes that  $R \sim N(r, C)$ , the entire distribution shifts. Such a shift may also affect observed volatility and covariance structure. When investors formulate their own expected returns  $r$ , several arguments may support adjusting the covariance matrix:

1. A shift in expected returns changes the location of the entire distribution.
2. Scenarios with higher expected returns may be associated with higher volatilities.
3. Market scenarios underlying return views may also alter the risk structure.

In practice, correlations are often assumed to remain relatively stable while volatilities can be adjusted. This preserves the fundamental dependence structure among assets while adapting the risk intensity of individual assets to the investor's views. The most consistent approach, however, is to adjust both volatilities and correlations to reflect the investor's return views, especially in tactical asset allocation, where return distributions frequently undergo regime changes.<sup>8</sup>

---

<sup>8</sup>This assumption should nevertheless be tested empirically and statistically.

### 3.4 Assess and, if Necessary, Stabilize the Covariance Matrix

The covariance matrix is central to optimization stability. Because the solution depends on  $C^{-1}$ , poorly conditioned matrices can produce unstable results. The condition number  $\kappa(C)$  is a key stability measure and should be assessed and, if necessary, improved through shrinkage or regularization. If  $\kappa(C)$  is large, a small relative error in  $C$  may result in a substantially amplified error in  $C^{-1}$ .<sup>9,10</sup>

A large condition number usually indicates very small eigenvalues, strong linear dependencies, high correlations, or severe multicollinearity. The following indicative rules of thumb can be used to assess conditioning:  $\kappa(C)$

- $\kappa(C) < 10$ : very good and unproblematic; the matrix is numerically very stable
- $10 \leq \kappa(C) < 100$ : good and acceptable; the matrix is numerically stable
- $100 \leq \kappa(C) < 1000$ : the matrix is already noticeably sensitive; plausibility check recommended
- $1000 \leq \kappa(C) < 10000$ : regularization recommended
- $\kappa(C) \geq 10000$ : regularization necessary

Regularization stabilizes the solution and reduces the amplification of estimation errors. It is an integral component of modern portfolio optimization. Regularization makes inversion more stable and limits extreme portfolio weights.

#### Basic form of regularization:

Effective covariance matrix regularization typically combines a ridge term with a shrinkage term:  $C_{\text{ridge}} = C + \delta I$   $C_{\text{shrink}} C_{\text{target}}$

$$C_{\text{effektiv}} = (1 - \alpha)(C + \delta I) + \alpha C_{\text{target}}$$

The parameter  $\delta$  primarily serves numerical stabilization by increasing the smallest eigenvalues of the covariance matrix. The parameter  $\alpha$  governs the trade-off between empirical information and structural model assumptions. The target matrix represents a robust, economically motivated approximation of the risk structure. The choice of these three components largely determines the stability and plausibility of the resulting portfolio allocation.  $C_{\text{target}}$

#### Choosing $\delta$ (ridge term):

The parameter  $\delta$  raises the small eigenvalues of the covariance matrix and improves numerical stability of inversion.  $\lambda_i() = \lambda_i(C) + \delta$ . A typical choice is  $\delta \approx \epsilon \cdot \lambda_{\max}(C)$ , with  $\epsilon \in [10^{-4}, 10^{-2}]$ .  $C_{\text{ridge}}$

#### Choosing $\alpha$ (shrinkage intensity):

The parameter  $\alpha$  controls the trade-off between the empirical covariance matrix and the target matrix and structurally reduces extreme correlations. Typical values for  $\alpha$  range from 0.1 to 0.3 (with large datasets) and from 0.3 to 0.7 (with small datasets).

<sup>9</sup>  $\kappa(C) = \|C\| \cdot \|C^{-1}\|$ .

<sup>10</sup> Shrinkage

### Choosing the target matrix: $C_{\text{target}}$

In the shrinkage approach, the target matrix acts as a structural reference model and incorporates economic knowledge, whereas  $C$  represents the raw empirical data. In this sense, shrinkage combines data and model. Typical target matrices include:  $C_{\text{shrunk}} \cdot C_{\text{target}}$

1. Diagonal target:  $\text{diag}(C)$  is highly stable, perfectly conditioned, and free of multicollinearity.  $C_{\text{target}}$
2. Constant-correlation target:  $DpD$  is robust, substantially reduces noise, and retains diversification.  $C_{\text{target}}$
3. Factor-model target:  $BFB^T + D$  is stable and, above all, economically interpretable.  $B$  contains the factor loadings,  $F$  is the covariance matrix of common factors, and  $D$  represents idiosyncratic risk.<sup>11</sup>  $C_{\text{target}}$

## 3.5 Choose Constraints Deliberately

Investment constraints do more than impose side conditions: they modify the effective return vector through  $M^T \xi^*$  and actively shape the solution. They define the feasible space and significantly influence the economic interpretation of the portfolio. Pragmatic and practice-oriented selection and formulation of investment constraints can therefore make a major contribution to stabilizing optimization results.

## 3.6 Conduct Sensitivity and Robustness Analyses

The stability of the optimal portfolio with respect to small changes in expected returns ( $\Delta r$ ) and changes in the covariance matrix ( $\Delta C$ ) must be assessed. Unstable solutions are not reliable in practice.<sup>17</sup>

## 3.7 Interpret and Validate the Portfolio

A mathematically optimal portfolio is not necessarily economically meaningful. Portfolio weights, concentration, and implementability must be critically evaluated. It is not merely advisable but essential to check the plausibility of the resulting portfolio and relate it to the underlying assumptions before implementation.

<sup>11</sup>For a comprehensive discussion, see Risk and Asset Allocation (Meucci, 2005, Springer Finance Series).

## 4 References

- Markowitz (1952) Portfolio Selection
- Markowitz (1959) Portfolio Selection Efficient Diversification of Investments
- Ledoit & Wolf (2004) Honey I Shrunk the Sample Covariance Matrix
- Fama & French (1993) Common Risk Factors in the Returns on Stocks and Bonds
- Risk and Asset Allocation (Meucci, 2005, Springer Finance Series)