Multidimensional extension of the Modern Portfolio Theory

Auteur : El Azri, Samy
Promoteur(s) : Boniver, Fabien
Faculté : HEC-Ecole de gestion de l'Université de Liège
Diplôme : Master en sciences de gestion, à finalité spécialisée en Banking and Asset Management
Année académique : 2024-2025
URI/URL : http://hdl.handle.net/2268.2/23770

MULTIDIMENSIONAL EXTENSION OF THE MODERN PORTFOLIO THEORY

Jury:

Supervisor -
Fabien Boniver

Reader -
Fanny Bricart

Master thesis presented by
Samy EL AZRI
To obtain the degree of
MASTER IN MANAGEMENT
with a specialization in
Banking and Asset Management
Academic year 2024/2025

Ecole de Gestion de l'Université de Liège

Acknowledgements

First and foremost, I would like to express my sincere gratitude to my supervisor, Fabien Boniver, for his constant availability, attentive guidance, and invaluable advice throughout this project. His rigorous and supportive supervision was essential to the successful completion of this thesis. Thank you for believing in me at the very beginning.

My deepest thanks also go to my parents for their unwavering support, constant encouragement, and reassuring presence throughout this demanding journey. Most of all, I am profoundly grateful to my grandfather, whom I called every day and who never failed to ask how the thesis was progressing. Your care and concern meant the world to me, your support was, without a doubt, the most meaningful of all.

I would like to sincerely thank Zainab, whose unwavering support and encouragement carried me through the most challenging phases of this thesis. In moments of doubt, she offered clarity and motivation. Her presence and emotional support were invaluable, especially in the final weeks when persistence mattered most.

Finally, I warmly thank my fellow students Noam and Camille. Sharing the experience of writing our theses side by side, supporting each other through the challenges and uncertainties, was an immense source of motivation and strength.

List of Abbreviations

Acc	Accumulated
AI	Artificial Intelligence
AMH	Adaptive Market Hypothesis
CAL	Capital Allocation Line
CAPM	Capital Asset Pricing Model
CFP	Corporate Financial Performance
Dist	Distributed
EIOPA	European Insurance and Occupational Pension Authorities
ESG	Environmental, Social, and Governance
MASD	Mean-absolute semi-deviation
ML	Machine Learning
MOEA	Multi-Objective Evolutionary Algorithm
MOO	Multi-Objective Optimization
MPT	Modern Portfolio Theory
MV	Mean-Variance
MVL	Mean-Variance-Leverage
MVS	Mean-Variance-Skweness
NAV	Net Asset Value
NSGA-II	Non-Dominated Sorting Genetic Algorithm II
SAA	Strategic Asset Allocation
SCR	Solvency Capital Requirements
SD	Standard Deviation
SDG	Sustainable Development Goal
SR	Sharpe Ratio
SRI	Socially Responsible Investments
TAA	Tactical Asset Allocation

List of Tables

List of Figures

Introduction

Background and Context

Modern Portfolio Theory (MPT), introduced by Harry Markowitz in the mid-twentieth century, revolutionized the way investors approach portfolio construction. By formalizing the trade-off between expected return and risk, MPT laid the foundation for decades of both academic finance and institutional asset management. At its core lies the concept of the efficient frontier: a set of portfolios that maximize return for a given level of risk, or equivalently, minimize risk for a given level of return. This two-dimensional framework, however, was built in a financial era that prioritized purely quantitative, market-driven outcomes, with minimal regard for externalities, regulatory complexities, or sustainability considerations.

Today, the context has changed. Investors, particularly institutional actors such as insurance companies, pension funds, and sovereign asset managers, must operate under a confluence of financial, regulatory, and ethical pressures. On one hand, financial institutions face increasingly stringent capital requirements under regulations such as Solvency II, which impose substantial capital charges for holding high-risk or volatile assets. On the other hand, there is a growing consensus, both societal and regulatory, that investors must account for environmental, social, and governance (ESG) factors in their decision-making processes. What was once considered a non-financial concern has become a critical component of fiduciary duty and long-term value preservation.

Despite these evolving priorities, the tools used in portfolio construction have remained largely two-dimensional. ESG considerations are often treated as afterthoughts, integrated through heuristic screens, exclusions, or scoring overlays, while capital requirements are managed post-optimization, as compliance checks rather than embedded objectives. This fragmented treatment fails to capture the complex trade-offs that institutional investors must navigate when balancing return, risk, sustainability, and capital efficiency. The efficient frontier, in its traditional form, does not fully represent the multidimensional nature of investment decision-making in the twenty-first century.

This thesis responds to that gap by proposing an extension of Modern Portfolio Theory into a 4D space. It incorporates two additional dimensions: Environmental, Social, and Governance (ESG) performance and Market Solvency Capital Requirement (SCR) into the optimization process, alongside the classical return and volatility. In doing so, it aims to offer a more holistic, data-driven framework that reflects the actual constraints and objectives facing modern institutional investors.

Problem Statement

The integration of sustainability and regulatory constraints into investment decisions is no longer optional, it is a structural reality. Yet, the methodologies used to incorporate these factors often remain superficial or externally imposed. ESG scores, when used, are typically integrated as exclusion criteria or as qualitative adjustments to asset allocation models. SCR, although central to the regulatory capital framework for European insurers, is often handled outside the optimization process, used to validate results rather than shape them. This leads to suboptimal portfolios, where trade-offs are not explicitly accounted for, and the opportunity to optimize capital efficiency or sustainability alignment is lost.

At the same time, there is a persistent perception that the inclusion of non-financial or regulatory objectives in optimization models will distort the efficient frontier, making portfolios less attractive in terms of return or risk. Another concern is computational feasibility: as more dimensions are introduced, optimization complexity increases exponentially, raising questions about whether such models can be implemented in practice without relying on heuristic or opaque algorithms.

These challenges underscore the lack of a unified, rigorous framework capable of incorporating ESG and SCR into portfolio optimization in a transparent and replicable way. What is needed is a model that does not treat these objectives as constraints to be managed separately, but as integrated, co-equal dimensions within the decision-making process. This thesis addresses this problem by constructing a four-dimensional efficient frontier, where ESG and Market SCR are embedded directly into the optimization criteria alongside return and volatility. It seeks to test whether this multidimensional framework can yield meaningful, differentiated portfolio solutions without undermining financial performance or computational efficiency.

Objectives

The primary aim of this research is to develop and evaluate a multidimensional extension of Modern Portfolio Theory that treats ESG and Market SCR as formal objectives within the portfolio optimization process. This framework is not intended to replace classical MPT but to enhance it, aligning it with the needs of investors who must consider not only return and risk, but also regulatory capital costs and sustainability mandates.

To achieve this aim, the thesis pursues several interrelated goals. First, it constructs a clean, analytically tractable dataset of 14 asset classes, spanning the period from January 31, 2019, to January 31, 2025. Each asset class is described by four attributes: expected return, standard deviation of returns (as a proxy for risk), ESG score (drawn from Refinitiv and Morningstar), and Market SCR (calculated using Solvency II standard formula methodology). These inputs form the foundation for building the optimization models.

Second, the thesis develops two complementary modeling frameworks. The first is an *a posteriori* model, in which portfolio combinations are generated through exhaustive enumeration and filtered using Pareto efficiency criteria. This model maps the full universe of non-dominated portfolios, free from investor preference biases, thereby establishing the structure of the multidimensional efficient frontier. The second is an *a priori* model, in which investor preferences are embedded directly into the optimization process via weighted scoring and lexicographic sorting. This model simulates how different investor types, such as return-maximizing, low-risk, ESG-focused, or SCR-conscious profiles, would navigate the frontier and select portfolios in line with their priorities.

Third, the thesis evaluates the impact of adding ESG and SCR as optimization objectives on portfolio composition, structure, and performance. It compares the resulting 4D portfolios against classical 2D efficient portfolios, analyzes shifts in asset allocation, and investigates the presence and magnitude of trade-offs among the four dimensions.

Finally, the study assesses the computational feasibility of the model by examining the volume of generated portfolios, runtime, and scalability under tranche-based enumeration. The objective is to test whether a multidimensional model, implemented transparently and deterministically, can serve as a viable tool for real-time portfolio construction in professional settings.

Research Question and Study Design

The central research question guiding this thesis is:

To what extent can ESG and Market SCR be integrated into portfolio optimization as formal dimensions, and how do these additions reshape the efficient frontier and optimal asset allocation?

This question is explored through a rigorous, multi-step study design that combines theoretical development, data integration, empirical modeling, and comparative analysis. The study begins with a review of existing extensions to MPT and relevant literature on ESG integration. It then constructs the necessary dataset by collecting and harmonizing asset-level data on return, volatility, ESG, and SCR. Using this dataset, the study implements two parallel optimization approaches: an *a posteriori* model that filters efficient portfolios from a full enumeration of possible allocations, and an *a priori* model that selects portfolios based on preference-weighted objective scores.

Portfolios are generated using a systematic enumeration technique with a 3.75% granularity, which ensures full coverage of the feasible allocation space while maintaining tractability. The 4D frontier is analyzed through numerical and visual tools, including 2D projections, radar plots, and allocation tables. Special attention is given to the performance and characteristics of portfolios aligned with specific investor profiles. Comparisons are also made between the 4D and 2D models to evaluate whether the inclusion of ESG and SCR produces substantively different portfolio outcomes and trade-offs.

Throughout, the analysis is guided by five core hypotheses: that the 4D framework provides a more holistic optimization structure (H1); that ESG and SCR can be effectively integrated without distorting the frontier (H2); that asset allocations shift meaningfully in response to multidimensional objectives (H3); that return and risk performance are preserved (H4); and that the model remains computationally tractable (H5). By the end of the study, these hypotheses are tested and interpreted in light of empirical results, with broader implications for both portfolio theory and investment practice.

1 Chapter – Literature Review

This literature is divided into three main chapters. The first chapter revisits they key concepts of Modern Portfolio Thoery (MPT) and its foundational models. This will help the reader understand what has been done with the classical theory, its drawbacks, and potential for improvements. The second chapter examines Socially Responsible Investments (SRI), ESG integration, and stock behaviours when regulatory risk is taken into account, with a particular focus on their implications for portfolio management. While it will not be the focus on the next parts of the thesis, understanding why we consider these factors might be insightfull. Finally, the third chapter reviews the existing literature on multidimensional portfolio optimization, including models that extend MPT by incorporating criteria such as ESG, leverage, and tail risk.

1.1. Introduction to Modern Portfolio Theory

Modern Portfolio Theory, first introduced by *Harry Markowitz (1952)* in his seminal paper "Portfolio Selection", represents one of the most significant advancements in the field of finance. By shifting the focus from individual asset selection to portfolio optimization, MPT introduced a systematic and quantitative approach to investing. Prior to MPT, investors primarily focused on picking individual securities that appeared attractive based on return potential. However, Markowitz's theory changed the perspective by asserting that it is not only the performance of individual assets that matters but how assets interact with one another when combined into a portfolio.

The key principle of MPT is the notion of diversification *(Rubinstein, 2002)*, which states that by combining a variety of assets with low or negative correlations, investors can reduce the overall risk of their portfolio. This approach works because when one asset performs poorly, others may perform well, thus offsetting the losses. This reduction in risk is particularly important because MPT identifies and separates risk into two categories: systematic risk (market-wide risk) and unsystematic risk (specific to individual assets). Through diversification, unsystematic risk can be minimized, but systematic risk, which is influenced by factors such as interest rates, inflation, or market-wide shocks, cannot be entirely eliminated *(cf. Section 1.5).*

At the core of MPT lies a mathematical framework known as mean-variance optimization. In this framework, risk is quantified as the variance or standard deviation of the returns of a portfolio (σ_P^2), while the expected return is calculated as the weighted average of the expected returns of individual assets (μ) within the portfolio. The key to portfolio optimization, according to MPT, is finding the optimal allocation of assets that maximizes return for a given level of risk or minimizes risk for a given level of return. This optimization is achieved through the efficient frontier, which represents a set of portfolios that offer the highest return for each level of risk. Mathematically, Markowitz's model optimizes portfolios by minimizing variance subject to a return constraint *(Lassance, 2021)*:

$$\min_{w \in W} \sigma_P^2 \ subject\ to\ \mu_P \geq \mu_0 \tag{1}$$

Where:

- σ_P^2 is the portfolio variance
- μ_P is the average expected returns of the assets composing the portfolio
- μ_0 is the minimum acceptable return

Mathematically, the average expected returns of the assets composing the portfolio is:

$$\mu_P = \sum_{i=1}^{n} w_i \times \mu_i \tag{2}$$

Where:

- w_i and μ_i represent the weight and the return of asset i, respectively.

Mathematically, the portfolio variance, denoted by σ_P^2, is expressed as:

$$\sigma_p^2 = \sum_{i=1}^{n}\sum_{j=1}^{n} w_i w_j Cov(r_i, r_j) \tag{3}$$

Where:

- w_i and w_j are the weights of assets i and j,
- $Cov(r_i, r_j)$ represents the covariance between the returns of assets i and j.

