Master of Science in
Engineering and Management

Department of Management and Production Engineering

Portfolio Optimization Using Tabu Search: A Risk–Return–ESG Approach

Supervisors:
Marco Ghirardi
Fabio Salassa

Candidate:
Hamdi Ahmed
Mohamed

Graduation Session March 2026

I express my sincere gratitude to my parents,

Kadra Hassan Elmi and *Ahmed Hassan Moyale*

for their support and sacrifices throughout my studies.

I thank my professors *Marco Ghirardi* and *Fabio Salassa*

for their availability and guidance during my last course and this thesis.

Torino, 17/02/2026

Hamdi Ahmed Mohamed

Abstract

Author: Hamdi Ahmed Mohamed
Title: Portfolio Optimization Using Tabu Search: A Risk–Return–ESG Approach
University: Politecnico di Torino
Programme: Master of Science in Engineering and Management
Major: Finance
Thesis supervisors: Marco Ghirardi; Fabio Salassa

Date: 17/02/2026 Number of pages: 46 + 3 Language: English

This thesis studies ESG-aware portfolio choice as a problem with three goals: increase expected return, lower volatility, and cut ESG risk. The setup stays close to practice: long-only portfolios, full investment, limits on single positions, and a fixed range for the number of holdings. The tests use FTSE MIB stocks. Historical prices provide return and risk estimates, and ESG risk scores represent non-financial exposure.

Tabu Search drives the optimization. It tries many portfolios, uses short-term memory to avoid repeating the same moves, and keeps only the best trade-off portfolios. The non-dominated solutions are shown on the Pareto front. The analysis also comments on diversification and concentration, and notes that changes in index membership can shift the investable set over time.

Keywords: portfolio management; portfolio optimisation; tabu search; multi-objective optimisation; ESG risk; Pareto front; FTSE MIB; diversification

Table of Contents

Chapter 1

Introduction

1.1 Portfolio Management

Portfolio management is the disciplined activity of building, steering, and revising a set of financial positions so the whole set fits an investor's objectives and constraints. The object is the portfolio as a system: correlations, liquidity, drawdown tolerance, and cash-flow timing matter as much as single-asset expected returns. In practice the manager trades across asset classes—equities, sovereign and corporate bonds, commodities, currencies, and pooled vehicles such as mutual funds and ETFs—and keeps an eye on frictions like transaction costs, taxes, bid–ask spreads, and trading capacity.

Different types of investors are *Individuals*, *Defined benefit pension plans*, *Endowments and foundations*, *Banks*, *Insurance companies*, *Investment companies*, and *Sovereign wealth funds*. They differ by their objectives, namely, Time Horizon, Risk Tolerance, Income Needs, and Liquidity Needs.

The portfolio management cycle can be described in five phases. *Security analysis* examines securities and narrows the investable set. *Portfolio analysis* then turns these into a risk summary of volatility, covariance, and tests whether the estimates and constraints fit together without conflicts.

Portfolio selection fixes the allocation rule and chooses weights that respect the stated risk tolerance and objectives. *Portfolio revision* rebalances as weights changes with time, information updates, or costs and limits start to bind. *Portfolio evaluation* closes the loop by judging realised results with attribution, turnover, tracking error, and drawdowns.

1.2 Portfolio Optimization

Portfolio optimization turns those decisions into a formal problem: choose weights that satisfy constraints and optimize one or more criteria. Linear programming (LP) covers many "clean" formulations: minimum-cost rebalancing with linear transaction costs, cash and sector bounds, and simple tracking-error proxies; practitioners often pair LP with linear risk surrogates (e.g., absolute deviation, piecewise-linear CVaR). Mixed-integer linear programming (MILP) enters when the policy has discrete structure: cardinality constraints (hold at most K assets), minimum-lot sizes, buy-in thresholds, or logical rules such as "if you hold asset i then also hold a hedge." For local search, the common workhorses are neighborhood heuristics (swap, add/drop, and weight-shift moves), tabu search, and simulated annealing; they handle messy constraint sets and nonconvex objectives, especially when the manager insists on integer-like features without paying the full MILP cost.

Multi-objective methods are based on finding the right trade-offs giving multiple objective functions. In the case of Portfolio Selection, the classic one opposes return to risk, yet it is common to add other criterias such as ESG scores, liquidity, or tail risk mesures. Stochastic optimization adds an uncertainty factor, estimating probabilities on possible future outcomes to make more accurate investment decisions.

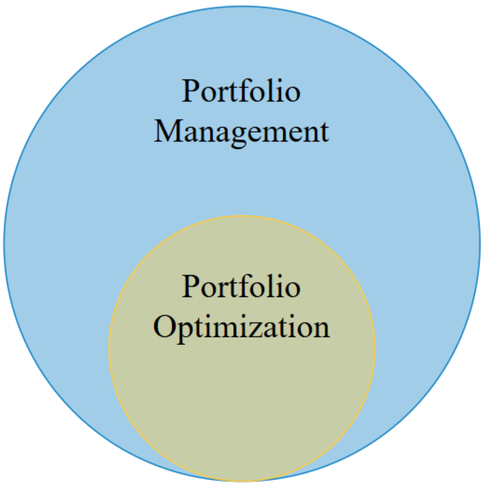

Figure 1.1

Chapter 2

Literature Review

2.1 Portfolio Theory

Return and risk

Return as an expectation

Markowitz [1] treats each security's return as a random variable R_i and defines its *expected return* as the statistical mean. In the discrete exposition, if a random variable Y takes values $y_1, \ldots, y_n$ with probabilities $p_1, \ldots, p_n$, then

$$E(Y) = \sum_{k=1}^{n} p_k \, y_k. \tag{2.1}$$

He deliberately links financial "yield" language to this expectation operator, noting that much of the finance vernacular aligns with "expected yield" even when writers do not state it explicitly.

Risk as a variance

Risk is calculated as the dispersion around the mean, measured by variance:

$$V(Y) = \sum_{k=1}^{n} p_k \, (y_k - E(Y))^2. \tag{2.2}$$

Markowitz also names standard deviation $\sigma = \sqrt{V}$ and the coefficient of variation σ/E as closely related dispersion measures. A useful interpretive point follows:

even if an investor cared about σ or σ/E rather than V, the relevant choice still lies on the efficient set produced by the mean–variance rule.

The mean–variance model

Portfolio mean

Let X_i be the portfolio weight on security i, with $\sum_{i=1}^{N} X_i = 1$ and (in Markowitz's baseline) $X_i \geq 0$ to exclude short sales. Portfolio return is

$$R = \sum_{i=1}^{N} X_i R_i. \tag{2.3}$$

Using linearity of expectation, expected portfolio return becomes

$$E(R) = \sum_{i=1}^{N} X_i \mu_i, \tag{2.4}$$

where $\mu_i = E(R_i)$. This simple additivity makes mean controllable through weights.

Portfolio risk and the covariance

Markowitz defines covariance as

$$\sigma_{ij} = E\Big[(R_i - \mu_i)(R_j - \mu_j)\Big], \tag{2.5}$$

and relates it to correlation ρ_{ij} via $\sigma_{ij} = \rho_{ij}\sigma_i\sigma_j$. The variance of a weighted sum expands into diagonal and cross terms:

$$V(R) = \sum_{i=1}^{N} \sum_{j=1}^{N} X_i X_j \sigma_{ij}. \tag{2.6}$$

Return–risk trade-off

Markowitz rejects the idea that the maximum-mean portfolio must also be the minimum-variance portfolio. Instead, he treats investing as a trade-off between expected return E and variance V: if you want a higher E, you typically accept a higher V, and if you want a lower V, you usually give up some E. He summarizes this logic with the E–V rule: pick portfolios that (i) minimize V for a given level of E (or higher), and (ii) maximize E for a given level of V (or lower). The portfolios that satisfy this rule form the efficient set in the mean–variance plane.

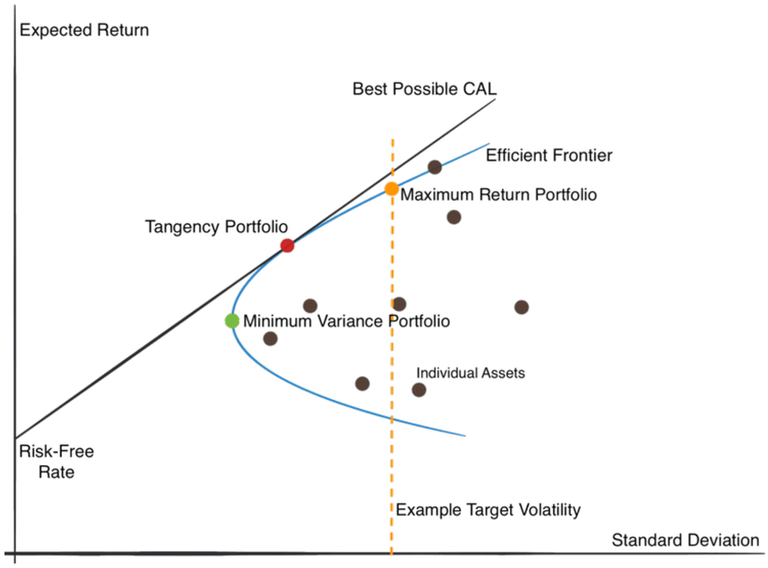

Figure 2.1: (Source: Quantpedia.com)

Figure 2.1 shows the efficient portfolios for a given level of volatility. The Capital Allocation Line is the straight line that shows all mixes of a risk-free asset and one risky portfolio; how expected return rises as you take more risk. The tangency portfolio gives the highest reward per unit of risk or the maximum Sharpe ratio,

$$\text{Sharpe Ratio} = \frac{\mathbb{E}[R_p] - R_f}{\sigma_p}.$$

2.2 ESG rating

Introduction

Abate et al. [2] frame ESG investing as a shift in the objective function of asset allocation. Traditional portfolio choice prioritizes risk-adjusted performance; ESG investing adds a second layer: environmental stewardship, social responsibility, and governance quality become explicit criteria alongside return and risk. The authors treat this as a structural change in the market rather than a temporary theme, driven by investor demand and by the operational availability of ESG ratings from specialized agencies. They also underline a practical tension, screening on ESG can shrink the investable universe, yet high-ESG firms may carry lower exposure to reputational and regulatory shocks.

Figure 2.2: (Source: FTSE Russell)

Definition

ESG stands for environmental, social, and governance factors, and investors use it alongside financial analysis to judge material risk and firm conduct. It helps analysts flag material non-financial risk and gauge how an issuer is governed. ESG works as a screening framework. Figure 2.2 visualises this three-part structure.

Rating systems

Mesure name	**Rating agency**	**Scale**
ESG Ratings	MSCI	CCC to AAA
ESG Risk Ratings	Morningstar Sustainalytics	0 to 100
ESG Scores	LSEG	0 to 100

Table 2.1: ESG rating systems

Table 2.1 summarises three ESG scoring systems. MSCI uses an industry-relative, ordinal ladder (AAA to CCC), which places each issuer against its sector peer set. Sustainalytics reports an unmanaged-risk score on a 0 to 100 scale and ties it to enterprise value at risk. LSEG converts disclosed policies, programmes, and performance signals into a 0 to 100 percentile rank.

2.3 Local and Tabu Search

Local Search

Local search basics

Local search treats an optimization problem as a *walk* through a set of candidate solutions. Schaerf [3] defines a search space S whose elements are *states* $s \in S$, each state encoding a potential solution. A state is *feasible* when it satisfies all constraints. The method also needs a neighborhood function $N(\cdot)$ that assigns to each state s a set of nearby states $N(s) \subseteq S$, built by applying small local changes called *moves*. Starting from an initial state s_0 (constructed or random), the algorithm repeats a loop that moves from the current state s_i to a neighbor $s_{i+1} \in N(s_i)$ according to a chosen selection strategy.

Iterations and stopping

An *iteration* is one pass through the search loop: the method lists the neighborhood, evaluates each candidate with the cost function, selects a move, and updates the incumbent. The notation $s' = s \otimes m$ denotes applying move m to solution s to obtain the successor s'.

The run may terminate at s_{final}, while the best solution encountered is s_{best}; they can diverge when the trajectory drifts into a weaker basin after crossing a better region. The cost function also changes across phases. In a feasibility stage, f measures constraint violation only, so the target is $f = 0$. In an optimization stage, f combines the objective with a violation penalty, which permits minor breaches when they purchase a material improvement.

Neighborhood moves

Neighborhood types fixes what "close" means in thoughout the search. The neighborhood uses a compact menu of moves. One move nudges a single weight up or down, then renormalizes the remaining weights so the budget constraint still holds. Another move inserts or deletes an asset to satisfy the cardinality cap. A third move transfers a fraction of weight from one asset to another, so it disturbs fewer holdings. After any move, a repair step restores feasibility: weights sum to 1 and respect any minimum bounds.

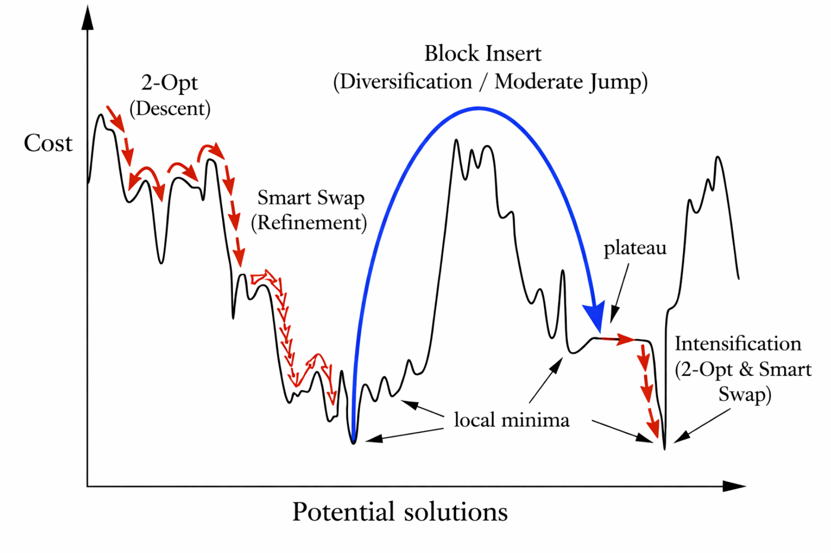

Figure 2.3: Local Search and Neighborhood moves

Tabu Search

Why tabu search

Tabu search keeps the local-search loop but relaxes the move rule. From a state s_i, it scans the neighborhood $N(s_i)$ and selects the neighbor with the lowest cost, even if that move increases f relative to s_i. This lets the walk climb out of local minima instead of stalling at the first non-improving step. The trade-off is cycling: once uphill moves are allowed, the path can loop back to states it has already visited.

Tabu memory and aspiration

To limit cycles, tabu search stores recent moves on a short-term tabu list and forbids immediate reversals. The run typically ends after a fixed number of *idle iterations*, i.e., iterations that do not improve the best solution found so far.

Figure 2.4 shows a local search sliding downhill into a local minimum and even cycling near flat surfaces, while Tabu Search breaks that trap by allowing controlled non-improving moves and memory-based exploration that can reach the global minimum.

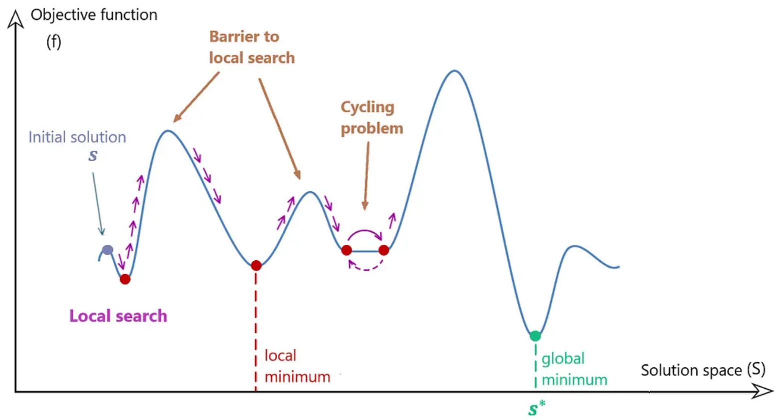

Figure 2.4: Local Search limitations (source: Medium.com)

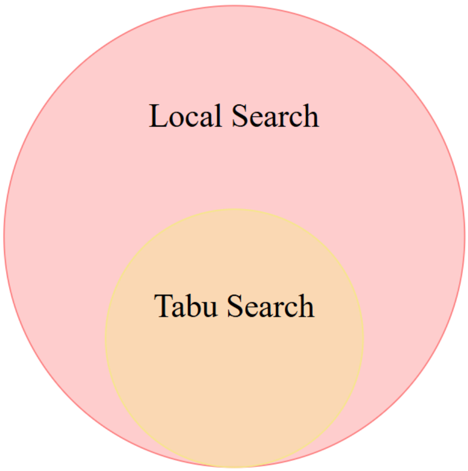

Figure 2.5

2.4 Pareto Front

Objective space

According to Grodzevich and Romanko [4], multi-objective optimization asks us to improve several criteria at the same time, written as min $f_1(x), \ldots, f_k(x)$ subject to $x \in \Omega$. A useful viewpoint is the *objective feasible region* $Z \subset \mathbb{R}^k$, obtained by mapping each feasible decision x to its objective vector $(f_1(x), \ldots, f_k(x))$. Once you work in this space, trade-offs become visible: moving in one direction helps one objective while hurting another.

Pareto optimality

A feasible point $x^\star \in \Omega$ is Pareto optimal if there is no other feasible x such that $f_i(x) \leq f_i(x^\star)$ for all i and $f_j(x) < f_j(x^\star)$ for at least one index j. The set of all Pareto optimal solutions P is the *efficient frontier*. From a mathematical angle, every point on P is acceptable but in practice, the decision maker still has to pick one final compromise.

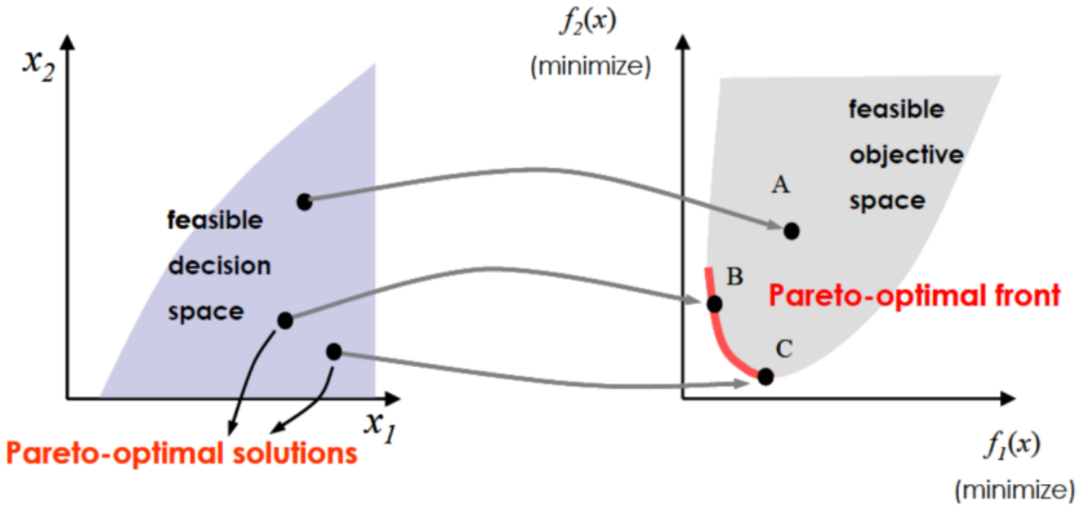

Figure 2.6

In the Figure 2.6, the left panel shows the feasible decision space in (x_1, x_2), and the arrows map feasible choices into the objective space $(f_1(x), f_2(x))$ on the right.

The red curve is the Pareto-optimal front: along this boundary you can't reduce f_1 without increasing f_2, or vice versa, so it marks the efficient trade-off set.

Weighted-sum method

The weighted-sum method turns a k-objective problem into one objective:

$$\min_{x\in\Omega} \sum_{i=1}^{k} w_i\, f_i(x) \qquad \text{with} \qquad w_i \geq 0,\ \sum_{i=1}^{k} w_i = 1.$$

When Ω and the objectives are convex and all weights satisfy $w_i > 0$, any minimizer is Pareto optimal. Under strict convexity, that minimizer is unique. Varying w shifts the selected trade-off along the efficient frontier.

Why normalization

Weights make sense only when objectives share comparable scales. If one f_i is numerically larger, it can dominate the sum even with a small w_i.

A clean split is $w_i = u_i\theta_i$: u_i reflects preference, while θ_i handles units. This scaling can mislead when θ_i comes from a weak reference value or from a single-objective optimum that ignores the Pareto geometry, since it can skew the implied trade-off.

Weighted-sum limits

With linear objectives and a simplex-type solver, the method can behave poorly. If the Pareto set has a flat face, the solver often returns a vertex rather than a point in the interior, even when w changes. As a result, tiny shifts in w may flip the solution between corners, while large shifts may leave the same corner selected.