As said before *(Rubinstein, 2002)*, the concept of covariance is central to MPT, as it captures the relationship between asset returns. If two assets are positively correlated (i.e., their returns tend to move in the same direction), their combined risk will be higher. On the other hand, assets with low or negative correlations help reduce the overall portfolio risk, even if the assets themselves are volatile. Formula 4 is giving the correlation between two assets return is given by:

$$Cov\left(r_i, r_j\right) = \frac{\sum(r_i - \overline{r_i}) \times \left(r_j - \overline{r_j}\right)}{N - 1} \tag{4}$$

Where:

- r_i and r_j are the returns of assets i and j
- $\bar{r}_i$ and $\bar{r}_j$ are the average of assets i and j
- N is the sample size

Thus, the correlation coefficient measures the relationship between asset returns. Diversification benefits arise from low or negative correlations.

1.2. The Efficient Frontier and Risk-Return Trade-Off

Markowitz's concept of the efficient frontier is one of the most important contributions of MPT. The efficient frontier represents the set of portfolios that offer the highest expected return for a given level of risk, or conversely, the lowest risk for a given level of return. Portfolios that lie on the efficient frontier are considered optimal because they provide the best possible return for each unit of risk. Any portfolio below the efficient frontier is suboptimal, as it offers lower returns for the same level of risk. The efficient frontier is typically visualized as a curve in a two-dimensional risk-return space. The x-axis represents the risk (usually the standard deviation or variance of returns), and the y-axis represents the expected return of the portfolio (Figure 1-1). Portfolios on the efficient frontier are considered optimal because they lie at the maximum return-to-risk ratio.

For example, in a portfolio containing two assets, one with a high return and another with low return, the combination of these two assets may lead to a more favorable risk-return trade-off than holding either asset alone. By allocating the appropriate weights to each asset, an investor can construct a portfolio that provides a higher expected return for a given risk level than if they were to invest in individual assets. Thus, the efficient frontier demonstrates that by carefully combining assets, an investor can achieve higher returns while minimizing overall portfolio risk. The tangent portfolio on this frontier achieves the highest Sharpe Ratio *(Hull, 2021)*.

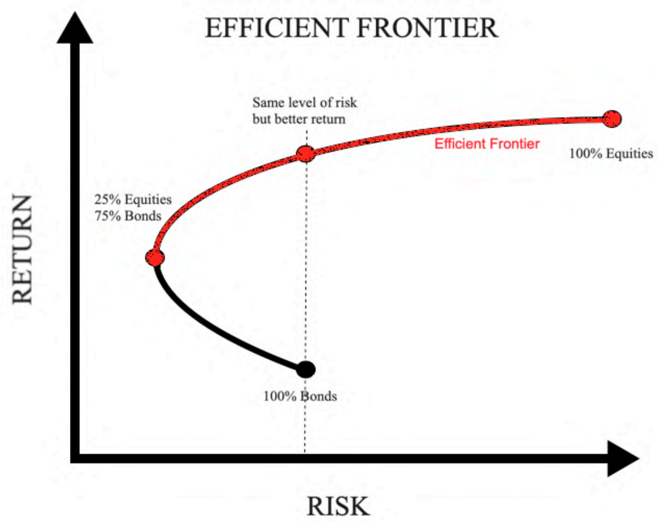

Figure 1-1: Efficient Frontier Representation (Fahy, 2024)

1.3. Sharpe Ratio

As mentionned in *Section 1.2,* the Sharpe Ratio, developped by William Sharpe in *1966,* is a performance indicator widely used in portfolio management. It is also referred to as the reward-to-variability ratio. It measures volatility-adjusted performance, which represents the average return obtained in excess of the risk-free interest rate per unit of total risk.

$$Sharpe\ Ratio\ (SR) = \frac{R_P - r_f}{\sigma_P} \quad (5)$$

Where:

- R_P is the return of the portfolio
- r_f is the risk-free interest rate
- σ_P is the standard deviation of the portfolio's excess return

SR can have a value ranging from $-\infty; \infty$. If the ratio is negative, we can assume that the portfolio of securities performs poorer than a risk-free investment while taking a greater risk. If the ratio has value ranging from 0 to 1, it means that the excess return over r_f is lower than the risk taken. Finally if SR is greater than 1, the portfolio delivers superior return for its level of volatility (*François & Hübner, 2024*). When plotting the Capital Allocation Line (CAL) against the efficient frontier, the intersection is the tangent portfolio. The point of intersection represents the portfolio with the highest Sharpe Ratio. In other words, the portfolio that provides the highest return for the lowest level of risk (*Appendix 1.1.1*).

1.4. MPT Assumptions and Criticisms

While MPT has fundamentally changed how investors approach portfolio construction, it is not without its assumptions and limitations. One of the key assumptions of MPT is that investors are rational and risk-averse. According to this assumption, investors will seek to maximize their utility by choosing the portfolio that offers the highest expected return for the least amount of risk, or vice versa. This assumption is important because it implies that investors' decisions are based solely on the expected return and risk of portfolios, without considering other factors like behavioral biases or emotions *(Mangram, 2013).*

Another important assumption of MPT is that asset returns follow a normal distribution. This assumption allows MPT to quantify risk as the variance or standard deviation of returns. In a normally distributed world, extreme events (often called "tail risks") are rare, and the probability of large market moves is low. However, Mangram *(2013)* showcases that financial markets often exhibit non-normal characteristics, such as fat tails and skewness, meaning that extreme market events (e.g., crashes) occur more frequently than MPT predicts. This limitation has been a major criticism of MPT, as it underestimates the likelihood and impact of such extreme events.

Furthermore, MPT assumes that markets are efficient, meaning that all information is reflected in asset prices. This assumption leads to the conclusion that there is no way for investors to consistently outperform the market because asset prices already incorporate all publicly available information. However, in reality, markets are not always efficient *(Kim and Francis, 2013).* Behavioral finance research has shown that markets can be influenced by cognitive biases, irrational behavior, and informational asymmetries, which can lead to mispricing of assets and deviations from the efficient market hypothesis.

1.4.1. Behavioral Finance

Behavioral finance challenges the assumptions of rationality and efficiency underlying MPT and the Capital Asset Pricing Model. It posits that investors do not always act in a rational manner due to cognitive biases, emotional reactions, and social factors. The emergence of behavioral finance has significantly influenced modern finance theory by introducing psychological insights into decision-making *(Lo, 2004).*

1.4.2. The Role of Behavioral Biases

Behavioral finance emphasizes the impact of psychological factors such as overconfidence, loss aversion, and herding behavior on investment decisions. These biases often lead to suboptimal decisions that deviate from the predictions of traditional finance models. For example, loss aversion, where investors feel the pain of losses more acutely than the pleasure of gains, can lead to risk-averse behavior, such as holding on to losing investments for too long in the hope of a recovery. Similarly, overconfidence can lead investors to overestimate their ability to predict market movements, leading to excessive trading and under-diversification *(Lo, 2004).*

1.4.3. The Adaptive Markets Hypothesis

The AMH, proposed by Andrew Lo *(2004)*, provides an alternative view to MPT's assumption of market efficiency. AMH suggests that financial markets are not always efficient and that market participants adapt their behavior over time based on changing conditions. Unlike MPT, which assumes that investors' decision-making processes are static, AMH posits that investors evolve their strategies in response to shifts in market conditions, learning from past experiences and adjusting to new information. AMH views financial markets as dynamic and evolving, with investors continually adapting to new information, economic conditions, and market events. According to AMH, the efficient frontier is not a fixed concept but rather a dynamic one that changes over time in response to market conditions and investor behavior.

1.5. The Capital Asset Pricing Model (CAPM)

The Capital Asset Pricing Model, developed by William Sharpe in 1964, is a fundamental extension of MPT that refines the relationship between an asset's expected return and its associated risk in a portfolio. CAPM builds on MPT's insights by introducing the concept of systematic risk, the portion of risk that cannot be diversified away, and helps determine the expected return of an asset based on its exposure to the broader market. CAPM has become a cornerstone of modern finance, offering a simple yet powerful framework for pricing risky assets and understanding the trade-offs between risk and return.

1.5.1. CAPM Overview

The core concept behind CAPM is that the expected return of an asset is directly related to its exposure to market risk, represented by the asset's beta (β). Beta measures the sensitivity of an asset's returns to changes in the overall market returns. If an asset has a beta greater than 1, it is more volatile than the market, meaning that it tends to move more dramatically in response to market changes. Conversely, an asset with a beta less than 1 is less volatile than the market. An asset with a beta of 1 moves in perfect correlation with the market *(François & Hübner, 2024).*

CAPM assumes that markets are efficient, all investors have access to the same information, and the relationship between risk and return is linear. The model asserts that the expected return on any asset can be calculated using the following formula:

$$R_i = R_f + \beta_i (R_m - R_f) \tag{7}$$

Where:

- R_i is the expected return of asset i,
- R_f is the risk-free rate (often represented by the return on government bonds),
- β_i is the asset's beta, measuring its sensitivity to the market,
- R_m is the expected return of the market, and
- $(R_m - R_f)$ is the market risk premium, which is the excess return over the risk-free rate that investors expect from the market.

In this equation, $(R_m - R_f)$ represents the market risk premium, which compensates investors for taking on the risk associated with the market as a whole. The asset's return is adjusted for its beta, reflecting the extent to which it moves with the market. CAPM implies that the higher an asset's beta, the higher the expected return, because investors demand higher returns to compensate for higher risk.

1.5.2. Systematic vs. Unsystematic Risk

A key contribution of CAPM is the distinction between systematic and unsystematic risk. Systematic risk refers to market-wide risks that affect all assets to some degree, such as changes in interest rates, economic recessions, or geopolitical events. This type of risk cannot be diversified away by holding a portfolio of assets, as it impacts the entire market. On the other hand, unsystematic risk is specific to individual assets or industries, such as the risk of a company's management issues, labor strikes, or regulatory changes. This type of risk can be mitigated through diversification, as different assets react differently to these factors. CAPM focuses solely on systematic risk because it is the only risk that compensates investors for taking on exposure to the broader market. According to CAPM, an asset's expected return is determined by its exposure to systematic risk, which is captured by its beta. Therefore, the greater an asset's beta, the higher the expected return, as investors require additional compensation for bearing more market risk.

1.5.3. CAPM Assumption

While CAPM provides a powerful framework for understanding risk and return, it is built upon several key assumptions, which have been the subject of much critique over the years. Some of the main assumptions of the model include:

- **Efficient Markets:** CAPM assumes that markets are efficient, meaning that all available information is already reflected in asset prices. In an efficient market, there are no mispricings, and assets are always fairly priced.

- **Rational Investors:** The model assumes that investors are rational and risk-averse, meaning they make decisions based purely on expected returns and risk, and seek to maximize utility.

- **No Transaction Costs or Taxes:** CAPM assumes that investors can buy and sell assets without incurring transaction costs or taxes, allowing them to freely diversify and rebalance their portfolios.

- **Homogeneous Expectations:** All investors are assumed to have the same expectations about asset returns, risks, and correlations. This assumption implies that all investors will hold the same optimal portfolio for a given level of risk.

While these assumptions simplify the mathematical modeling process, they do not necessarily reflect real-world conditions. The assumption of market efficiency, for instance, has been widely debated, as real markets often exhibit inefficiencies, such as information asymmetry or investor behavioral biases *(cf. Section 1.4).* Moreover, the assumption of rational behavior conflicts with insights from behavioral finance, which shows that investors are often influenced by emotions and cognitive biases, leading them to make decisions that do not always align with the CAPM framework.

1.5.4. Criticisms of CAPM

Despite its theoretical appeal and widespread use, CAPM has faced significant criticisms, particularly in its empirical validity and assumptions. One of the main critiques is that the model oversimplifies the relationship between risk and return by relying on beta as the sole measure of systematic risk. While beta is useful for understanding an asset's sensitivity to market movements, it fails to capture other relevant risk factors that may affect asset prices *(Hübner et al, 2014).*

Another criticism is the empirical failure of CAPM to fully explain asset returns. Researchers have found that factors such as company size (the size effect) and value versus growth stocks (the value effect) play significant roles in asset returns but are not accounted for in the CAPM model. These anomalies have led to the development of alternative models such as the Fama-French three-factor model, which incorporates size and value factors in addition to market risk *(Fama and French, 2015).*

1.6. Strategic Asset Allocation

Strategic Asset Allocation (SAA) is a key concept in portfolio management that focuses on determining the optimal allocation of assets across various asset classes based on an investor's risk tolerance, financial goals, and investment horizon. Unlike tactical asset allocation (TAA), which is more short-term and market-driven, SAA aims to create a long-term asset mix that will achieve the desired financial objectives while minimizing risk *(Black and Litterman, 1992).*

SAA involves optimizing the mix of asset classes (such as equities, bonds, real estate, and commodities) to maximize return for a given level of risk or to minimize risk for a given expected return. The application of SAA is particularly important for institutional investors, such as pension funds and insurance companies, which need to balance their financial goals with long-term investment horizons.

1.6.1. Theoretical Foundations of SAA

SAA is based on the principles of MPT, which advocates for diversification to reduce unsystematic risk. In MPT, the risk of a portfolio is not simply the sum of the risks of individual assets, but is determined by the correlation between those assets. By holding assets that are not perfectly correlated, an investor can reduce overall portfolio risk *(cf. Section 1.1).*

One way to determine the theoretical basis of SAA stems from the mean-variance optimization framework, where the objective is to allocate capital to various assets in a way that maximizes the portfolio's return while minimizing its risk. As explained in *section 1.2,* the efficient frontier is a key concept in MPT and SAA. It represents the set of portfolios that provide the highest expected return for a given level of risk. The efficient frontier is derived by plotting the expected return of portfolios against their risk. Portfolios lying on the efficient frontier are considered optimal because they offer the best possible return for a given level of risk *(Zhang, 2024).*

1.6.2. Mathematical Formulation of SAA

The portfolio's expected return (R_p) is the weighted sum of the expected returns of the individual assets in the portfolio. The mathematical formulation can be found in *section 1.1* with Formula 2. The portfolio's risk (variance), denoted by σ_p^2, is computed as the weighted sum of the variances of the individual assets and the covariances between the returns of the assets. Again, refer to *section 1.1* with Formula 3. The optimization problem for SAA can be formulated as follows:

$$Maximize: R_p = \sum_{i=1}^{n} w_i \times R_i \tag{6}$$

$$Subject\ to: \sum_{i=1}^{n} w_i \times \sigma_i^2 \leq Risk\ Tolerance \tag{7}$$

Where:

- $Risk\ Tolerance$ represents the investor's maximum acceptable level of risk.

This mathematical formulation allows investors to optimize their asset allocation by balancing the expected return and risk, taking into account the correlations between asset returns. The challenge lies in accurately estimating the expected returns, variances, and covariances of the assets in the portfolio, which requires historical data and economic forecasts *(Campbell et al., 2003).*

1.6.3. Factors Affecting Asset Allocation Decisions

Several factors influence asset allocation decisions in SAA *(Meucci, 2005, pp. 239–243):*

- **Risk Tolerance**: This refers to the degree of risk an investor is willing to take on. Investors with higher risk tolerance may allocate more to volatile asset classes like equities, while those with lower risk tolerance might allocate more to bonds or other lower-risk investments.

- **Investment Horizon**: Long-term investors, such as pension funds, typically have a higher capacity to take on risk, as they can afford to wait out short-term market fluctuations. In contrast, investors with shorter time horizons, such as those saving for retirement in the near future, may prefer more conservative portfolios with lower risk.

- **Financial Objectives**: The investor's goals, such as income generation, capital preservation, or growth, will affect the allocation. For example, income-focused investors may prefer to allocate more to bonds and dividend-paying stocks, while growth-focused investors might allocate more to equities.

- **Market Conditions**: The broader economic environment, including factors like interest rates, inflation, and economic growth, can impact asset returns and influence asset allocation decisions. For instance, during periods of low interest rates, bonds may offer lower returns, prompting investors to allocate more to equities or other assets.

- **Liquidity Needs**: Some asset classes, such as private equity or real estate, are less liquid than stocks and bonds. Investors with short-term liquidity needs may prefer more liquid assets like stocks or money market instruments.

1.6.4. Challenges and Limitations of SAA

While SAA is a powerful tool for portfolio management, it comes with several challenges *(Meucci, 2005, pp. 239–243):*

- **Forecasting Risk and Return**: SAA relies heavily on the estimation of expected returns, volatilities, and correlations, which are based on historical data and future projections. These forecasts are inherently uncertain, and inaccurate estimates can lead to suboptimal asset allocations.

- **Changing Market Conditions**: SAA assumes that the relationships between asset returns (i.e., their correlations and variances) remain relatively stable over time. However, market conditions can change rapidly, and the correlations between asset classes may shift during times of economic crisis or market turmoil. For example, during the 2008 financial crisis, many assets that were previously uncorrelated, such as stocks and bonds, moved in the same direction, leading to higher portfolio risk than expected.

- **Behavioral Biases**: Investors are often influenced by cognitive biases, such as overconfidence or loss aversion, which can lead to suboptimal decisions. These biases can cause investors to deviate from their optimal allocation, particularly in times of market stress or volatility.

- **Liquidity Constraints**: Some asset classes, such as real estate or private equity, may not be as liquid as others. If an investor needs to sell assets quickly to meet liquidity requirements, these less liquid assets could pose challenges.

2 Chapter

2.1. Social Responsible Invetsments and Regulatory Risks for Financial Institutions

Portfolio optimization has traditionally centered on the trade-off between risk and return, with seminal contributions by Harry Markowitz *(1952)* introducing the mean-variance framework (*cf. Section 1.1*). Over time, this model evolved to accommodate real-world frictions and investor preferences. In recent decades, two dimensions, ESG factors and market regulatory risks, have emerged as critical considerations in portfolio design *(Gasser et al, 2017; Official Journal of the European Union, 2025).*

Socially Responsible Investing is grounded in stakeholder theory and ethical investing paradigms, with foundational work from scholars such as Friede, Busch, and Bassen (*2015*), who conducted a meta-analysis demonstrating a positive relationship between ESG performance and financial returns. The emergence of the triple bottom line and concepts like “doing well by doing good” challenged the traditional view that ESG investing entails a performance sacrifice.

Market reulatory risks, on the other hand, has been historically viewed as an exogenous constraint or risk. For instance, Amihud and Mendelson (*1986*) were among the first to model liquidity as a determinant of asset pricing. The incorporation of liquidity risk into portfolio optimization became especially relevant after the 2008 financial crisis and the introduction of regulatory regimes like Solvency II. These regulatory regimes introduced stress factors surch as liquidity, currency and many other risks an asset may be subject to. The aim was to oblige financial institutions to take these into accounts when constructing their portfolios. These stress factors indicate the amount of capital required to hold the asset in the portfolio. It is an additional capital buffer that is being added depending on the asset’s exposure to the relevant stress factors (*EIOPA, 2024*). The intersection of ESG and regulatory risk within portfolio theory gives rise to new models that generalize Markowitz's 2D efficient frontier to 3D and 4D spaces, introducing multidimensional optimization problems. These theoretical extensions are increasingly being tackled using tools from operations research and machine learning. Key debates in this domain include:

- Whether ESG integration enhances or detracts from financial performance.
- The optimal method to quantify market liquidity in portfolio models.
- The trade-offs between risk, return, regulatory risk, and ethical considerations.

2.2. Background, Debates, and Areas of Further Discussion

The literature on ESG integration in portfolio optimization is rich but diverse, with studies using a wide array of methodologies and arriving at varying conclusions. However, market regulatory risks integration was not studied as much as ESG integration. This may be due to the difficulty and the niche segment where this dimension may be useful. It is therefore useful to structure the review according to three core themes: methodologies, major findings, and ongoing debates.

2.2.1. Theme 1: Methodologies Used in the Field

A majority of studies employ quantitative approaches, often grounded in classical or extended portfolio theory. Techniques such as mean-variance optimization, stochastic programming, multi-objective optimization, and robust optimization have been used to incorporate ESG scores into portfolio models. For example, Bianchi and Bigio (*2022*) utilize a robust multi-objective approach to integrate ESG preferences with return and risk, showing how portfolio outcomes change based on different ESG weightings.

Other studies apply simulation-based methods, such as Monte Carlo simulations, to test the impact of ESG filters on portfolio volatility and downside risk. These methods provide the advantage of modeling uncertainty and dynamic changes in asset characteristics, which is particularly important for ESG data that are often updated periodically and influenced by external shocks (e.g., regulatory changes, stakeholder activism) *(Silva et al, 2021).*

Qualitative approaches have also been applied in institutional and behavioral finance studies. These investigate how asset managers incorporate ESG principles into investment decisions or how investors perceive the regulatory risks of sustainable assets. Although less common, such approaches provide valuable insights into decision-making processes and strategic alignment with sustainability goals *(Bianchi and Bigio, 2022; Brière and Ramelli, 2021a).*

More recently, machine learning techniques have been adopted to address the high dimensionality of integrating multiple ESG factors. Tools like neural networks and clustering algorithms are now being used to identify latent relationships between ESG indicators and market signals, offering more adaptive and data-driven optimization strategies *(Bengio et al., 2021)*.

Despite the methodological richness, challenges remain. ESG data are often inconsistent across providers, leading to model sensitivity depending on the source of ratings *(Berg et al, 2019)*. Similarly, regulatory risk metrics are typically very complicated to compute and assess and may not reflect real-time market behavior. These limitations complicate cross-study comparisons and reduce the reliability of optimization outputs *(Amihud and Mendelson, 1986)*.

2.2.2. Theme 2: Major Findings and Contributions

There is a growing body of empirical work that finds a positive or neutral relationship between ESG integration and financial performance. Friede et al *(2015)*, in their meta-study covering over 2,000 empirical studies, concluded that the majority found a non-negative link between ESG criteria and corporate financial performance (CFP), with a significant portion identifying a positive correlation. This has been interpreted as evidence that ESG factors can act as proxies for better risk management, stronger governance, and long-term value creation.

Several studies have also examined the implications of incorporating liquidity constraints into portfolio optimization. Amihud and Noh *(2002)* demonstrate that portfolios subject to liquidity constraints tend to deviate significantly from the traditional mean-variance efficient frontier, favoring more liquid assets and reducing exposure to small-cap or high-volatility securities. This shift, while potentially lowering expected returns, offers better resilience in stressed market conditions.

Other findings emphasize the complementarity between ESG and liquidity. For example, studies by Nofsinger and Varma *(2014)* show that during periods of market downturns, ESG portfolios tend to outperform conventional portfolios due to investor preference for transparency, governance, and risk-averse behavior, which indirectly enhances liquidity by attracting stable capital. Moreover, sustainable indices often exhibit lower turnover, suggesting lower transaction costs and more stable liquidity profiles.

Nonetheless, there are studies that challenge these optimistic interpretations. For instance, Trinks et al. *(2018)* argue that exclusion-based ESG screening may result in sectoral biases and reduced diversification, which could impair risk-adjusted returns. Similarly, concerns remain over "greenwashing" and the inconsistent application of ESG criteria across firms and industries.

2.2.3. Theme 3: Ongoing Debates and Emerging Areas

One of the most salient debates concerns the heterogeneity of ESG ratings. Research by *(Berg et al, 2019; LaBella et al, 2019)* reveals that ESG scores from different rating agencies often diverge significantly due to differences in scope, measurement techniques, and weighting schemes. They describe the phenomenon of “aggregate confusion”, highlighting that ESG ratings across providers diverge due to differences in methodology, rather than inaccuracy. This lack of standardization undermines the credibility of ESG-based portfolio strategies and makes comparative analysis across studies difficult. It also complicates regulatory oversight and investor communication.

Another active area of debate revolves around the time horizon of ESG benefits. While many proponents argue that ESG integration enhances long-term performance, skeptics question its efficacy over shorter investment cycles, particularly for strategies requiring frequent rebalancing or driven by quarterly benchmarks *(Garcìa et al, 2019)*.

Emerging trends in the field include the growing use of AI and big data analytics to enhance ESG modeling. For example, real-time sentiment analysis from news and social media platforms is increasingly being used to augment ESG scores and detect shifts in market perception. At the same time, the integration of climate risk and carbon intensity metrics is gaining prominence as regulators and institutional investors push for climate-aligned portfolios *(Lahmiri and Bekiros, 2019)*.

3 Chapter

3.1. Overview of Methodologies

The methodologies used in ESG and market regulatory risk portfolio optimization are diverse and evolving. A closer comparative analysis reveals the advantages and limitations of different approaches, their suitability to the research questions posed, and their capacity to model the increasing complexity of modern investment environments.

The traditional mean-variance optimization framework, originating with Markowitz *(1952)*, remains a foundational reference point. It is widely used for its simplicity and tractability, especially in introductory studies. However, its assumptions, including normally distributed returns, quadratic utility, and the absence of market frictions, are increasingly criticized as unrealistic, especially when modeling ESG and liquidity constraints. This method often fails to account for illiquidity premiums, time-varying ESG scores, or nonlinear trade-offs between sustainability and return.

To address these shortcomings, many recent studies have employed multi-objective optimization (MOO) frameworks, such as the Non-dominated Sorting Genetic Algorithm II (NSGA-II), which allow for simultaneous optimization across multiple criteria. These techniques are particularly useful for visualizing and quantifying trade-offs between return, risk, ESG performance. For example, Lahmiri and Bekiros *(2019)* demonstrated the application of evolutionary algorithms to model conflicting investor preferences in SRI portfolios. MOO models are often better suited for dealing with the non-compensatory nature of ESG factors, where a deficit in one dimension (e.g., poor governance) cannot always be offset by strength in another (e.g., high returns).

Another widely adopted approach involves robust and stochastic optimization, which explicitly accounts for uncertainty in model parameters. This is especially important for ESG and regulatory risks variables, which are known to be volatile and influenced by exogenous shocks. Gasser et al, *(2017)* explore robust ESG investing under parameter ambiguity, showing that robust portfolios are less sensitive to ESG data inaccuracies and better perform under model misspecification.