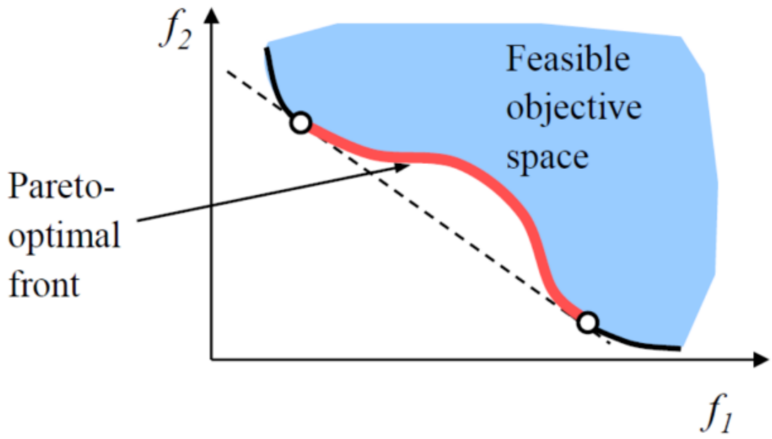

Figure 2.7

Figure 2.7 shows the feasible objective space (blue) in the (f_1, f_2) plane and the Pareto-optimal front (red), where any gain in one objective forces a loss in the other. The dashed weighted-sum line acts as a supporting line, it slides until first contact, so it selects only the contact points (white circles) and ignores the inward-bending part of the front. Because the front bends inward, the weighted sum can miss good trade-offs even when they are Pareto optimal.

Two practical fixes keep the weighted-sum setup but change how the solver behaves. Interior-point methods can return interior efficient points when they exist, rather than sticking to vertices. In some cases, squaring a linear objective adds curvature and produces more spread-out efficient solutions, but the reformulation must still reflect the original decision goal.

2.5 Diversification

Diversification Ratio

Definition

Diversification is a portfolio-level attribute: it describes how combining imperfectly co-moving assets can compress total risk relative to holding exposures in isolation. In that sense, diversification is encoded in the covariance structure of returns, not in any single security's volatility taken alone. Choueifaty and Coignard [5] treat diversification as a measurable criterion for construction, built from second moments.

The diversification ratio shows how much the portfolio's risk drops compared with holding each asset on its own. It's mesured as

$$D(P) = \frac{P^\top \Sigma}{\sqrt{P^\top V P}}.$$

Given an N-asset portfolio. Let $\Sigma = (\sigma_1, \ldots, \sigma_N)^\top$ be the vector of individual asset volatilities, let V be the covariance matrix, and let $P = (w_{p1}, \ldots, w_{pN})^\top$ be the weight vector with $\sum_{i=1}^{N} w_{pi} = 1$.

Interpretation of the Diversification Ratio

This view shifts attention from return forecasts to the dependence structure of the asset set, i.e., covariances and correlations. The diversification ratio captures that idea. Its denominator, $\sqrt{P^\top V P}$, uses the covariance matrix only, while the numerator, $P^\top \Sigma$, sums stand-alone volatilities and needs no expected-return input. Maximising $D(P)$ therefore favours portfolios with low total volatility relative to the average marginal risk of their holdings.

Estimation error

Mean–variance optimisation needs two estimates: expected returns and the covariance matrix. Expected returns are the weak link. They jump around across samples and depend heavily on the chosen window, so the optimiser can mistake noise for a real signal and push weights to extremes. Covariances are usually steadier and easier to pin down with decent precision. This gap explains why the diversification ratio is an attractive objective: it leans on volatilities and covariances, not on expected returns. It focuses on inputs we can estimate more safely, so the portfolio is less likely to follow noisy return guesses and more likely to gain from assets not moving together.

Herfindahl–Hirschman Index

Definition

The Herfindahl–Hirschman Index (HHI) measures concentration from shares. It rises when fewer firms hold most of the market. The index squares each share and sums the results. Rhoades [6] defines

$$HHI = \sum_{i=1}^{n} (MS_i)^2 ,$$

where MS_i is the market share of firm i and n is the number of firms.

Interpretation of the HHI

Squaring makes large shares count more than small ones, so dominant incumbents push the index up. High HHI signals a concentrated market and weaker competitive pressure; low HHI points to dispersed shares and stronger rivalry.

Stocks use

You can repurpose the HHI to measure how concentrated a stock index is by treating each constituent's index weight as a "share".

For an equal-weight index, every share is identical and the HHI gets smaller as its number increases while for a capitalization-weighted index, those weights already represent each firm's slice of total index capitalization. The index becomes

$$HHI = \sum_{i=1}^{n} w_i^2 .$$

Correlation

As Saxo Group [7] defines it, correlation is a statistical coefficient that tracks co-movement between two variables. In portfolio work, it describes how the returns of two assets move relative to each other. The coefficient lies in the closed range $[-1,1]$: values near $+1$ signal strong comovement, values near -1 signal strong opposite movement, and values near 0 mean no clear linear link.

A standard definition uses the *correlation coefficient*:

$$\rho_{X,Y} = \frac{\mathrm{cov}(X,Y)}{\sigma_X\,\sigma_Y},$$

with $\mathrm{cov}(X,Y)$ the covariance and σ_X, σ_Y the standard deviations. Scaling removes units, so the correlations are comparable between asset pairs. As $|\rho|$ moves toward 1, the linear co-movement tightens.

Portfolio work uses correlation to manage variance and harvest diversification. Low or negative correlations reduce risk because shocks do not propagate across holdings in lockstep. A correlation matrix gathers these coefficients and reveals the dependence structure, such as tight clusters and weaker links elsewhere.

Chapter 3

Research Discussion

3.1 Problem formulation

In words, the problem is to construct portfolios from a given equity universe, using return, variance and ESG scores as inputs. The decision variables are the continuous weights w_i and the binary indicators z_i that encode whether an asset enters the portfolio. Full investment and the no short-selling restriction keep the allocation realistic: all wealth is deployed, no leverage is allowed, and each position corresponds to a non negative share of total capital.

Cardinality and quantity constraints gives boundaries to our local search. The portfolio must contain between 15 and 25 different assets, which gives a reasonable amount of stocks given our choosen index, FTSE MIB. Once an asset is selected, its weight must lie between 2.5% and 10% of the portfolio, so every chosen stock carries are significant while no single one dominates. The research explores: a three-objective risk–return–ESG model. This setup allows us to evaluate the best solutions on a three-dimensional Pareto front.

3.2 Mathematical model

The portfolio is described by a vector of decision variables

$$\mathbf{w} = (w_1, \ldots, w_n)^\top,$$

where w_i is the portfolio weight assigned to asset i, and n is the number of assets in the investable universe. Each asset has an expected return, a contribution to portfolio variance, higher-order moments, and an environmental, social and governance (ESG) score.

Decision variables and data

The model uses the following inputs:

- $\mathbf{r} = (r_1, \ldots, r_n)^\top$: vector of expected asset returns.
- $\Sigma \in \mathbb{R}^{n \times n}$: covariance matrix of asset returns.
- $\mathbf{e} = (e_1, \ldots, e_n)^\top$: vector of ESG risk scores, one for each asset.
- $\mathbf{z} = (z_1, \ldots, z_n)^\top \in \{0,1\}^n$: vector of binary selection variables, where $z_i = 1$ if asset i is included in the portfolio and $z_i = 0$ otherwise.

Objective functions

Expected return. The portfolio's expected return is the weighted sum of individual expected returns:

$$f_1(\mathbf{w}) = R_P(\mathbf{w}, \mathbf{r}) = \mathbf{w}^\top \mathbf{r}, \tag{3.1}$$

which we seek to maximise.

Variance (risk). Portfolio risk is quantified through variance:

$$f_2(\mathbf{w}) = \Sigma_P(\mathbf{w}, \mathbf{r}) = \mathbf{w}^\top \Sigma \mathbf{w}, \tag{3.2}$$

which is minimised.

ESG score. The ESG profile of the portfolio is modelled as a weighted average of asset-level scores:

$$f_3(\mathbf{w}) = ESG_P(\mathbf{w}, \mathbf{e}) = \mathbf{w}^\top \mathbf{e}, \tag{3.3}$$

which is minimised to favour allocations with the least ESG risk.

Constraints

Cardinality constraint. The portfolio size is constrained through a two-sided cardinality bound:

$$15 \;\le\; \sum_{i=1}^{n} z_i \;\le\; 25. \tag{3.4}$$

This inequality forces any feasible portfolio to hold at least fifteen and at most twenty-five distinct assets, without imposing any sector restrictions.

Quantity constraint. To avoid vanishingly small positions and highly concentrated bets, the continuous weights w_i are linked to the binary variables via a quantity constraint. In a typical 20-stock portfolio an equal-weight allocation would assign $1/20 = 0.05$ of the wealth to each asset. I allow some dispersion around this reference level and impose

$$0.025\, z_i \;\le\; w_i \;\le\; 0.10\, z_i, \qquad i = 1, \ldots, n. \tag{3.5}$$

If $z_i = 0$, (3.5) yields $w_i = 0$, so asset i is excluded. If $z_i = 1$, the position in asset i must lie between 2.5% and 10% of the total portfolio value. These bounds are compatible with the full investment constraint

$$\sum_{i=1}^{n} w_i = 1$$

and the no–short–selling requirement $w_i \ge 0$ for all i, which remain as in the base formulation.

Full investment. The full wealth of the investor is allocated across the n assets:

$$\sum_{i=1}^{n} w_i = 1. \tag{3.6}$$

No short selling. Short positions are excluded, so all weights are non-negative:

$$w_i \geq 0, \quad i = 1, \dots, n. \tag{3.7}$$

Model

The final optimization problem combines these three objectives with the constraints. In compact form the model is

$$\begin{aligned}
\max_{w} \quad & f_1(w), \\
\min_{w} \quad & f_2(w),\ f_3(w), \\
\text{s.t.} \quad & \sum_{i=1}^{n} w_i = 1, \\
& w_i \geq 0, \qquad i = 1, \dots, n, \\
& 15 \leq \sum_{i=1}^{n} z_i \leq 25, \\
& 0.025\, z_i \leq w_i \leq 0.10\, z_i, \qquad i = 1, \dots, n, \\
& z_i \in \{0,1\}, \qquad i = 1, \dots, n.
\end{aligned}$$

3.3 Stock index

The **FTSE MIB** is the flagship equity index of the Italian market. It tracks 40 large, liquid, free-float-adjusted companies listed on Euronext Milan, so it works as the closest thing to a "thermometer" for Italy's blue-chip segment. In practice, it also acts as a reference benchmark for asset managers who run Italy-focused mandates, and it supports a full ecosystem of indexed products: ETFs, index funds, and derivatives linked to the same basket of constituents.

The index gives quick access to Italy's largest listed firms. It usually trades with tighter spreads than many single names, which helps keep transaction costs down. Holding the index also mixes multinationals with domestic issuers, so it approximates the Italian equity market without assembling each stock.

Its main drawback is concentration risk. A few sectors often dominate the weights, especially financials and utilities. That tilt raises sensitivity to macro shocks such as rate moves, shifts in the Italian sovereign spread, policy changes, and commodity-price swings.