A growing subset of literature explores the use of machine learning (ML) and artificial intelligence (AI). Techniques such as decision trees, support vector machines, and neural networks are increasingly employed to identify hidden patterns in ESG data and predict market shocks events. For instance, ML models can enhance the predictive accuracy of ESG scores by integrating real-time data from news sentiment analysis, corporate disclosures, and macroeconomic signals *(Bengio et al., 2021)*. These methods also help in feature selection and dimensionality reduction when dealing with high-frequency or unstructured data.

3.2. Extension of MPT

Despite these criticisms, MPT remains a cornerstone of modern portfolio management. Over the years, various extensions and modifications to MPT have been developed to address its limitations. One such extension is the Capital Asset Pricing Model (CAPM), which builds on MPT by introducing the concept of systematic risk and offering a way to price individual assets based on their sensitivity to the market as a whole. CAPM helps investors understand the relationship between the expected return of an asset and its risk relative to the market, making it a useful tool for asset pricing and portfolio optimization *(cf. Section 1.5)*.

Other extensions of MPT have sought to include additional factors such as liquidity constraints, ESG considerations, and macroeconomic variables. These extensions allow investors to optimize portfolios not only for risk and return but also for ethical, regulatory, and social considerations. For example, ESG investing, which incorporates factors like corporate governance, social responsibility, and environmental impact, has become increasingly important as investors seek to align their portfolios with their values.

Furthermore, the Solvency Capital Requirement (SCR) is a regulatory measure that ensures financial institutions, particularly insurance companies, maintain adequate capital reserves to cover potential market risks. Thus, it is closely related to liquidity, currency, and other risks constraints imposed on portfolios. The inclusion of a market regulation risk dimension into MPT helps institutional investors meet regulatory requirements while also optimizing their portfolios.

3.2.1. Metaxiotis – A Mean-Variance-Skewness Portfolio Optimization Model

In his 2019 paper, Metaxiotis proposed to extend Markowitz's classical model by incorporating skweness with new algorithms. As said before, MPT assumes asset returns follow the Gaussian distribution *(cf. Section 1.4*). This enables investors to describe returns with the first and second moments of distributions, namely returns and volatility. Nevertheless, it has been demonstrated that asset returns may not follow the normal distribution *(Chunhachinda et al., 1997)*. Thus, the author tells us that skewness, which represents the assymetry level observed in a probability distribution *(Lassance, 2021)*, is also characterizing asset returns. The objective of the paper was to maximize return and skewness and minimize risk *(cf. formulas 8 – 11)*. The Mean-Variance-Skewness (MVS) portfolio optimization model considers three parameters simultaneously with following relationships:

Optimize: $f(w) = \left(f_1(w), f_2(w), f_3(w)\right),$ (8)

Maximize portfolio return: $f_1(w) = \sum_{i=1}^{N} w_i \bar{r}_i,$ (9)

Minimize portfolio risk: $f_2(w) = \sum_{i=1}^{N} \sum_{j=1}^{N} w_i w_j \sigma_i \sigma_j \rho_{ij},$ (10)

Maximize portfolio skewness: $f_3(w) = \sum_{i=1}^{N} \sum_{j=1}^{N} \sum_{k=1}^{N} w_i w_j w_k S_{ijk},$ (11)

Where:

- w and $\bar{r}_i$ represent the weight of asset i and the return on asset i, respectively;
- σ_i represents the standard deviation of stock returns i;
- ρ_{ij} is the correlation between asset i and j and $-1 \leq \rho_{ij} \leq 1$;
- S_{ijk} represents the coskewness between the returns of assets i, j, and k, respectively.

The software NSGA-II develops and elaborates algorithms, incorporates constraints, and obtains results for complex problems. By applying and adapting this software, Metaxiotis tried to validate the model. In other words, he applied NSGA-II, a multi-objective evolutionary algorithm (MOEA), to find efficient frontiers in an MVS setting. He used the FTSE-100 databse to get the relevant input. The author contributed to the literature by proposing an extension of the Mean-Variance model. Furthermore, bu using a MOEA he could handle the difficulty of introducing a third moment in the model. It has allowed him to overcome limitations of the classical MV model that assumes normally distributed returns.

The results of his work showed that incorporating skewness provides more realistic portfolio choices than traditional MV models. The MVS optimization framework yields diversified portfolios with asymmetry preferences (investor's skewness appetite) (Figure 3-1). The NSGA-II algorithm successfully finds high-quality Pareto-optimal frontiers for the MVS model, validated on FTSE-100 stocks. And lastly, more evaluations (100,000) improve performance metrics *(Appendix 1.1.2).*

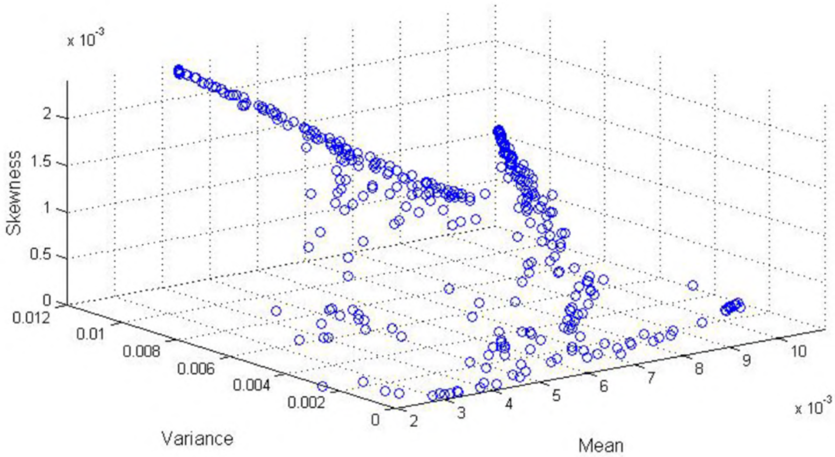

Figure 3-1 The derived efficient frontiers for the MVS portfolio optimization model (Metaxiotis, 2019)

3.2.2. Jacobs and Levy – Mean-Variance-Leverage Model

Jacobs and Levy (*2013*) followed the same principe as Metaxiotis *(cf. Section 3.2.1)* but decided to use leverage as additional parameter. The authors introduce leverage aversion in portfolio theory, extending Markowitz's Mean-Variance model into a Mean-Variance-Leverage optimization framework (MVL). Their motivation stems from the argument that the traditional model does not capture investors aversion to leverage risk, which plays a crucial role in real-world investment decisions. According to them leverage risk includes risks and costs of margin calls – which can force borrowers to liquidate securities at adverse price du to iliiquidity – losses exceeding the capital invested, and the possibility of bankruptcy. In other words, Jacobs and Levy want to build a model that implements investors sacrifice with regards to returns not only to reduce volatility but also to limit leverage. Thus, they introduced a third dimension to the classical theory: a leverage aversion term.

To summarise, the model:

- Explain the unique risks of leverage, including margin calls, bankruptcy risks, and systemic market disruptions.
- Demonstrate the limitations of the conventional MV model when applied to leveraged portfolios.
- Introduce the MVL optimization framework, incorporating leverage aversion directly into the utility function.
- Provide empirical illustrations and efficient frontiers under different leverage tolerance scenarios.
- Propose practical methods for integrating leverage aversion in portfolio optimization and decision-making.

By adding a leverage aversion term to the utility function, the model accounts for both volatility risk and leverage risk, offering a more realitsic representation of portfolio selection behaviour.

$$\text{Maximize utility: } U = \alpha_P - \frac{1}{2\tau_V}\sigma_P^2 - \frac{1}{2\tau_L}\sigma_T^2\Lambda^2, \tag{12}$$

Where:

- α_P is the portfolio expected active return;
- σ_P^2 is the variance of portfolio active return;
- τ_V is the investor's risk tolerance for the variance of portfolio active return;
- $\sigma_T^2\Lambda^2$ is the variance of the leveraged portfolio's total return;
- τ_V is the investor's tolerance for leverage risk.

To solve the MVL Optimization Problem, which is represented by a quartic mathematical problem (Formula 13 and 14 give the utility function to be maximized) – in contrast the the MV Optimization Problem, which is a quadratic mathematical problem – the authors take the quartic of the active-weight variable x_i, including fourth-order cross-products. They then use fixed-point iteration, which allows a quadratic solver to be applied iteratively.

$$U = \sum_{i=1}^{N} \alpha_i x_i - \frac{1}{2\tau_V}\sum_{i=1}^{N}\sum_{j=1}^{N} x_i \sigma_{ij} x_j - \frac{1}{2\tau_L}\left(\sum_{i=1}^{N}\sum_{j=1}^{N}(b_i + x_i)\, q_{ij}\, (b_j + x_j)\right)\left(\sum_{i=1}^{N}|b_i + x_i| - 1\right)^2 \tag{13}$$

Where:

- α_i, b_i, x_i are the expected active return, weight in the benchmark, active weight of security i, respectively.
- σ_{ij} and q_{ij} are the covariance between the active returns and the covariance between the total returns of securities i and j, respectively.

This equation may be rewritten as the set of the following two equations:

$$U = \sum_{i=1}^{N} \alpha_i x_i - \frac{1}{2\tau_V}\sum_{i=1}^{N}\sum_{j=1}^{N} x_i \sigma_{ij} x_j - \frac{1}{2\tau_L}\,\sigma_T^2\left(\sum_{i=1}^{N}|b_i + x_i| - 1\right)^2 \tag{14}$$

$$\sigma_T^2 = \sum_{i=1}^{N}\sum_{j=1}^{N}(b_i + x_i)\, q_{ij}\, (b_j + x_j) \tag{15}$$

Now it is possible to chose an initial estimate of σ_T^2 and to use this as a constant to maximize the utility function in equation set 14 and 15. This maximization provided estimates of the x_i's, which they used to compute a new estimate of σ_T^2 using the second equation in equation set 14 and 15. With the new estimate of σ_T^2, they repeated the optimization to find new estimates of the x_i's. They repeated this iteration until successive estimates of σ_T^2 differed by a *de minimis* amount *(Jacobs and Levy, 2013)*.

Mean-Variance Optimization typically operates under two extreme assumptions: either zero leverage tolerance (i.e., long-only portfolios) or unlimited leverage tolerance. In the absence of a leverage constraint, this optimization framework can yield highly leveraged portfolios, which may not be suitable for all investors. When a leverage constraint is imposed, the resulting mean-variance optimal portfolio may align with the preferences of a leverage-averse investor, but only by coincidence. To accurately identify the optimal portfolio for an investor who is averse to leverage, two approaches can be employed. Both methods ultimately lead to the same optimal portfolio:

Generate a range of mean-variance optimal portfolios under varying leverage constraints and evaluate them using the investor's specific mean-variance-leverage utility function. Apply mean-variance-leverage optimization directly, which incorporates both volatility and leverage aversion into the decision-making process.

It is essential to recognize that both volatility tolerance and leverage tolerance play crucial roles in portfolio selection. Just as investors are willing to forgo some expected return to mitigate volatility risk, they may also sacrifice expected return to reduce exposure to leverage risk. Mean-variance-leverage optimization addresses this trade-off by balancing expected return against both sources of risk. Consequently, leverage aversion can significantly influence an investor's portfolio choices. The authors concluded that the MVL model allowed investors to consider their tolerance to volatility but also their leverage tolerance and were able to address one of the shortcomings of the MV model *(Appendix 1.1.3).*

3.2.3. Gasser et al. – Social Responsibility and Portfolio Optimization

Another relevant research, conducted by Gasser et al. *(2017)*, tries to incorporate social responsibility to the first two parameters of the classical framework. When making investment decisions, investors who pay attention to social impact should have a multi-criteria decision model. That is, the authors build such type of decision model by establishing preference parameters for return, risk and social responsibility (Formula 17). When taking these preferences into account, the authors created an *a priori* model to find the optimal portfolio for a specific investor. However, when the model is free of preferences, an *a posteriori* model illustrates all optimal portfolios on the Capital Allocation Plane (CAP) (Figure 3-2), which permits investors to choose portfolio that suits their preferences while considering all three optimization parameters and selecting their asset allocation.

Original Markowitz model: $\max\,(\alpha\mu - \beta\sigma^2)$, (16)

Extended Tri-Criterion model: $\max\,(\alpha\mu + \gamma\theta - \beta\sigma^2)$, (17)

Subject to:

Portfolio weight: $\sum_{i=1}^{n} w_i = 1$, (18)

No short-selling assumed (non-negative weights),

Where:

- μ is the expected return;
- σ^2 is the variance;
- α is the return preference parameter;
- β is the risk preference parameter;
- θ is the social responsibility score;
- γ is the social responsibility preference parameter;
- w_i is the portfolio weights for $i = 1, \ldots, n$.

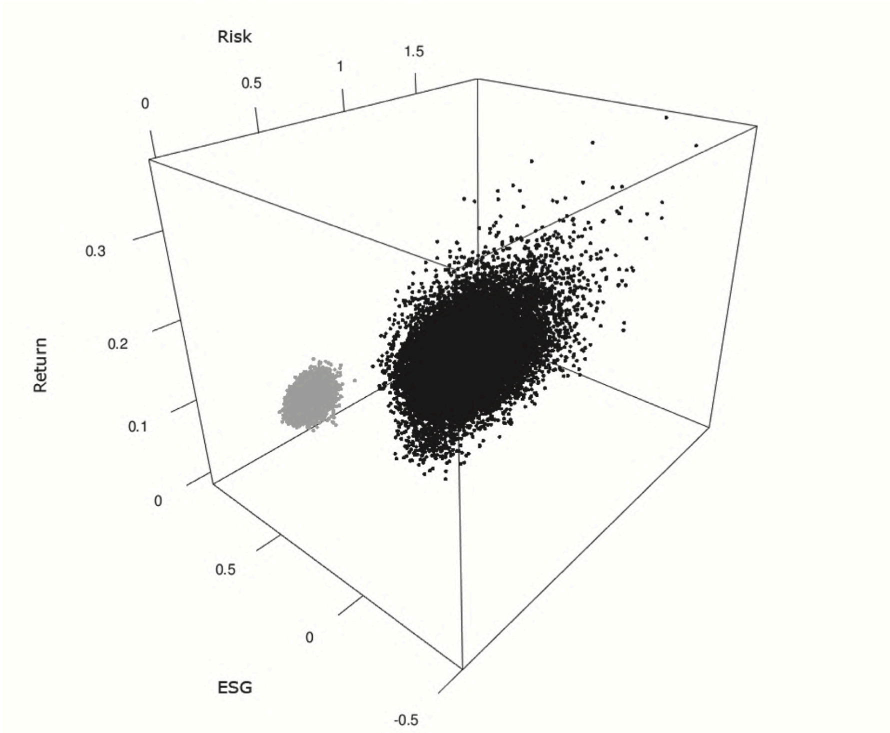

Figure 3-2: Capital Allocation Plane with view on ESG, Return, and Risk (Gasser et al, 2017)

In their paper, Gasser et al. conducted an empirical analysis and applied their a posteriori model to a sample of 6,231 international listed stocks by incorporating the ESG score, retrieved from the ASSET4 ESG database. Results are summarised here under:

- Adding social responsibility as a third criterion produces a three-dimensional efficient frontier.
- There is a statistically significant trade-off: higher Sharpe Ratio (SR) scores result in lower returns but lower risk.
- The CAP provides a visual decision tool for balancing investor preferences across return, risk, and social responsibility.
- Empirical simulations validate the practical feasibility of SR inclusion in portfolio optimization, but with acknowledged trade-offs.

3.2.4. Garcìa et al. – Fuzzy Multicriteria Model

Following the same approach, Garcìa et al. *(2019)* state that the portfolio selection problem could be solved by applying quantitative methods. In their approach, the authors propose a fuzzy multi-obective approach, whose purpose was to optimize the expected return, the expected ESG score, and the downside risk of a given portfolio, subject to real-world constraints such as budget, floor-ceilling and cardinality. García et al. (*2019*) assessed the portfolio's performance and ESG score using the credibility mean. To measure downside risk, they employed the credibility mean-absolute semi-deviation (MASD), arguing that this approach was more suitable than variance since investors tend to be more sensitive to losses than to gains. The authors employ L-R power fuzzy numbers to handle the uncertainty in asset returns and ESG scores and apply credibilistic theory for risk and return measurement. The Sortino ratio (instead of Sharpe ratio to identify the efficient frontier) is adapted into a credibilistic environment, and the Non-dominated Sorting Genetic Algorithm II (NSGA-II) is used to solve the multi-criteria optimization problem.

Their empirical analysis was based on data from Bloomberg's ESG database, combined with information from the Dow Jones Industrial Average (DJIA). We will not go deeper into technical analyses of this method as it is out of scope and typically follows the same approach than Metaxiotis *(cf. Section 3.2.1)*.

Results of this analysis showed that the socially responsible portfolio offered a slightly higher return and ESG score for a higher Value at Risk. With their study, they demonstrated that:

- The fuzzy multi-criteria approach provides superior portfolio selections by integrating ESG criteria with traditional financial objectives.
- The credibilistic Sortino ratio effectively selects optimal portfolios considering downside risk in an uncertain environment.
- The NSGA-II algorithm enables efficient multi-objective optimization, balancing return, risk, and ESG goals.
- The model offers promising results for SR investors and outperforms traditional strategies.
- The use of Bloomberg on the US Dow Jones was relevant for obtaining an ESG score.

3.2.5. Pedersen et al. ESG-Efficient Frontier and ESG Adjusted CAPM

This paper introduces a theoretical and empirical framework for responsible investing by integrating ESG scores into portfolio optimization. It defines the ESG-efficient frontier, which illustrates the maximum Sharpe ratio attainable for each level of ESG. The study models three types of investors (Unaware, Aware, and Motivated regarding ESG), and derives an ESG-adjusted CAPM, showing how ESG affects expected returns, asset prices, and portfolio selection. Empirically, the authors build ESG-efficient frontiers using large datasets and evaluate the costs and benefits of ESG investing, offering a four-fund separation theorem for optimal portfolio construction.

Pedersen et al. *(2021)* propose a novel portfolio theory that incorporates ESG factors directly into the investment process. Their model demonstrates that ESG scores play two critical roles:

- Informational Role: ESG scores provide signals about firm fundamentals, such as profitability and risk.
- Preference Role: ESG scores reflect investor preferences for socially responsible firms.

They introduce the concept of the ESG-efficient frontier, an extension of the traditional mean-variance frontier, which shows the highest attainable Sharpe ratio for each level of ESG score (Figure 3-3). Portfolios on this frontier achieve an optimal balance between return, risk, and ESG performance. They identify four-fund separation in optimal portfolios, involving combinations of the risk-free asset, tangency portfolio, minimum variance portfolio, and the ESG-tangency portfolio. Pedersen et al. *(2021)* extend the Capital Asset Pricing Model (CAPM) to incorporate ESG considerations, resulting in an ESG-adjusted CAPM. This model illustrates how equilibrium asset prices and expected returns adjust depending on the prevalence of different investor types:

- Type-U investors: Ignore ESG scores and focus solely on traditional risk-return trade-offs.
- Type-A investors: Are aware of ESG scores and integrate them into risk-return assessments.
- Type-M investors: Have explicit preferences for high ESG scores and accept lower expected returns for socially responsible portfolios.

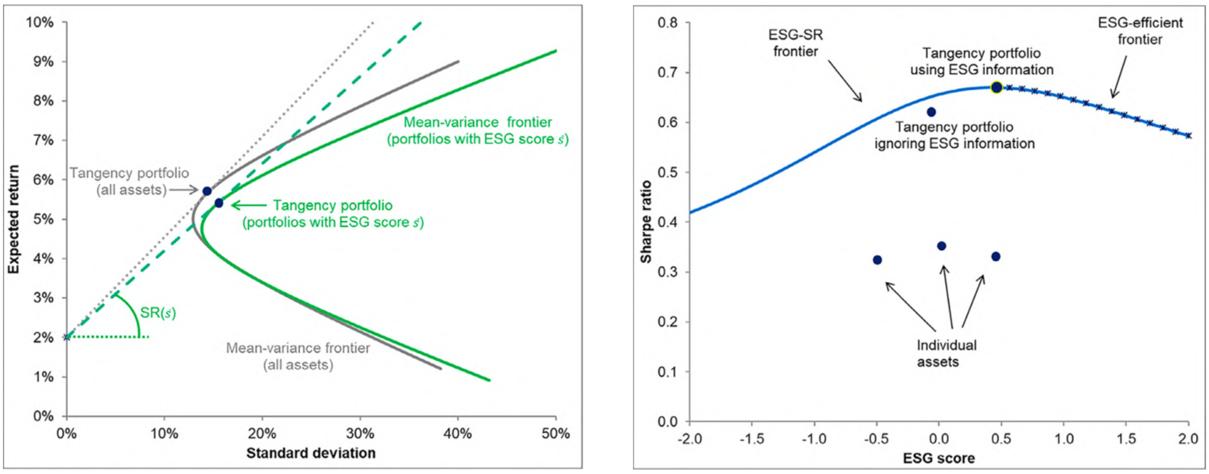

Figure 3-3: Mean-Variance frontier for all assets and portfolios and ESG efficient frontier

The equilibrium outcomes vary based on the composition of these investor types. When type-M investors dominate, high-ESG stocks can trade at premiums, resulting in lower expected returns due to heightened demand. However, when ESG scores predict higher future profitability and investor awareness is limited, high-ESG stocks can offer superior returns (Figure 3-3).

Pedersen et al. empirically estimate the ESG-efficient frontier using large datasets, including carbon intensity measures, sin stock indicators, governance quality, and MSCI ESG ratings. Their findings reveal: Governance (G) measures are strong predictors of future profitability and stock returns, suggesting an informational advantage for ESG-aware investors; Environmental (E) measures, such as carbon intensity, have weaker predictive power for returns, but investors can still construct low-carbon portfolios with minimal Sharpe ratio sacrifice; Social (S) measures show mixed results. Non-sin stocks often underperform due to high investor demand and lower profitability. Their empirical results confirm the theoretical predictions that ESG preferences and information jointly shape asset prices and investor portfolios. They also highlight the importance of ESG dimensions in explaining cross-sectional return differences. Pedersen et al.'s work contributes to both theoretical and empirical ESG investing literature:

- **Theoretical contributions**: They formalize investor preferences for ESG and derive an ESG-SR frontier that guides optimal portfolio choice, offering a more nuanced view than previous segmentation models.
- **Empirical contributions**: Their framework reconciles contradictory empirical findings on ESG performance. It explains why some ESG strategies outperform and others underperform

4 Chapter – Empirical Work

As of now, the thesis will focus on the empirical work provided within the research process. Thus, this section is divided into 3 main chapters.

We will discuss various elements of the chosen methodology to understand what the objectives and the desired outcomes are. We will also develop the research question and the hypotheses that will serve us during the elaboration of the following chapters. Then, it is important to discuss the structure of the research, data collection tools and methods, and data interpretation logics.

We will elaborate on the optimization algorithm and identify the right methodology to construct the desired outcomes. We will give explanation on the code used to create the program prototype, its use, and its purpose. We will understand how it operates and provides outputs. Then, we will provide a comparison between the data observed by using the classical theory and the data observed when using the optimization model in a 4D seeting. We will also build two optimization models – *a posteriori* and *a priori* – and compare them.

Thirdly, we will discuss limits of the analysis and further improvements. More specifically, we will go through the different input variables and try to understand the differences between the classical theory and the optimization model. Also, we will have a closer look to the results and we will explain how variables are influenced one by another.

4.1. Thesis Typology

The typology of this thesis is analytical because it is heavily based on quantitative methods, using mathematical programming, statistical modeling, and Python-based simulations to optimize portfolios. The thesis is building an efficient frontier model in multiple dimensions – return, risk, ESG, and Market SCR – which involves data analysis, optimization algorithms, and programming.

Moreover, it is theoretically-extending due to its core objective which is to move further in the Modern Portfolio theory. It is also an applied research project, grounded in quantitative finance and portfolio optimization by using real asset classes and their actual data. The thesis also tries to develop a prototype tool to model portfolio allocation in a multidimensional framework. This makes the thesis applied in nature, especially relevant to asset managers or institutional investors. Lastly, it is partially comparative due to the comparison between the classical 2D MPT model with a 4D optimized version. These typologies have strongly influenced the methodology adopted for this thesis.

This chapter discusses the theme of the thesis and the hypotheses from which we started to answer the research question. Further in the section, we will describe the data collection tools and methodology that we were ablet to define through the hypotheses, concepts, and dimensions

4.2. Research Question

The title of the thesis is: *"Extensions to the Modern Portfolio Theory".* The research question for the thesis is thus straightforward but challenging to quantify. Indeed, extending is a very broad term and could imply extending to anything possible. Therefore, by narrowing the scope, the following research question is the most appropriate:

"To what extent can the Modern Portfolio Theory be extended into a multidimensional framework that simultaneously integrates return, risk, ESG performance, and market SCR constraints, while remaining computationally tractable and practically applicable to institutional portfolio management and how does a multidimensional portfolio optimization model incorporating ESG scores and market SCR risk compare to the classical Markowitz model in terms of efficiency and asset allocation outcomes?"

This research question incorporates the term "multidimension", meaning we want to move from a 2D Cartesian model to a multi-dimensional (4D) model. It simultaneously integrates return, risk, ESG performance, and market SCR constraints, which are the higher-dimensional extensions to the original Modern Portfolio Theory. These two factors were chosen for two main reasons. First, market SCR comes out as particularly important today, especially since the implementation of Solvency II regulations *(EIOPA, 2024)*. This metric is an additional risk measure of an asset class as it computes the price relating to an asset's exposure towards several risks. Second, according to Lecomte *(2024)*, CEO of Belfius Asset Management, ESG-sensitive assets seem to have a significant impact on portfolio performance without exceeding the risk required by the investor. Moreover, the European market, today, has approximately $12.8 trillion in ESG assets under management *(Lecomte, 2024).*

A series of sub-questions also helped to better grasps the real meaning of the work. These sub-questions directly steem from the general research question.