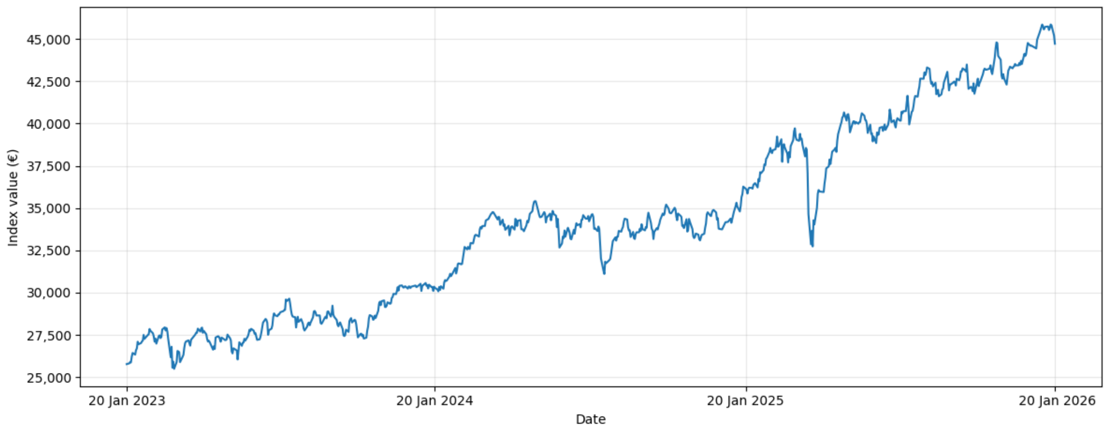

Figure 3.1: FTSE MIB Stock Index Value (Souce: Yahoo Finance)

Figure 3.1 plots the FTSE MIB index value over 20 January 2023 to 20 January 2026, which is the thesis sample window. The series tends upward overall, with a few sharp draws along the path.

Table 3.1: FTSE MIB constituents (40 stocks)

Ticker	Company Name	Max	Min	Adj Close	Std Dev	ESG Score
A2A.MI	A2A S.p.A.	2.72	1.15	2.45	0.35	20.1
AMP.MI	Amplifon S.p.A.	34.88	12.80	13.43	6.28	14.5
G.MI	Assicurazioni Generali S.p.A.	36.13	14.70	34.23	6.61	14.1
AZM.MI	Azimut Holding S.p.A.	37.16	16.01	36.20	5.13	20.4
BPE.MI	BPER Banca SpA	12.20	1.82	11.75	2.72	14.9
BMED.MI	Banca Mediolanum S.p.A.	19.92	6.28	19.34	3.76	24.0
BMPS.MI	Banca Monte dei Paschi di Siena S.p.A.	9.35	1.54	8.86	2.24	18.0
BPSO.MI	Banca Popolare di Sondrio S.p.A	17.57	3.05	16.99	3.77	25.6
BAMI.MI	Banco BPM S.p.A.	13.20	2.51	12.51	3.20	16.0
BC.MI	Brunello Cucinelli S.p.A.	130.60	67.63	82.16	12.96	21.4
BZU.MI	Buzzi S.p.A.	54.50	18.63	48.36	10.07	27.1
CPR.MI	Davide Campari-Milano N.V.	12.69	5.08	5.99	2.36	19.9
DIA.MI	DiaSorin S.p.A.	121.01	58.36	74.40	11.08	19.4
ENEL.MI	Enel SpA	9.18	4.32	8.94	1.20	17.4
ENI.MI	Eni S.p.A.	16.72	9.92	16.38	1.32	25.3
RACE.MI	Ferrari N.V.	484.13	217.91	289.70	64.91	21.3
FBK.MI	FinecoBank Banca Fineco S.p.A.	22.84	9.80	22.25	3.20	11.4
HER.MI	Hera S.p.A.	4.25	2.11	4.21	0.59	14.7
INW.MI	Infrastrutture Wireless Italiane S.p.A.	10.82	7.36	7.36	0.71	15.3
IP.MI	Interpump Group S.p.A.	53.09	25.93	46.28	5.07	27.5
ISP.MI	Intesa Sanpaolo S.p.A.	6.04	1.70	5.76	1.27	9.5
IG.MI	Italgas S.p.A.	10.54	4.19	10.05	1.51	15.2
IVG.MI	Iveco Group N.V.	19.25	6.69	18.85	4.04	22.3
LDO.MI	Leonardo S.p.a.	60.00	8.59	59.56	15.21	22.3
MB.MI	Mediobanca Banca di Credito Finanziario S.p.A.	21.26	7.18	17.05	3.72	12.6
MONC.MI	Moncler S.p.A.	67.55	45.28	49.57	5.20	8.5
NEXI.MI	Nexi S.p.A.	8.30	3.63	3.67	1.01	23.8
PIRC.MI	Pirelli & C. S.p.A.	6.42	3.77	6.22	0.71	7.2

Continued on next page

Ticker	Company Name	Max	Min	Adj Close	Std Dev	ESG Score
PST.MI	Poste Italiane S.p.A.	22.39	7.39	21.99	4.41	13.4
PRY.MI	Prysmian S.p.A.	95.80	32.72	94.12	16.86	18.0
REC.MI	Recordati Industria Chimica e Farmaceutica S.p.A.	59.41	35.22	47.48	5.12	25.0
STMMI.MI	STMicroelectronics N.V.	49.02	16.49	23.47	9.04	11.1
SPM.MI	Saipem SpA	2.76	1.05	2.76	0.45	20.8
SRG.MI	Snam S.p.A.	5.79	3.72	5.73	0.52	16.6
STLAM.MI	Stellantis N.V.	23.33	7.04	8.08	3.89	21.4
TIT.MI	Telecom Italia S.p.A.	0.58	0.21	0.56	0.09	19.9
TEN.MI	Tenaris S.A.	18.31	10.48	18.09	1.87	21.6
TRN.MI	Terna S.p.A.	9.40	6.10	9.09	0.78	9.6
UCG.MI	UniCredit S.p.A.	72.63	12.49	71.22	17.13	10.6
UNI.MI	Unipol Assicurazioni S.p.A.	20.78	3.76	19.88	5.38	18.4

Table 3.1 lists the ticker universe fed into the tabu search. Itreports four summaries: $\max_{t \in T} P_{i,t}$ (Max), $\min_{t \in T} P_{i,t}$ (Min), the last observed adjusted close in the window (Last Adj. Close), and $\mathrm{sd}(P_{i,t})_{t \in T}$ (Std. Dev.). For each stock i, let $P_{i,t}$ be the adjusted close on day t. Let T be the set of trading days shared by all 40 stocks from 20 Jan 2023 to 20 Jan 2026.

The ESG risk score spans from 7.2 to 27.5, with a mean near 17.9. That spread is large enough to create real allocation trade-offs once ESG becomes a third objective alongside risk and return. The table works as a compact input sheet for screening, scalarisation, or Pareto-based selection.

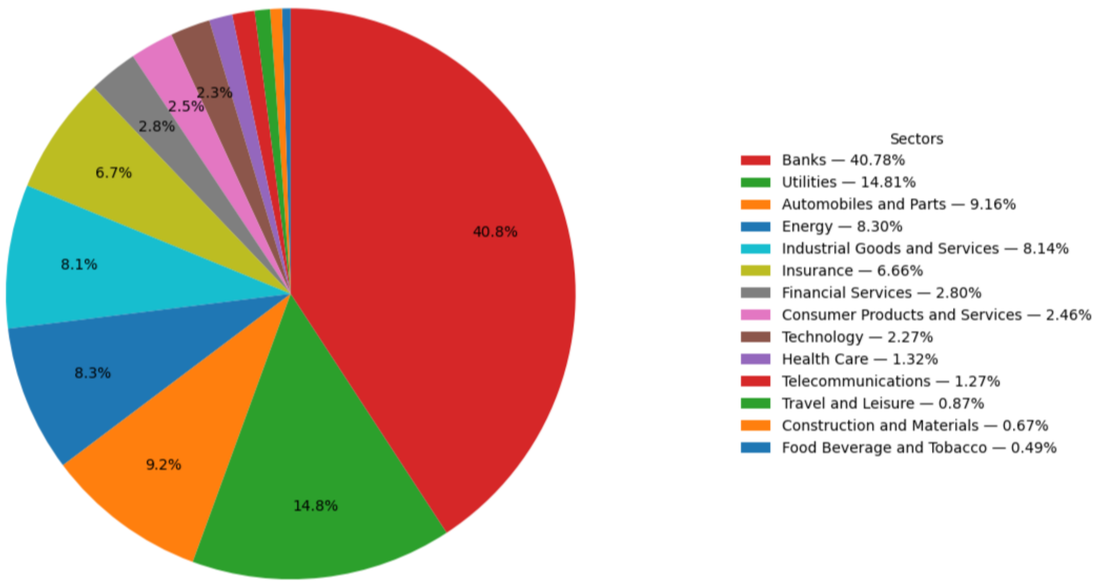

Figure 3.2: FTSE MIB sectors (Source: FTSE Russell as at 31 December 2025)

Figure 3.2 shows strong concentration in the index. Banks take about 40.8% of the index, while Utilities add about 14.8% and Automobiles and Parts about 9.2%. Most other sectors can be considered to have limited influence on index moves.

Chapter 4

Methodology

4.1 Introduction

This chapter opens the black box behind the Tabu Search procedure by pairing two views of the same method: the pseudocode, which pins down the steps and bookkeeping, and the flow chart, which makes the control logic easy to track at a glance. The pseudocode states how the algorithm turns multiple portfolio goals into a single score at each run through a weighted sum, keeps the inputs on a comparable scale, and records non-dominated portfolios in an archive as the search progresses. It also spells out the feasibility rules that matter in practice, full investment, long-only weights, selection limits, and weight bounds that apply only when an asset is actually chosen. The flow chart then follows one search cycle in motion: generate neighbours, screen moves against the tabu list, allow an aspirational override when a forbidden move improves the best-so-far solution, update the tabu tenure and the incumbent portfolio, and stop when the iteration cap or a stall counter signals that the trajectory has gone flat.

4.2 Algorithm

Pseudocode 1: Weighted-sum local search for ESG-aware portfolio optimization

Require: $r \in \mathbb{R}^n$ (expected returns), $\Sigma \in \mathbb{R}^{n\times n}$ (covariance), $e \in \mathbb{R}^n$ (ESG risk)
Require: $N_{\min} = 15$, $N_{\max} = 25$, $w_{\min} = 0.025$, $w_{\max} = 0.10$; grid step $\Delta = 0.05$
Require: Max iterations $I_{\max}$; tabu tenure T; stall limit I_{stall}
Ensure: Approximate Pareto archive $\mathcal{P}$ of feasible portfolios (w, z) and objective triples (f_1, f_2, f_3)

Objectives:
$f_1(w) \leftarrow w^\top r$ ▷ maximize
$f_2(w) \leftarrow w^\top \Sigma w$ ▷ minimize
$f_3(w) \leftarrow w^\top e$ ▷ minimize

Constraints:
$\sum_{i=1}^n w_i = 1, \quad w_i \geq 0, \quad z_i \in \{0,1\}$
$N_{\min} \leq \sum_{i=1}^n z_i \leq N_{\max}$
$w_{\min} z_i \leq w_i \leq w_{\max} z_i$ for all i

Step 1: scale factors from extreme runs
$(w^{(1)}, z^{(1)}) \leftarrow \textsc{SolveExtreme}(\max f_1(w)$ s.t. constraints$)$
$(w^{(2)}, z^{(2)}) \leftarrow \textsc{SolveExtreme}(\min f_2(w)$ s.t. constraints$)$
$(w^{(3)}, z^{(3)}) \leftarrow \textsc{SolveExtreme}(\min f_3(w)$ s.t. constraints$)$
$OF_1 \leftarrow (w^{(1)})^\top r$
$OF_2 \leftarrow (w^{(2)})^\top \Sigma w^{(2)}$
$OF_3 \leftarrow (w^{(3)})^\top e$
$\varepsilon \leftarrow 10^{-12}$

Step 2: weight grid and Pareto archive
$V \leftarrow \textsc{WeightGrid}(\Delta)$ ▷ $V = \{\, v \in \{0, \Delta, 2\Delta, \ldots, 1\}^3 \; : \; v_1 + v_2 + v_3 = 1 \,\}$
$\mathcal{P} \leftarrow \emptyset$
for each $v = (v_1, v_2, v_3) \in \mathcal{V}$ **do**
    $\alpha_1 \leftarrow 100 \cdot \dfrac{v_1}{|OF_1| + \varepsilon}$
    $\alpha_2 \leftarrow 100 \cdot \dfrac{v_2}{|OF_2| + \varepsilon}$

$\alpha_3 \leftarrow 100 \cdot \dfrac{v_3}{|OF_3| + \varepsilon}$
$(w^\star, z^\star) \leftarrow \textsc{TabuSearchWS}(\alpha_1, \alpha_2, \alpha_3, r, \Sigma, e, N_{\min}, N_{\max}, w_{\min}, w_{\max}, I_{\max}, T, I_{\text{stall}})$
$f_1^\star \leftarrow (w^\star)^\top r$
$f_2^\star \leftarrow (w^\star)^\top \Sigma w^\star$
$f_3^\star \leftarrow (w^\star)^\top e$
$\mathcal{P} \leftarrow \textsc{UpdatePareto}(\mathcal{P}, (w^\star, z^\star, f_1^\star, f_2^\star, f_3^\star))$
end for
return $\mathcal{P}$

procedure $\textsc{TabuSearchWS}(\alpha_1, \alpha_2, \alpha_3, r, \Sigma, e, N_{\min}, N_{\max}, w_{\min}, w_{\max}, I_{\max}, T, I_{\text{stall}})$
$F(w) \leftarrow -\alpha_1\, w^\top r + \alpha_2\, w^\top \Sigma w + \alpha_3\, w^\top e$ ▷ minimize
Initial solution (20 assets, equal weights)
$N_0 \leftarrow 20$
choose $S \subset \{1, \ldots, n\}$ with $|S| = N_0$
for $i = 1$ to n **do**
$z_i \leftarrow 1$ if $i \in S$, else 0
$w_i \leftarrow \frac{1}{N_0}$ if $i \in S$, else 0
end for
$(w, z) \leftarrow \textsc{Repair}(w, z, N_{\min}, N_{\max}, w_{\min}, w_{\max})$
$(w_{\text{best}}, z_{\text{best}}) \leftarrow (w, z)$
$F_{\text{best}} \leftarrow F(w)$
TabuList $\leftarrow \emptyset$
stall $\leftarrow 0$
for *iter* $= 1$ to $I_{\max}$ **do**

Neighborhood generation (three move types)
$\mathcal{N} \leftarrow \emptyset$
Swap move:
pick $i \notin S$, $j \in S$, set $S' = (S \cup \{i\}) \setminus \{j\}$
Add/Remove move:
if $|S| < N_{\max}$ add one asset to S; if $|S| > N_{\min}$ remove one asset from S
Shift-weight move:
pick $i, j \in S$, choose feasible δ, set $w_i \leftarrow w_i + \delta$, $w_j \leftarrow w_j - \delta$
$(w', z', move) \leftarrow \textsc{BestAdmissible}(\mathcal{N}, \text{TabuList}, F_{\text{best}}, F(\cdot))$
$(w', z') \leftarrow \textsc{Repair}(w', z', N_{\min}, N_{\max}, w_{\min}, w_{\max})$
$(w, z) \leftarrow (w', z')$
if $F(w) < F_{\text{best}}$ **then**
$(w_{\text{best}}, z_{\text{best}}) \leftarrow (w, z)$
$F_{\text{best}} \leftarrow F(w)$
stall $\leftarrow 0$ ▷ aspiration satisfied

else
$stall \leftarrow stall + 1$
end if
UPDATETABU(TabuList, move, T)
if $stall \geq I_{\text{stall}}$ **then**
break
end if
end for
return $(w_{\text{best}}, z_{\text{best}})$
end procedure

procedure REPAIR($w, z, N_{\min}, N_{\max}, w_{\min}, w_{\max}$)
$S \leftarrow \{i : z_i = 1\}$
$N \leftarrow |S|$
if $N < N_{\min}$ **then**
add arbitrary assets to S until $|S| = N_{\min}$
else if $N > N_{\max}$ **then**
remove arbitrary assets from S until $|S| = N_{\max}$
end if
$S \leftarrow \{i : i \in S\}$ ▷ refresh set
for $i = 1$ to n **do**
if $i \in S$ **then**
$z_i \leftarrow 1$
else
$z_i \leftarrow 0$
$w_i \leftarrow 0$
end if
end for
Feasible weight assignment
$N \leftarrow |S|$
for each $i \in S$ **do**
$w_i \leftarrow \frac{1}{N}$ ▷ for $N \in$ [15,25], this lies in [$0.04, 0.066\overline{6}$]
end for
return (w, z)
end procedure

Interpretation

In the above pseudocode, *Require* states the contract at the entry point: the data vectors and matrix (r, Σ, e) that define return, risk, and ESG-risk, plus the feasibility knobs that shape the investable set (cardinality bounds $N_{\min}, N_{\max}$, box bounds $w_{\min}, w_{\max}$, grid step Δ, and the search budgets $I_{\max}$, tabu tenure T, and stall limit I_{stall}). *Ensure* fixes the expected deliverable: an approximate Pareto archive P that stores feasible portfolios (w, z) together with their objective triples (f_1, f_2, f_3), not a single "best" point. The *Objectives* are explicit and asymmetric: $f_1(w) = w^\top r$ is maximized (reward), while $f_2(w) = w^\top \Sigma w$ and $f_3(w) = w^\top e$ are minimized (variance and ESG exposure). The *Constraints* couple continuous weights and binary selections: full investment $\sum_i w_i = 1$, long-only $w_i \geq 0$, inclusion flags $z_i \in \{0,1\}$, a cardinality band $N_{\min} \leq \sum_i z_i \leq N_{\max}$, and gated bounds $w_{\min} z_i \leq w_i \leq w_{\max} z_i$ so excluded assets carry zero weight and included assets stay within position limits.

Step 1 solves three single-objective problems (max return, min variance, min ESG risk). It takes the reference values (OF_1, OF_2, OF_3) to scale the weights, and uses a small ε to avoid division by zero.

Step 2 samples preferences on a fixed grid. It builds a simplex lattice V of triples (v_1, v_2, v_3) with entries in $\{0, \Delta, 2\Delta, \ldots, 1\}$ and the constraint $v_1 + v_2 + v_3 = 1$. It then loops over each $v \in V$, sets

$$\alpha_k \propto \frac{v_k}{|OF_k| + \varepsilon},$$

and may scale all α_k by a common constant (e.g., 100) for numerical convenience. For each grid point, tabu search minimises the weighted-sum scalarisation

$$F(w) = -\alpha_1 w^\top r + \alpha_2 w^\top \Sigma w + \alpha_3 w^\top e.$$

The search explores three neighborhood families with distinct economic meaning. The *swap move* alters composition while preserving cardinality: it ejects one selected asset $j \in S$ and admits one unselected asset $i \notin S$, so the portfolio can rotate across names without drifting outside $[N_{\min}, N_{\max}]$. The *add/remove move* changes breadth: if $|S| < N_{\max}$ it adds an asset (diversification), and if $|S| > N_{\min}$ it removes one (concentration), letting the search traverse different sparsity regimes. The *shift-weight move* keeps the asset set fixed but reallocates mass between two selected holdings $i, j \in S$ via a feasible transfer δ, sharpening or flattening exposures while respecting bounds. After selecting the best admissible neighbor under tabu restrictions (with aspiration via F_{best}), the algorithm applies a *Repair* procedure to re-impose feasibility: it first forces $|S|$ into $[N_{\min}, N_{\max}]$ by adding or removing assets, then refreshes (w, z) so excluded assets have $z_i = 0$ and $w_i = 0$, and finally assigns equal weights $w_i = 1/|S|$ over S, which automatically fits the stated $w_{\min}, w_{\max}$ range for $|S| \in$ [15,25].

After the search, the algorithm computes the objective vector (f_1, f_2, f_3) for the best incumbent and calls `UpdatePareto`. This routine keeps only non-dominated portfolios in P, dropping any portfolio that is weakly worse in all objectives and strictly worse in at least one. The remaining set gives a grid-based approximation of the efficient frontier.

Implementation

The Python implementation in a Jupyter Notebook builds three inputs: the return vector r, the covariance matrix Σ, and the ESG-risk vector e. It downloads adjusted close prices for the FTSE MIB universe, computes log-returns, and annualises the sample mean and covariance so r and Σ share consistent units. It then loads ESG risk scores from a local CSV and aligns them with the ticker order to form e. During the search, it evaluates (f_1, f_2, f_3) for the best incumbent and updates the Pareto archive via a dominance check, keeping only non-dominated portfolios across the grid.