Theoretical Foundations – sub-questions
Q1. What are the theoretical underpinnings and limitations of the classical Modern Portfolio Theory?
Q2. How have researchers and practitioners attempted to extend MPT to incorporate sustainability and liquidity dimensions?

For answers to these questions refer to *Chapter 1, 2, and 3.*

Data and Metrics – sub-questions
Q3. How can ESG scores and SCR be quantitatively integrated into a portfolio optimization model?
Q4. What data limitations exist when integrating ESG and SCR scores into portfolio construction?

Model Construction and Feasibility – sub-questions
Q5. What kind of optimization techniques are suitable for a 4D portfolio model?
Q6. Is the proposed multidimensional model computationally tractable for real-world application?

Empirical Comparison and Evaluation – sub-questions
Q7. How does the asset allocation produced by the 4D model differ from that of the traditional 2D model (return-risk)?
Q8. Does the 4D model generate portfolios with superior or more balanced outcomes in terms of ESG and market SCR, without significantly compromising return and risk levels?

4.3. Hypotheses

In answering these sub-questions, several assumptions will be used as a starting point for the reasoning throughout this thesis. When possible, the hypotheses will be either refuted or confirmed.

Theoretical Value
H1: Extending the Modern Portfolio Theory to include ESG and market SCR dimensions offers a more holistic framework for portfolio optimization aligned with institutional investor objectives.

Data Integration
H2: ESG and SCR scores can be effectively integrated into optimization models as quantitative constraints or scoring functions without distorting the optimization landscape.

Allocation Outcomes
H3: The 4D optimized portfolios will exhibit significantly different asset allocations compared to 2D portfolios, particularly favoring assets with higher ESG scores and lower market SCR.

Trade-off Management
H4: Multidimensional optimization does not significantly compromise return or risk performance when incorporating ESG and regulatory capital risk as additional constraints or objectives.

Computational Feasibility
H5: The proposed 4D optimization model remains computationally tractable and suitable for practical use in real-time portfolio construction.

Useful dimensions and concepts for answering those hypotheses have been defined in *Chapter 1, 2, and 3*.

As the integration of ESG criteria into portfolio optimization has already been the subject of numerous research studies, the hypotheses formulated in this thesis are situated within a body of existing literature. Rather than introducing an entirely novel concept, this work seeks to test the impact of integrating ESG and Market SCR criteria into the traditional MPT framework originally proposed by Harry Markowitz.

Consequently, the hypotheses are not intended to be descriptive, in the sense of merely identifying or outlining observed trends or behaviors (e.g., the increasing popularity of ESG investing). Instead, they are explicitly explanatory in nature. They aim to evaluate whether and how the inclusion of ESG and SCR metrics alters the structure and outcomes of optimal portfolios, particularly with respect to the shape of the efficient frontier, the distribution of asset weights, and the trade-offs between return, risk, sustainability, and regulatory capital risk.

In line with the typology proposed by Paquet et al. *(2016)*, explanatory hypotheses are those that test a presumed causal relationship between variables. Here, the central relationship under investigation is between the integration of ESG and SCR dimensions and the resulting changes in portfolio composition and performance, as compared to the traditional 2D return-risk optimization.

Therefore, this research adopts an explanatory approach, supported by quantitative modeling and empirical simulations, to test whether a multidimensional optimization model can offer both theoretical enrichment and practical value for institutional portfolio construction.

5 Chapter

5.1. Objective Function Definition

The definition of the objective function is the cornerstone of any portfolio optimization framework. It formalizes the investor's goal in a mathematical expression and determines how candidate portfolios are evaluated and selected. In classical MPT, this goal is defined in terms of a trade-off between two variables: maximizing expected return and minimizing risk. This approach, while foundational and analytically elegant, is inherently limited in scope. It presumes that the universe of investor preferences and constraints can be sufficiently captured by a two-dimensional trade-off between risk and return. However, in the evolving landscape of asset management, where sustainability, regulatory requirements, and risk diversification play increasingly critical roles, there is a compelling need to extend the objective function beyond this simplistic formulation. The purpose of this thesis is to construct and evaluate a multidimensional extension of the modern portfolio theory, which incorporates two additional dimensions of practical relevance: the Environmental, Social, and Governance score and the Market Solvency Capital Requirement. These metrics are not merely auxiliary variables; they are essential drivers of portfolio choice in today's investment environment. ESG considerations are central to sustainable finance initiatives, while SCR is a regulatory measure of risk exposure imposed by Solvency II for insurers and financial institutions operating under strict capital adequacy requirements.

This thesis addresses this shortcoming by proposing a multidimensional extension of the classical MPT. Specifically, the objective function is expanded to account for four critical dimensions:

- Expected Return (maximize),
- Volatility (minimize),
- ESG Score (maximize),
- Market SCR (minimize).

This shift transforms the optimization problem from a univariate or bivariate problem into a multi-objective decision-making framework. This framework accommodates multiple, and often conflicting, goals: (i) maximizing expected return, (ii) minimizing volatility, (iii) maximizing ESG performance, and (iv) minimizing Market SCR exposure. Each of these objectives reflects a different facet of investor preference and regulatory reality. The challenge lies not only in formulating such a function mathematically, but also in implementing it computationally in a way that allows for meaningful empirical comparison. To solve it, this thesis develops and implements two complementary strategies: a preference-free *a posteriori* approach and a preference-driven *a priori* approach inspired by Gaser et al. *(2013)* as defined in *Section 3.2.3*. Both serve to operationalize the objective function under different assumptions about investor input, decision-making constraints, and use cases.

5.1.1. A Posteriori Approach – Preference-Free, Exploratory Optimization

The *a posteriori* approach refrains from embedding any investor preferences directly into the objective function. Instead, it seeks to identify the entire set of Pareto-efficient portfolios, each representing a unique trade-off among the four target dimensions *(Garcìa et al, 2019)*. This model preserves the multidimensional nature of the optimization problem by not collapsing the objectives into a single scalar value. Rather than searching for a single "optimal" portfolio, it aims to characterize the full efficient frontier in four dimensions: return, volatility, ESG score, and Market SCR.

To operationalize this approach, the thesis implements a combinatorial grid search that constructs every feasible portfolio using a tranche-based enumeration method. Each of the 14 asset classes is allocated at least one tranche of 3.75% (i.e., 1/27th of the total portfolio), and the remaining budget is distributed across assets using bounded integer compositions *(Boniver, 2025)*. This deterministic method produces a complete and replicable set of feasible portfolios (Figure 5-1).

```
def generate_all_portfolios(expected_returns, covariance_matrix, esg_scores,
spread_vector, equity_vector, currency_vector, scr_corr_matrix):
""""
Vectorized full enumeration of portfolios using 3.75% tranches with fast matrix ops.

Returns:
np.ndarray: Array of shape (n_portfolios, 4) with [Return, Volatility, ESG, SCR]
List[np.ndarray]: List of normalized weight vectors (each of size n_assets)
""""
n_assets = len(expected_returns)
n_tranches = 27  # Each tranche is 3.75%, so 27 tranches = 100%
tranche_size = 0.0375

min_tranches = np.full(n_assets, 1, dtype=int)  # 1 tranche = 3.75%
adjusted_total = n_tranches - np.sum(min_tranches)
upper_bounds = np.full(n_assets, n_tranches, dtype=int) - min_tranches  # full flexibility

def bounded_compositions(total, bounds):
if len(bounds) == 1:
if 0 <= total <= bounds[0]:
yield [total]
return
for i in range(min(total, bounds[0]) + 1):
for tail in bounded_compositions(total - i, bounds[1:]):
yield [i] + tail
```

Source: Own Analysis

Figure 5-1: Portion of code for the set-up of the entire set of portfolios from the script **`Portfolio_construction.py`**

Each portfolio is evaluated across all four objectives, and non-dominated solutions are extracted using the principle of Pareto optimality *(Garcìa et al, 2019)*. A portfolio is said to be Pareto-efficient if no other portfolio exists that improves one objective without worsening at least one of the others *(cf. Appendix 1.2.1)*. Formally, the Pareto-efficient set $\mathcal{P}^*$ is defined as:

$$\mathcal{P}^* = \{w \in W \mid \notin w' \in W \text{ such that } f_i(w') \geq f_i(w)\ \forall i \text{ and } f_j(w') > f_j(w) \text{ for some } j\} \quad (19)$$

Where:

- $f_1(w) = \mu(w)$ — Expected return
- $f_2(w) = -\sigma(w)$ — Risk (to be minimized)
- $f_3(w) = ESG(w)$ — ESG performance
- $f_4(w) = -SCR(w)$ — Market SCR (to be minimized)

Each portfolio w is a vector of weights across 14 asset classes. The feasible set W is constrained by common rules (non-negativity, full investment, etc., covered in *Section 5.2*).

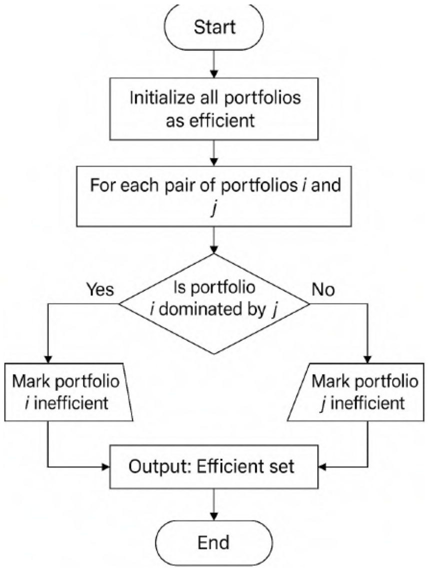

Source: Own analysis

Figure 5-2: Filtering efficient portfolios algorithm according to the Pareto Efficiency methodology

The filtering of efficient portfolios is refined using a lexicographic sorting *(Zhang et al. 2025)* and greedy scan procedure, incorporating small epsilon tolerances *(IBM Corporation, 2022)* along each objective axis (Figure 5-3). These steps are essential to avoid overfitting the Pareto concept to minor numerical noise. Without them, the sheer volume of portfolios, each slightly differing in performance metrics due to rounding or allocation granularity, would result in an overwhelming number of points being classified as "efficient".

These epsilon tolerances can be adapted as required by the investor. They represent the minimum a candidate must overperform another in each dimension to be classified as dominating or not.

Lexicographic sorting is a method used to identify truly non-dominated portfolios when dealing with multiple objectives. After generating the full set of feasible portfolios, lexicographic sorting ranks them sequentially by each objective, according to a strict priority order. For example, the sorting procedure may first select portfolios with the highest return. Among those, it selects the ones with the lowest volatility. Then, from that subset, it chooses the portfolios with the highest ESG score, and finally the ones with the lowest SCR. This hierarchical filtering continues until no further refinement is possible (*cf. Appendix 1.2.2* for a detailed explanation of the lexicograpic sorting procedure).

This technique is crucial in high-dimensional optimization because it avoids falsely labeling nearly-optimal portfolios as efficient. Without lexicographic refinement, over 95% of the generated portfolios might appear Pareto-efficient due to minor numerical differences. The lexicographic ordering prioritizes dominant objectives while the greedy scan removes near-duplicates and portfolios offering marginally different trade-offs. This produces a compact, interpretable, and analytically useful approximation of the 4D efficient frontier.

```python
def filter_efficient_portfolios(portfolios_4d):
    """
    Efficient filtering using lexicographic sorting + greedy scan
    with epsilon tolerance across all metrics.

    Args:
        portfolios_4d (np.ndarray): shape (N, 4), [Return, Volatility, ESG, SCR]

    Returns:
        np.ndarray: efficient portfolios
        np.ndarray: indices in original set
    """
    # Epsilon margins for soft filtering
    eps_return = 0.001
    eps_vol    = 0.001
    eps_esg    = 0.5
    eps_scr    = 0.01

    # Normalize sorting order: maximize return/ESG, minimize vol/SCR
    sort_data = np.copy(portfolios_4d)
    sort_data[:, 1] *= -1  # Volatility
    sort_data[:, 3] *= -1  # SCR

    sort_idx = np.lexsort((sort_data[:, 3], -sort_data[:, 2], sort_data[:, 1], -sort_data[:,
0]))
    sorted_portfolios = portfolios_4d[sort_idx]

    efficient_mask = np.ones(len(sorted_portfolios), dtype=bool)
    current_best = sorted_portfolios[0]

    for i in range(1, len(sorted_portfolios)):
        candidate = sorted_portfolios[i]

        dominated = (
            candidate[0] <= current_best[0] + eps_return and
            candidate[1] >= current_best[1] - eps_vol and
            candidate[2] <= current_best[2] + eps_esg and
            candidate[3] >= current_best[3] - eps_scr and
            (
                candidate[0] < current_best[0] - eps_return or
                candidate[1] > current_best[1] + eps_vol or
                candidate[2] < current_best[2] - eps_esg or
                candidate[3] > current_best[3] + eps_scr
            )
        )
```

Source: Own Analysis

Figure 5-3: Portion of code for the filtering of the entire set of portfolios from the script `frontier.py`

The *a posteriori* approach offers several advantages. It enables the investor to explore structural trade-offs without imposing potentially arbitrary or subjective preferences. For instance, one can investigate whether ESG performance is inversely related to return, or whether portfolios with low SCR cluster in specific regions of the frontier. This method is particularly valuable when preferences are undefined or exploratory analysis is desired. It supports decision-makers by presenting a complete spectrum of efficient outcomes, rather than pre-selecting a single “best” portfolio.