4.3 Tabu Search process

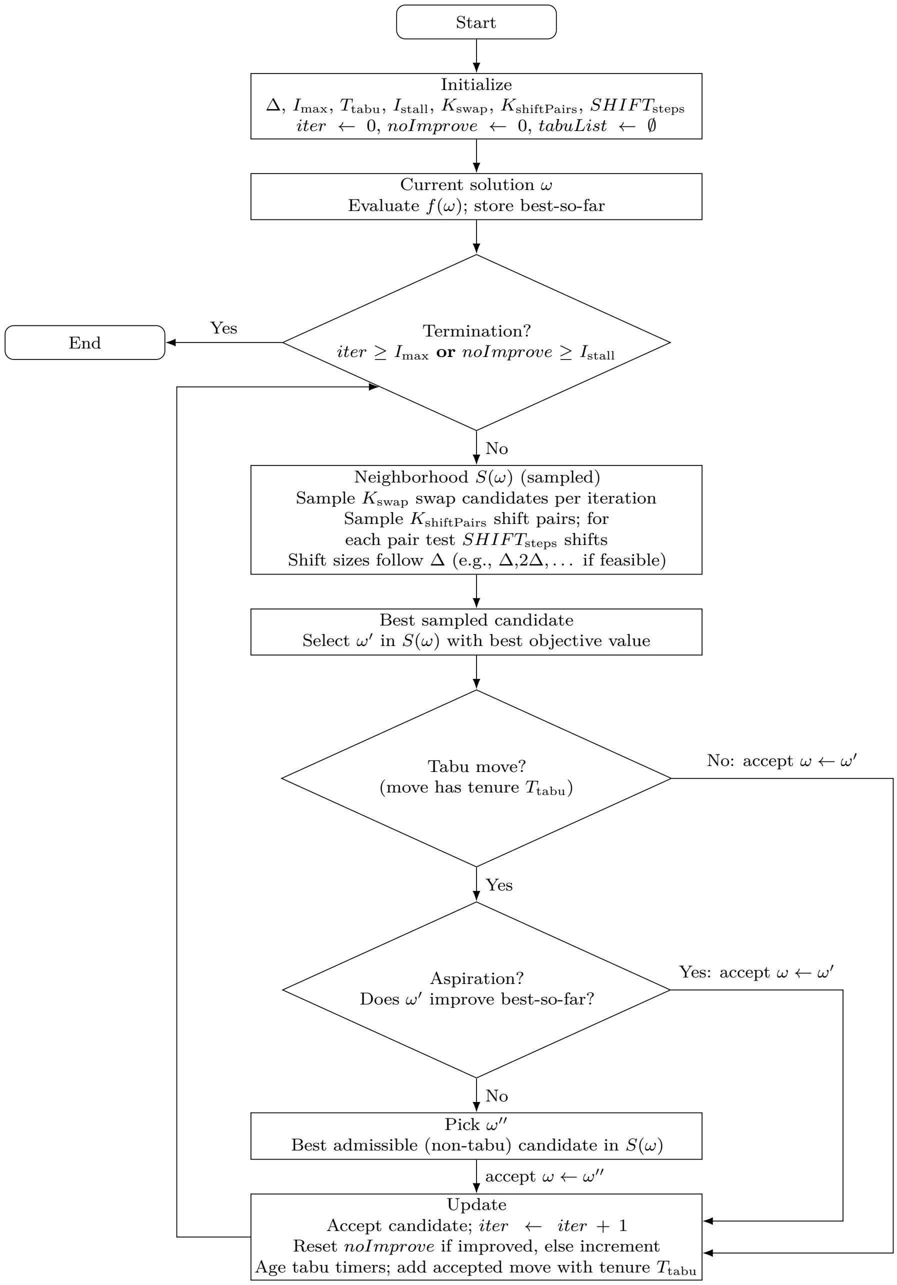

Figure 4.1: Tabu Search flow chart

Interpretation

Flow chart 4.1 starts by fixing the search settings and counters, then it evaluates an initial portfolio ω and stores the best objective value found so far. Two stop checks gate the run: the iteration cap I_{max} and the stall cap I_{stall}, which counts how long the search goes without improving the best score. If neither limit triggers, the method builds a sampled neighborhood $S(\omega)$ rather than scanning every possible move. It draws K_{swap} swap candidates and $K_{\text{shiftPairs}}$ shift pairs, then tests up to $SHIFT_{\text{steps}}$ discrete transfers per pair, with move size aligned to the grid step Δ. This sampling budget keeps the per-iteration workload predictable while still exposing diverse local perturbations.

After ranking the sampled candidates, the algorithm proposes ω' as the best move in $S(\omega)$ and tests it against the tabu list. With tenure T_{tabu}, tabu status blocks recently used move patterns to limit short cycles and force exploration. If ω' is tabu, the aspiration rule can still accept it when it improves the best-so-far objective value; otherwise the search takes ω'', the best admissible move in the sample.

The update step applies the accepted move, increments the iteration count, and resets the no-improvement counter only after a new best value; if not, the stall counter moves toward I_{stall}. It also ages the tabu list and records the accepted move with tenure T_{tabu}. The loop ends when the stopping test triggers, and the procedure returns the best solution found during the run.

Chapter 5

Results and Analysis

5.1 Search setup

Optimization parameters

Parameter	Value
Grid step size (Δ)	0.05
Maximum iteration	200
Tabu tenure (iterations)	10
Stall limit (no-improve iterations)	50
Swap candidates per iteration	80
Shift pairs per iteration	40
Shift steps per pair	9

Table 5.1: Optimization parameters

These settings steer the Tabu Search loop and the way it samples the neighbourhood. The maximum iteration caps the run length, while the stall limit ends the search early when the best weighted-sum score does not improve for many iterations. The tabu tenure keeps recent moves forbidden for a while, which cuts short cycling and pushes exploration; the aspiration rule still allows a tabu move when it beats the current best score.

The remaining parameters shape how many candidate moves you actually evaluate per iteration. The swap candidates per iteration fixes the sampling budget for swap moves, so the method does not scan an oversized combinatorial set. The shift pairs per iteration and the shift steps per pair control the shift-weight operator: first choose a limited number of asset pairs, then test a small ladder of feasible weight transfers, which sets the move granularity and runtime. Grid step size sets the outer weight grid used to generate the weighted-sum coefficients; with a coarse step you cover the trade-off surface quickly, but you also accept a rougher sweep of preference weights.

Objective scaling

Reference Portfolio	**Objective Value**
Maximum Return (OF1)	0.434
Minimum Variance (OF2)	0.015
Minimum ESG Risk (OF3)	12.563

Table 5.2: Extreme objective functions

Table 5.2 shows the extreme reference portfolios and the corresponding value for the objective function of interest. OF1 sits at the top of the return scale with a value of 0.434, but this figure is not on the same numerical axis as the other two criteria. OF2's variance optimum is less than 0.02 and OF3's ESG score is larger than either return or variance.

The objectives use different scales, so a raw sum would let the largest-magnitude term dominate. Normalising by the extreme values puts the criteria on comparable ranges. The weighted sum can then reflect preferences rather than units.

Grid exploration

Once we have normalised the objectives, we can generate a grid of weight triples (v_1, v_2, v_3) such that $v_1 + v_2 + v_3 = 1$ and $v_i \geq 0$. A discrete step size, for instance $\Delta = 0.05$, generates a finite but comprehensive set of weight combinations that cover the simplex of possible preferences. Each triple is then transformed into normalised weights $(\alpha_1, \alpha_2, \alpha_3)$ using the extreme values from the table above. For

$\Delta = 0.05$, a standard combinatorial count gives 231 distinct triples. Each weighting scheme drives a separate optimisation run, producing portfolios that approximate different regions of the trade-off surface.

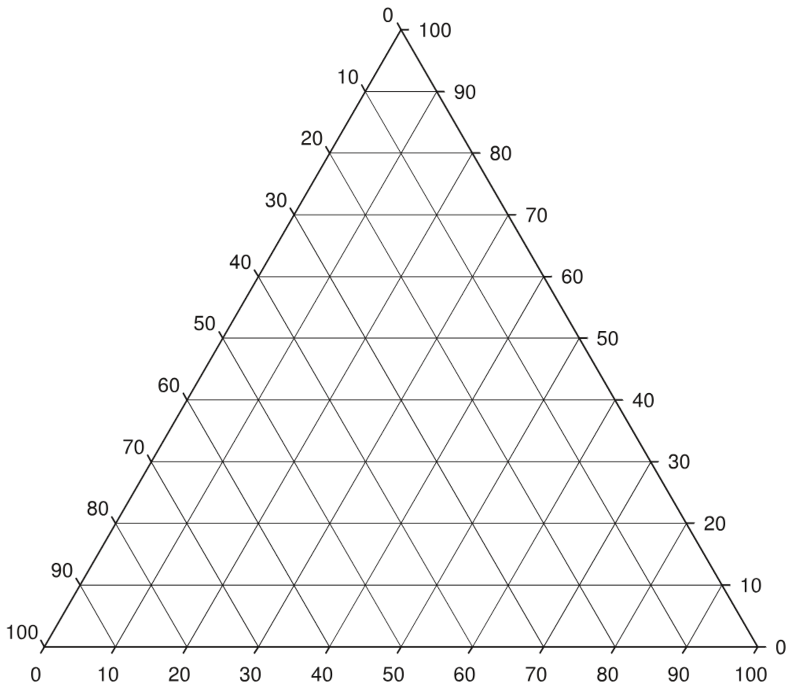

Figure 5.1: Grid Triangle (Source: Wikipedia.org)

Figure 5.1 shows how the weight space for three objectives can be partitioned into discrete regions. Each intersection of grid lines corresponds to a distinct triple. This picture makes tangible how moving along one set of parallel lines increases one component at the expense of the others, exposing the geometric structure behind the weighted-sum grid search.

5.2 Pareto Front

Pareto archive

The *Pareto archive* is the filtered list of portfolios that survive a dominance check across the three objectives: expected return f_1 (to maximize), variance f_2 (to minimize), and ESG risk score f_3 (to minimize). After each weighted-sum run, the algorithm compares the new candidate with the archive and discards it if an archived portfolio is at least as good on all three objectives and strictly better on one. If the candidate is not dominated, it enters the archive and may expel weaker entries. With $\Delta = 0.05$, the weight grid contains 231 distinct triples (v_1, v_2, v_3), yet the archive ends at size 229, so only a small fraction of grid runs collapse into dominated outcomes.

A *Pareto point* here is one archived portfolio a concrete configuration (w, z) that sits on the non-dominated surface for return, risk, and ESG risk. More return requires higher f_2 or f_3, and lowering f_2 or f_3 requires giving up some return.

Pareto results

The Pareto Front in Figure 5.2 appears as a convex surface in the three-dimensional objective space. Expected return f_1 spans roughly from 0.10 to 0.50. Variance f_2 ranges between about 0.015 and 0.050. ESG risk f_3 lies approximately between 11 and 17. No portfolio improves one dimension without worsening at least one of the others.

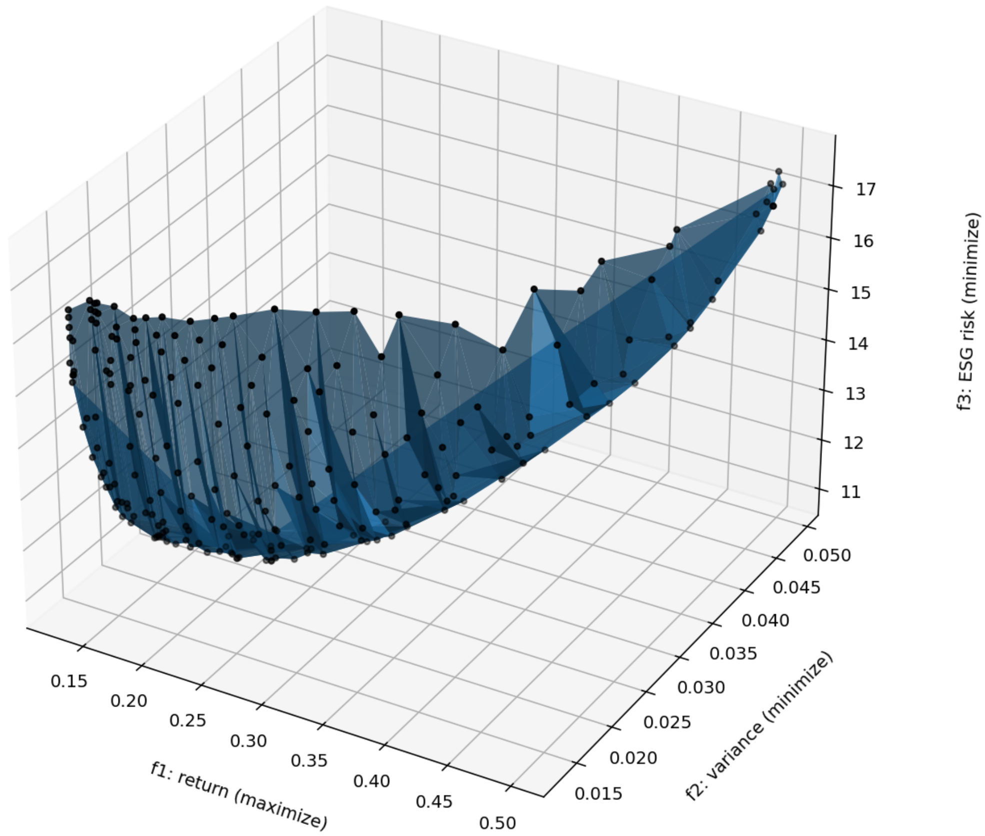

Figure 5.2: Pareto Front plot

Portfolio choice depends on the investor's preference ordering. A risk-averse profile will remain near $f_2 < 0.02$, accepting lower return. An investor targeting high performance may move toward $f_1 > 0.45$, while tolerating variance above 0.04 and ESG risk above 16. If ESG exposure carries binding constraints, then portfolios with f_3 close to 11–13 become feasible candidates. The front defines the admissible set while preferences select the final allocation.

5.3 Computational features

Computational time

Table 5.3 report the time spent *building* neighborhood candidates, averaged per run and summed over the full weight grid (231 runs).

Neighborhood	Mean run (s)	Sum (s)
Shift-weight	3.130	723.126
Swap	1.073	247.960
Add–Remove	0.403	92.982

Table 5.3: Neighborhood timing summary

The table cover neighborhood *generation* time, not the whole Tabu loop. The total Tabu Search lasted 1262.595 seconds. While neighborhoods time amounts to 84.28% of it, the remaining runtime is taken by scoring, ranking, tabu checks, and aspiration logic.

Time complexity

Scoring one neighbour costs $O(n^2)$, dominated by the dense quadratic term $w^\top \Sigma w$. Total runtime scales with the number of neighbours evaluated per iteration, the tabu-search budget I, and the number of weight triples $|V|$ (here $|V| = 231$).

Iteration count

Statistic	Iterations
Average	119
Minimum	76
Maximum	200

Table 5.4: Iteration summary

Table 5.4 sketches the run-length profile of the Tabu loop: the mean sits at 119 iterations, the shortest run stops at 76, and the longest hits the ceiling at 200.

Improvement count

Metric	Shift-weight	Swap	Add/Remove
Total improvements	6181	1988	1142
Improvements per run	26.758	8.606	4.944

Table 5.5: Improvements per neighborhood

Table 5.5 divides the improvements by neighborhood operator. Shift-weight dominates the upgrade count (6181 in total, 26.758 per run), so most progress comes from granular reallocations inside a fixed asset set. Swap plays a supporting role (1988 in total, 8.606 per run), while Add/Remove stays low (1142 in total, 4.944 per run).

5.4 Diversification

Diversification ratio

The estimated diversification ratio is $\mathbf{D}(\mathbf{P}) = \mathbf{1.878}$. The weighted average standalone volatility is $w^{\top}\Sigma = 0.282$, while the portfolio volatility is $\sqrt{w^{\top}Vw} = 0.150$. The spread shows a strong covariance benefit: co-movement cuts total risk well below what isolated volatilities suggest.

The value also points to uneven dependence within the FTSE MIB set. A market factor exists, but sector splits and idiosyncratic shocks still offset part of it. If $D(P)$ drifts toward one during a drawdown, correlations are tightening and the diversification premium is fading, even if $w^{\top}\Sigma$ remains high.

Herfindahl–Hirschman Index

Each constituent's index weight are treated as a share w_i (the published weight in % divided by 100) and computed

$$\mathrm{HHI} = \sum_{i=1}^{40} w_i^2.$$

Using the FTSE MIB indicative weights as at 30 September 2025 published by FTSE Russell. The result is **HHI = 0.073**.

The value exceeds the equal-weight baseline $1/40 = 0.025$, so the index is concentrated. Squaring makes the tilt clear: Intesa Sanpaolo (14.49%) and Unicredit (14.44%) contribute heavily, and Enel (10.31%) and Ferrari (8.08%) reinforce the skew.

Correlation

Figures A.1–A.4 show correlation heatmaps for the FTSE MIB universe. The structure is block-like by sector. Banks (*BMPS, BPSO, BAMI, BPE, ISP, UCG, MB*) move together, with many pairs in the 0.6–0.8 range; *ISP–UCG* is about 0.81. Utilities and regulated networks (*A2A, ENEL, HER, IG, TRN*) form another tight block, often above 0.6. Adding more names from the same block increases count, not diversification.

Cross-block tiles show the links between clusters. *SRG* aligns with the network/utility group (around 0.81 with *TRN*), so it offers little offset when rates or regulation reprice cash flows. Several industrial and smaller names sit nearer 0.1–0.3, while negative entries (down to about -0.06) are sparse and too small to rely on. Two dependence channels dominate: a financial/sovereign-spread channel in banks, and a rate-sensitive regulated channel in utilities and grids.

5.5 Limitations

A weakness of using the FTSE MIB over three years is constituent change: the index does not behave like a sealed basket, and ticker continuity can break when reviews, mergers, spin-offs, or liquidity screens reshuffle membership. That creates small but real measurement frictions. For the Tabu Search, the constituents are treated as a fixed snapshot even though the investable set keeps changing.

The tickers in Table 3.1 match the FTSE MIB membership from 23 Dec 2024 to 19 Sep 2025. The basket later changed: on 22 Sep 2025, *Lottomatica Group* replaced *Pirelli & C*, and on 22 Dec 2025, *Fincantieri* replaced *Interpump Group.* A stricter setup should time-stamp membership and use each stock only for the dates when it is in the index.

Appendix A

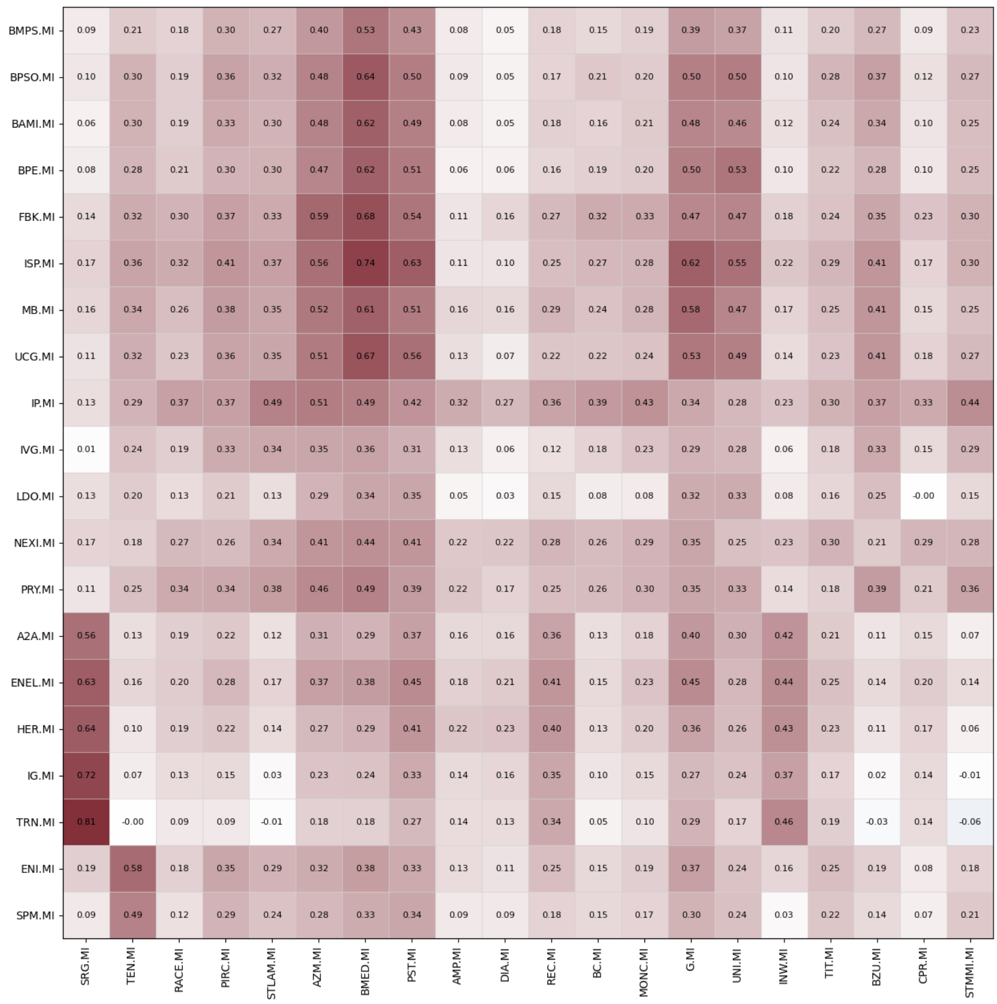

Figure A.1: FTSE MIB Correlation Heatmap Quadrant 1 (Top-Right)