5.1.2. A Priori Approach – Preference-Driven, Scalar Optimization

The *a priori* approach, by contrast, begins with the assumption that investor preferences can be explicitly quantified. Instead of treating all four dimensions as separate and equally valid optimization targets, it aggregates them into a single scalar-valued objective function using weighted coefficients that reflect the relative importance of each criterion.

The general form of the objective function is:

$$\max_{w} S(w) = \alpha \times \mu(w) - \beta \times \sigma(w) + \gamma \times ESG(w) - \delta \times SCR(w) \tag{20}$$

Where:

- $\alpha, \beta, \gamma, \delta \in [0, 1]$ are weights that reflect investor's preferences over return, risk, ESG, and SCR
- The sum of weights may be normalized (e.g., $\alpha + \beta + \gamma + \delta = 1$), though other weighting schemes are also feasible.
- $ESG(w) = \sum w_i \times ESG_i$
- $SCR(w) = \sum w_i \times SCR_i$

Before applying the weighted scoring function in the *a priori* model, all four objectives (return, volatility, ESG score, and Market SCR) are first standardized using Z-score normalization *(Nevil, 2019)*. This step is essential because these metrics differ significantly in scale and dispersion: for example, expected returns typically range between 0.03 and 0.07, whereas ESG scores span 0 to 100, and SCR values may fall within 0.1 to 0.5. Applying weights directly to raw values would therefore distort the relative importance of each dimension, giving undue influence to higher-magnitude variables. Z-score normalization resolves this by transforming each metric into a standardized scale with a mean of zero and a standard deviation of one. The formula used is:

$$Z(x_i) = \frac{x_i - \bar{x}}{\sigma} \tag{21}$$

Where:

- x_i is the value of the metric for a given portfolio,
- $\bar{x}$ is the mean across all portfolios,
- σ is the corresponding standard deviation.

This procedure ensures that the weighting scheme reflects true investor preferences, not differences in measurement units or data dispersion. It also allows for a fair and consistent comparison of portfolio performance across all four dimensions. With this, the objective function could be rewritten as follows:

$$\max_{w} S(w) = \alpha \times Z(\mu(w)) - \beta \times Z(\sigma(w)) + \gamma \times Z(ESG(w)) - \delta \times Z(SCR(w)) \tag{22}$$

Where:

- Z denotes the normalized Z-score of each metric across all portfolios in the dataset.

Once all portfolios have been evaluated across the four dimensions, a single composite score is calculated for each portfolio. After the score is computed, the model simply selects the portfolio with the highest score as the optimal solution (Figure 5-4). The score function transforms the multi-objective decision problem into a single-objective ranking. Once $S(w)$ is computed for all portfolios $w \in W$, the optimal portfolio is selected as:

$$w^* = \arg.\max_{w \in W} S(w) \tag{23}$$

This approach guarantees that w^* is the portfolio that best aligns with the investor's weighted preferences, while respecting all feasibility constraints (e.g., full allocation, non-negativity, tranche granularity). This approach does not involve mathematical optimization in the traditional sense (e.g. solving a system via optimization algorithms), but instead performs an exhaustive ranking of all feasible portfolios and chooses the top-ranked one. This makes the selection process both transparent and deterministic: given the same preferences and portfolio universe, the outcome will always be the same. It also ensures that the selected portfolio respects all predefined constraints (such as weight bounds and full allocation), since it is chosen from a set that was already constructed to be fully feasible.

```
# === Compute scores based on investor profile ===
def compute_scores(df_norm, weights):
    df = df_norm.copy()
    df['score'] = (
        weights['return']     * df['Return_norm'] -
        weights['volatility'] * df['Volatility_norm'] +
        weights['esg']        * df['ESG_norm'] -
        weights['scr']        * df['SCR_norm']
    )
    return df

# === Example investor profiles ===
investor_profiles = {
    "balanced":         {"return": 0.25, "volatility": 0.25, "esg": 0.25, "scr": 0.25},
    "esg_focused":      {"return": 0.2,  "volatility": 0.2,  "esg": 0.5,  "scr": 0.1},
    "scr_conservative": {"return": 0.2,  "volatility": 0.2,  "esg": 0.1,  "scr": 0.5},
    "return_max":       {"return": 0.6,  "volatility": 0.2,  "esg": 0.1,  "scr": 0.1},
    "low_risk":         {"return": 0.2,  "volatility": 0.5,  "esg": 0.2,  "scr": 0.1},
    "pure_return":      {"return": 1.0,  "volatility": 0.0,  "esg": 0.0,  "scr": 0.0},
    "pure_volatility":  {"return": 0.0,  "volatility": 1.0,  "esg": 0.0,  "scr": 0.0},
    "pure_esg":         {"return": 0.0,  "volatility": 0.0,  "esg": 1.0,  "scr": 0.0},
    "pure_scr":         {"return": 0.0,  "volatility": 0.0,  "esg": 0.0,  "scr": 1.0}
}
```

Source: Own Analysis

Figure 5-4: Portion of code for computing the Z-Score of each portfolio from the script **`Optimizer_a_priori.py`**

By adjusting these weights, different investor profiles can be modeled:

- A conservative investor may heavily penalize volatility and SCR.
- A sustainable investor may assign a high value to ESG performance.
- A return-seeking investor may overweight expected return while accepting higher risk and lower ESG scores.

This model is computationally efficient and aligns with real-world implementation, especially for institutional investors operating under clear mandates. However, it assumes that preferences are both known and stable, which is not always the case.

5.1.3. Comparative Insights – Complementary Strengths of A Posteriori and A Priori Approaches

The *a posteriori* and *a priori* models offer two distinct paths through the problem of multidimensional portfolio selection. They are not mutually exclusive but rather complementary tools, each suited to different stages of the investment decision-making process (Table 5-1 and Table 5-2)

The *a posteriori* model is:

- Exploratory: It reveals the structure of possible trade-offs without commitment.
- Preference-free: It does not require explicit weighting, avoiding premature subjectivity.
- Rich in insight: It produces an entire surface of efficient outcomes, ideal for visualization and policy exploration.

The *a priori* model is:

- Directive: It outputs a single, immediately actionable portfolio.
- Aligned with mandates: Useful when working within regulatory or institutional constraints.
- Flexible: Can easily incorporate changes in investor attitudes or policy adjustments.

Table 5-1: Features comparison for the a posteriori and the a priori model

Feature	A Posteriori Model	A Priori Model
Preference Input	Not required	Explicit weights required
Output	Set of efficient portfolios (frontier)	One optimal portfolio
Use Case	Research, policy, sensitivity analysis	Real-world asset allocation
Computation	Very high (full enumeration + Pareto filtering)	Very high (full enumeration + scalar scoring)
Transparency	High	High
Customization	Ex-post selection	Ex-ante preference modeling

From a practical investment workflow, these models can be sequenced: an investor may begin with an *a posteriori* analysis to explore trade-offs, then calibrate weights for an *a priori* optimization that reflects informed preferences. This dual-track framework enhances both the analytical rigor and practical applicability of the multidimensional model developed in this thesis.

Table 5-2: Summary of the model used

Model	Optimization Type	Output	Method
Classical (2D)	Return & risk only	One portfolio	SLSQP (Sharpe maximization)
A Posteriori (4D)	Preference-free	Set of non-dominated portfolios	Full enumeration + Pareto filtering
A Priori (4D)	Preference-based	One optimal portfolio	Full enumeration + scalar score maximization

5.2. Constraints and Multi-Criteria Approach

In portfolio optimization, constraints play a foundational role by defining the boundaries of the feasible investment universe. They ensure that the resulting portfolios are not only theoretically optimal but also economically viable, operationally implementable, and consistent with real-world regulatory and institutional frameworks. Constraints formalize the admissibility criteria by which portfolios are constructed, thereby filtering out solutions that would be irrelevant or infeasible in practice.

In this thesis, a unified set of baseline constraints is applied across all optimization models, including the classical 2D mean-variance framework, the *a posteriori* simulation-based approach, and the *a priori* preference-driven scoring model. The use of common constraints ensures methodological consistency, realism, and comparability between models, which is essential for evaluating the trade-offs and advantages of each approach in a controlled and coherent manner.

Extending the classical MPT from two to four objectives fundamentally alters the nature of the optimization problem. While the classical formulation is based on a scalar trade-off, maximizing expected return for a given level of volatility or minimizing volatility for a given return, the inclusion of ESG performance and Market Solvency Capital Requirement introduces two additional, structurally distinct and operationally relevant criteria. The optimization problem no longer resides in a two-dimensional utility space, but in a four-dimensional objective space defined as:

$$\max_{w \in W} \big(f_1(w), -f_2(w), f_3(w), -f_4(w)\big) \quad (24)$$

Where:

- $f_1(w) = \mu(w) - Expected\ return$
- $f_2(w) = \sigma(w) - Risk\ (to\ be\ minimized)$
- $f_3(w) = ESG(w) - ESG\ score$
- $f_4(w) = SCR(w) - Market\ Solvency\ Capital\ Requirements$

In this expanded framework, no single portfolio is likely to dominate all others across all dimensions. The solution must therefore shift from finding the optimal portfolio to navigating a set of non-dominated portfolios, each offering a different trade-off. This is the foundational principle behind multi-objective optimization, and it underpins the dual-model architecture developed in this thesis.

The *A Posteriori* Model: Exploration of the Efficient Set

The *a posteriori* model is designed to explore the entire efficient frontier. It evaluates millions of portfolios generated through a constrained, tranche-based enumeration process and filters them using Pareto optimality with epsilon tolerances, followed by lexicographic sorting to manage near-indistinguishable portfolios. Importantly, this model does not aggregate the objectives; it maintains them as independent optimization targets, and thus avoids imposing trade-offs (*cf. Section 5.1.1* for a full example of the code snippet used).

This model is especially powerful in contexts where the investor's preferences are either unknown, evolving, or subject to external negotiation (e.g., ESG committees, regulatory stress-testing, scenario analysis). It enables users to visualize the geometry of trade-offs in the feasible space and perform post hoc selection based on qualitative overlays.

The *A Priori* Model: Preference-Driven Selection

In contrast, the *a priori* model embraces the scalarization principle of multi-objective optimization. Investor preferences are encoded as normalized weights over the four dimensions and combined into a single composite score that defines the objective function. This collapses the multi-objective problem into a scalar-valued optimization problem and yields a single, implementable portfolio (*cf. Section 5.1.2* for a full example of the code snippet used).

This structure is particularly aligned with practical portfolio management, where mandates often require a fixed allocation rather than a spectrum. It is computationally efficient and interpretable, but also sensitive to weight specification, which means that robust calibration and sensitivity analysis are essential.

5.2.1. Common Constraints Applied to All Models

To ensure realism, consistency, and institutional feasibility, all portfolio models developed in this thesis, whether classical, *a posteriori*, or *a priori*, are subject to a unified set of core constraints. These constraints are designed to enforce diversification, avoid uninvestible weight structures, and create a bounded and interpretable solution space. Rather than relying on continuous random sampling, which would yield an uncountably infinite number of portfolios and highly granular (yet practically meaningless) weights, this thesis adopts a tranche-based approach that discretizes the weight space into predefined increments of investible proportions.

Tranche-Based Portfolio Enumeration

Portfolios are generated by allocating capital in discrete units (tranches) of fixed size $\Delta = 3.75\%$. Each asset is assigned an integer number of tranches such that the total number of tranches across all assets equals a fixed budget (e.g., 27 tranches × 3.75% = 101.25%). This ensures portfolios are composed of investible allocations like 3.75%, 7.5%, or 11.25%, which align with real-world asset mandates and internal portfolio block constraints.

Let $w_i \ in \ \{1, 2, \dots, T_i{}^{max}\}$ denote the number of tranches assigned to asset i, and let Δ be the tranche size. Then, the raw portfolio weights before normalization are:

$$w_i{}^{raw} = \Delta \cdot t_i \ with \sum_{i=1}^{n} w_i{}^{raw} > 1 \tag{25}$$

Where:

- t_i is the number of tranches assigned to asset i.

To ensure full capital allocation without leverage, we normalize all raw weights across the portfolio:

$$w_i = \frac{w_i^{raw}}{\sum_{j=1}^{n} w_j^{raw}} \ so\ that \sum_{i=1}^{n} w_i = 1 \tag{26}$$

This post-normalization guarantees that all final portfolio weights are strictly positive, sum to 100%, and remain interpretable.

Full Investment Constraint – Capital Completeness

This constraint ensures that 100% of available capital is allocated across the portfolio:

$$\sum_{i=1}^{n} w_i = 1 \tag{27}$$

It is enforced directly through the normalization process applied after tranche assignment. From a financial perspective, this constraint guarantees that no capital is left uninvested (i.e., no idle cash), nor is capital borrowed to generate leverage. For institutional investors such as insurance companies or pension funds this requirement aligns the optimization framework with actual portfolio construction rules.