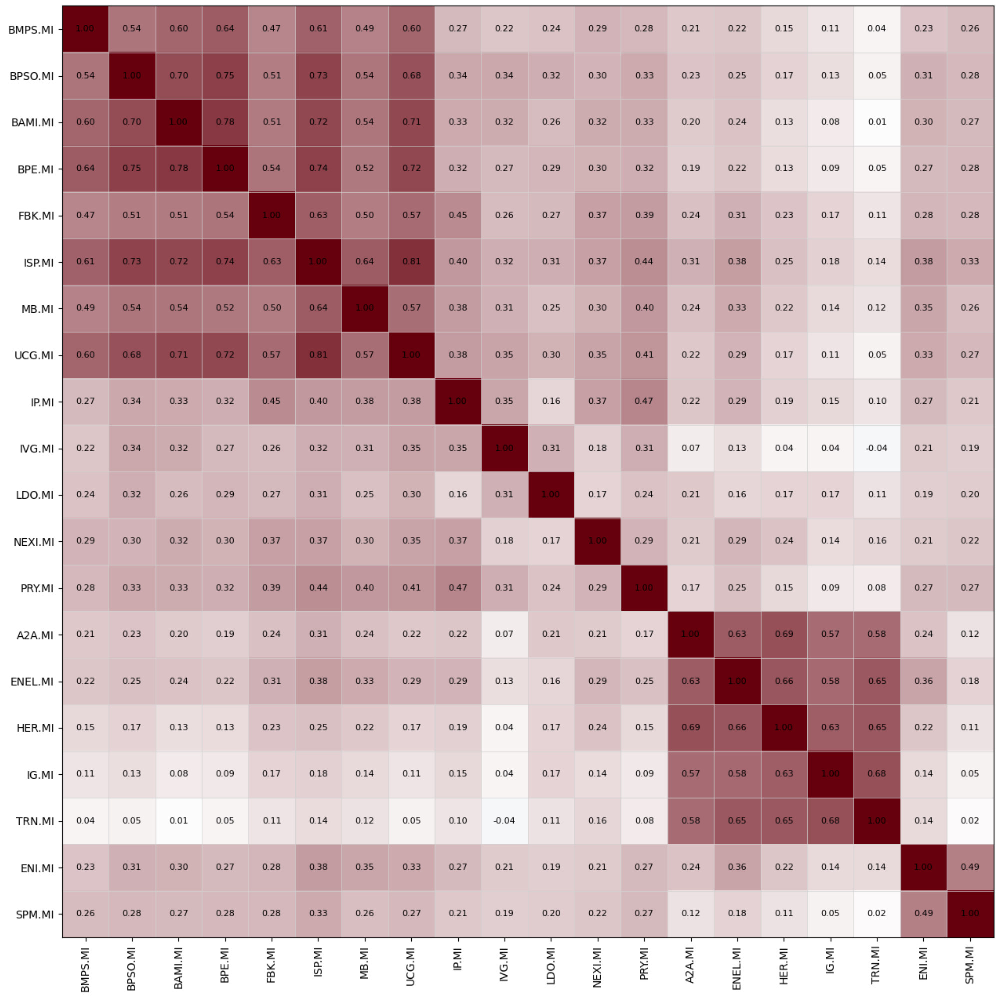

Figure A.2: FTSE MIB Correlation Heatmap Quadrant 2 (Top-Left)

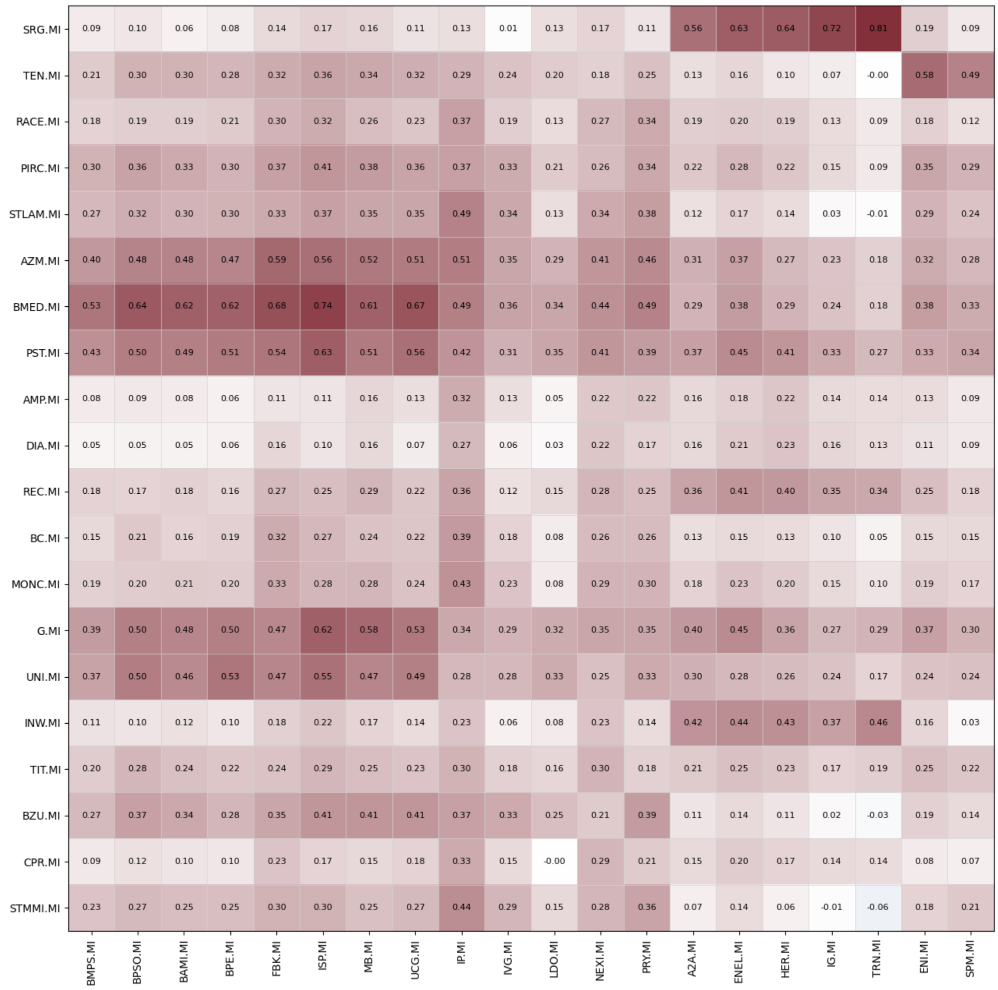

Figure A.3: FTSE MIB Correlation Heatmap Quadrant 3 (Bottom-Left)

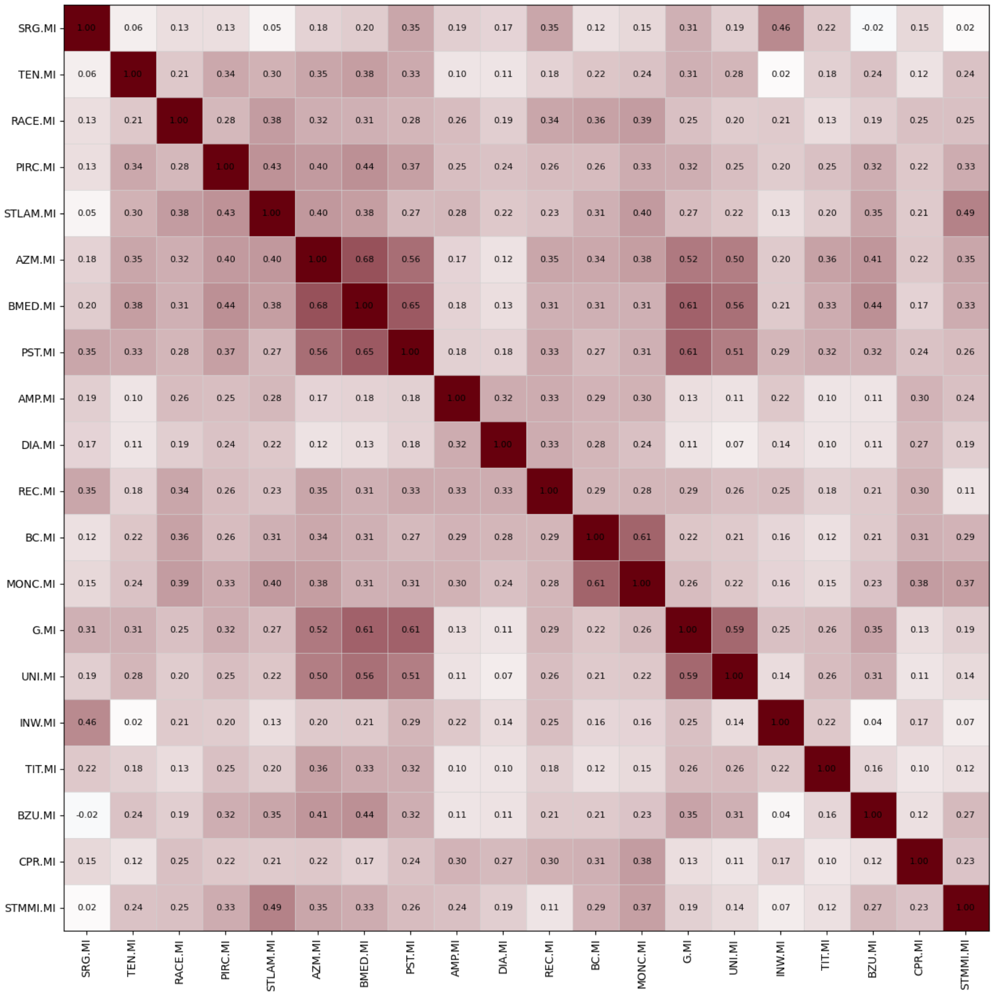

Figure A.4: FTSE MIB Correlation Heatmap Quadrant 4 (Bottom-Right)

Bibliography

[1] H. Markowitz. «Portfolio Selection». In: *The Journal of Finance* 7.1 (Mar. 1952), pp. 77–91 (cit. on p. 3).

[2] G. Abate, I. Basile, and P. Ferrari. «ESG Ratings and Financial Performance: An Empirical Analysis». In: *Int. J. Financial Stud.* 13.4 (Dec. 2025) (cit. on p. 6).

[3] A. Schaerf. «Local Search Techniques for Constrained Portfolio Selection Problems». In: *Computational Economics* 20.3 (Dec. 2002), pp. 177–190 (cit. on p. 8).

[4] O. Grodzevich and O. Romanko. «Normalization and Other Topics in Multi-Objective Optimization». In: *Proceedings of the First Fields–MITACS Industrial Problems Workshop*. The Fields Institute for Research in Mathematical Science. Toronto, Ontario, Canada, Aug. 2006, pp. 89–101 (cit. on p. 11).

[5] Y. Choueifaty and Y. Coignard. «Toward Maximum Diversification». In: *The Journal of Portfolio Management* 35.1 (Oct. 2008), pp. 40–51 (cit. on p. 14).

[6] S. A. Rhoades. «The Herfindahl–Hirschman Index». In: *Federal Reserve Bulletin* 79.3 (Mar. 1993), pp. 188–189 (cit. on p. 15).

[7] Saxo Group. *How correlation impacts diversification: A guide to smarter investing*. Saxo Bank (web page). URL: `https://www.home.saxo/en-mena/learn/guides/diversification/how-correlation-impacts-diversification-a-guide-to-smarter-investing` (cit. on p. 16).