From a computational standpoint, enforcing the sum-to-one constraint ensures all portfolios lie on the unit simplex, simplifying comparison across models and making portfolio returns and risk figures directly interpretable *(Winkel et al, 2023)*.

Minimum Allocation Constraint – No Asset Omission

A key modeling choice in this thesis is to ensure that every asset class is represented in all portfolios. This is enforced by assigning at least one tranche to every asset before normalization:

$$w_i^{raw} \geq \Delta \Longrightarrow \ w_i \geq \frac{\Delta}{\sum_j w_j^{raw}} \geq 0 \tag{28}$$

This guarantees that no asset receives a zero weight in the final portfolio. This modeling choice is significant for several reasons:

- **Diversification**: It prevents extreme corner solutions in which capital is entirely concentrated in a few low-volatility or high-return assets.
- **Policy alignment**: Many institutional mandates require exposure to a full strategic asset mix, with minimum floor allocations per asset class.
- **Analytical stability**: It ensures that all four performance dimensions (Return, Volatility, ESG, SCR) reflect inputs from the entire asset universe, avoiding distortions caused by exclusion.

The constraint thereby improves realism and interpretability while still allowing the optimization process to explore a rich space of feasible allocations.

Non-Negativity Constraint – Long-Only Investing

Short-selling is not permitted in any of the models developed in this thesis:

$$w_i \geq 0 \ \forall i \tag{29}$$

This constraint reflects the long-only nature of most institutional portfolios. It is especially relevant in the context of ESG and SCR integration:

- **For ESG**: Negative weights would imply inverse exposure to sustainability metrics, which is economically and ethically inconsistent with sustainable investing goals.
- **For SCR**: Negative exposures could artificially reduce total capital requirements, violating Solvency II logic.
- **For allocation transparency**: Non-negativity produces cleaner, more intuitive outputs where weights correspond to real, positive asset exposures.

Moreover, combined with the minimum tranche constraint, this ensures that all portfolios are strictly positive and fully diversified.

Optional Performance Filter – Post-Processing for Relevance

Although not a constraint in the optimization engine itself, an optional return filter is applied during post-processing:

$$\mu(w) = \sum_{i=1}^{n} w_i \cdot \mu_i \geq 0 \tag{30}$$

This step removes portfolios with negative expected return, which, while mathematically feasible, are rarely of interest in practical investment analysis. This filter:

- Enhances the clarity of visualizations such as the efficient frontier.
- Avoids statistical distortions in descriptive analyses.
- Ensures that only economically meaningful portfolios are retained for downstream comparison and profile selection.

It is important to emphasize that this step does not affect the generation or optimization process, and is used solely to clean the dataset for final reporting.

5.2.2. A Posteriori Model – Multi-Criteria without Preferences

In the *a posteriori* framework, constraints serve a structural role: they define the admissible space of feasible portfolios but do not encode any investor-specific preferences. Unlike the *a priori* model, which embeds priorities via a scalar objective function, the *a posteriori* method retains the full multidimensionality of the problem, treating expected return, volatility, ESG score, and Market SCR as distinct and equally important objectives.

Portfolio generation is conducted via a tranche-based enumeration algorithm *(cf. Section 5.2.1)*, ensuring that:

- All weights are strictly positive (minimum allocation per asset),
- The portfolio is fully invested (weights sum to 100%), and
- All weights are realistically discretized in units of 3.75% to reflect real-world investibility.

These constraints do not affect the optimization logic directly. Instead, the model evaluates all admissible portfolios using a two-phase selection mechanism:

Pareto Filtering with Epsilon Tolerance

To account for numerical noise and practical indistinguishability between portfolios, the filtering process allows small tolerances ε on each dimension. For instance, a portfolio is considered dominated only if another performs strictly better in at least one dimension and no worse within tolerances in the others (*cf. Section 5.1.1* for a full example of the code snippet used). This ensures the resulting frontier is robust to near-duplicates and avoids artificial inflation of the Pareto set.

Lexicographic Sorting

Among non-dominated portfolios, a sorting heuristic is applied that hierarchically prioritizes the four dimensions (maximize return → minimize volatility → maximize ESG → minimize SCR). This mechanism imposes no preference weights but is crucial to resolve equivalence across portfolios that fall within ε-tolerances, resulting in a parsimonious and interpretable efficient frontier *(cf. Appendix 1.2.2)*.

The outcome is a preference-free set of optimal portfolios, each representing a distinct trade-off across the four dimensions. The role of constraints here is to bound the simulation space, enforcing realism and consistency, but not to guide selection. Final choice remains with the investor, who may interpret or select among the frontier portfolios based on qualitative, strategic, or regulatory objectives not captured by the model itself.

5.2.3. A Priori Model – Multi-Criteria with Embedded Preferences

The *a priori* model adopts a fundamentally different approach by embedding investor preferences directly into the portfolio selection process. Rather than constructing an entire Pareto frontier of non-dominated portfolios, the model seeks to identify a single optimal portfolio per investor profile by maximizing a scalarized score that aggregates all four dimensions. Formally, each portfolio w is evaluated based on a weighted scoring function as seen in Equation 20. $\alpha, \beta, \gamma, \delta$ in [0, 1] are the user-defined weights representing the relative importance of each dimension. The negative signs for volatility and SCR reflect their undesirable nature. Prior to scoring, each objective is standardized using Z-score normalization to ensure comparability of scales across metrics.

Instead of solving this scalar function through numerical optimization, the model evaluates the score for every portfolio in a precomputed, finite set of feasible portfolios. These portfolios are generated once via a tranche-based simulation that already satisfies core investment constraints *(cf Section 5.2.1)*. The portfolio with the highest score is selected for each investor profile.

The same foundational constraints applied in the *a posteriori* model hold here:

- **Full investment**: All capital must be allocated.
- **Long-only**: No short positions are permitted.
- **Minimum allocation per asset**: Every asset receives a non-zero weight (≥ one tranche).

In the *a priori* context, these constraints also interact with the preference structure. For example, a profile that heavily penalizes SCR might favor a portfolio with low regulatory capital charges, but if that portfolio violates a return threshold or falls outside the investible weight space, it will be excluded. As a result, the final selection is shaped not only by investor preferences but also by the geometry and diversity of the feasible region.

This approach is computationally efficient, transparent, and aligned with real-world practices, especially within institutional asset management where mandates are well-defined. It offers clear traceability from preferences to portfolio selection. However, it remains sensitive to input weights: small variations in α, β, γ, or δ may lead to markedly different portfolio compositions, emphasizing the need for careful calibration.

5.3. Dimensional Analysis in Portfolio Theory

The selection of ESG and SCR as optimization objectives is not arbitrary. Both are now embedded in the strategic priorities of asset managers, insurers, and pension funds. ESG integration reflects the growing requirement to align investments with long-term sustainability goals and stakeholder expectations. It has moved beyond exclusion-based screening *(La Torre, 2025)* or post-hoc reporting and increasingly informs strategic asset allocation. From a modeling perspective, ESG offers a qualitatively different dimension: it is neither purely financial nor reducible to traditional measures of risk or return.

In this thesis, ESG scores are computed at the asset class level using a weighted aggregation of issuer-level ESG ratings sourced from Refinitiv and Morningstar, based on data availability *(cf. Section 6.3* and *Section 6.4)*. These scores are treated as continuous optimization targets, enabling portfolios to be selected based on their overall ESG footprint, rather than through exclusion lists or hard thresholds.

The second added dimension, Market SCR, originates from regulatory risk assessment frameworks such as Solvency II. It quantifies the capital buffer required to protect against asset-specific market shocks, and thus captures the regulatory cost associated with holding risky assets. Unlike volatility, which measures market-implied dispersion, SCR reflects risk through the lens of regulatory prudence. The two are positively related in many cases, but not always tightly correlated, especially due to the use of stress-based approximations and asset aggregation rules in the regulatory methodology. In this thesis, SCR values are calculated using a weighted approach based on predefined shock factors and a regulatory correlation matrix, as defined by the European Insurance and Occupational Pensions Authority (EIOPA). The overall SCR for a portfolio is obtained using the square root of the sum of pairwise correlated SCR contributions across asset classes *(cf. Section 5.4.2)*. This approach, implemented in the Python module `scr_calculator.py`, captures diversification effects and allows capital efficiency to be modeled directly. Figure 5-5 provides insights into the code used to compute the SCR score.

```python
import numpy as np

def   compute_scr_diversified(weights,   spread_vector,   equity_vector,   currency_vector,
correlation_matrix):
    """
    Optimized computation of Market SCR (after diversification) for a given portfolio.

    Args:
        weights (np.ndarray): Portfolio weights, shape (n_assets,)
        spread_vector (np.ndarray): Spread risk loads per asset, shape (n_assets,)
        equity_vector (np.ndarray): Equity risk loads per asset, shape (n_assets,)
        currency_vector (np.ndarray): Currency risk loads per asset, shape (n_assets,)
        correlation_matrix (np.ndarray): 3x3 correlation matrix between stress types

    Returns:
        float: Diversified SCR score for the portfolio
    """
    # Combine stress vectors into a matrix: shape (3, n_assets)
    stress_matrix = np.vstack((spread_vector, equity_vector, currency_vector))

    # Compute total exposure to each stress type: shape (3,)
    S = stress_matrix @ weights

    # Compute the diversified SCR using the regulatory correlation matrix
    return np.sqrt(S @ correlation_matrix @ S)
```

Source: Own Analysis

Figure 5-5: Portion of code for computing the SCR Score of each portfolio from the script scr_calculator.py

Fom an optimization standpoint, the inclusion of ESG and SCR expands the efficient frontier from a two-dimensional curve into a complex Pareto surface, which represents portfolios that are non-dominated across all four criteria. This structural expansion introduces not only computational challenges, but also richer decision possibilities. The inclusion of ESG and Market SCR transforms the efficient frontier from a two-dimensional curve into a multi-dimensional surface. This expansion increases the potential for investor differentiation, as portfolios can now be optimized across regulatory capital efficiency and sustainability performance in addition to traditional financial objectives.

In conclusion, the dimensional expansion to include ESG and Market SCR transforms the optimization process from a narrowly defined financial task to a more comprehensive and adaptable framework. Each added dimension introduces a non-redundant layer of decision complexity and provides additional levers for aligning portfolios with institutional values and external constraints.

5.4. Construction of the 4D Efficient Frontier Model

Constructing an efficient frontier in four dimensions requires a significant methodological departure from the classical two-dimensional MPT framework. Unlike the traditional risk–return trade-off curve, the 4D efficient frontier exists in a multidimensional space that reflects the complexity of modern investment mandates. This section outlines the simulation, evaluation, and filtering procedures used to construct the 4D frontier in this thesis and summarises all information collected throught the section discussed before.

5.4.1. Portfolio Generation: Tranche-Based Enumeration

The first step is the generation of a large, yet finite, set of feasible portfolios. To achieve this, the thesis adopts a tranche-based enumeration method rather than continuous random sampling. Each portfolio is constructed by allocating investment capital in fixed increments of 3.75% (one "tranche") across the 14 asset classes. Every asset is required to receive at least one tranche, enforcing diversification and preventing asset omission. This approach results in a discrete grid of possible portfolio combinations, which is both computationally manageable and aligned with real-world investment practices where portfolios are often implemented in blocks.

After generating all valid combinations, a normalization step is applied to ensure the sum of weights equals 100%. This ensures that the portfolios are fully invested, long-only, and composed of strictly positive, interpretable weights.

5.4.2. Portfolio Evaluation: Computing 4D Metrics

For each generated portfolio, the four objective values are calculated using preloaded asset-level data. This includes:

Return and Volatility

Let $\mu \in \mathbb{R}^n$ be the expected return vector and $\sigma \in \mathbb{R}^{n \times n}$ the covariance matrix. Then for each portfolio weight vector w:

$$Return(w) = \mu^T w$$

$$Volatility(w) = \sqrt{w^T \sum w}$$

ESG Score and Market SCR

ESG scores are computed as weighted averages:

$$ESG(w) = \sum_{i=1}^{n} w_i \cdot ESG_i$$

Market SCR is computed using the Solvency II aggregation formula over three risk types (spread, equity, currency):

$$SCR(w) = \sqrt{S(w)^T \cdot \rho \cdot S(w)} \; where \; S(W) = \begin{bmatrix} w^T \cdot Spread\ load \\ w^T \cdot Equity\ load \\ w^T \cdot Currency\ load \end{bmatrix}$$

Where:

- ρ is the 3x3 correlation matrix between risk types (Figure 5-6)

	Spread	Equity	Currency
Spread	1	0.75	0.25
Equity	0.75	1	0.25
Currency	0.25	0.25	1

Figure 5-6: Correlation matrix between risk types

The code snippet used to compute the SCR score looks as follows:

```
    Returns:
        float: Diversified SCR score for the portfolio
    """
    # Combine stress vectors into a matrix: shape (3, n_assets)
    stress_matrix = np.vstack((spread_vector, equity_vector, currency_vector))

    # Compute total exposure to each stress type: shape (3,)
    S = stress_matrix @ weights

    # Compute the diversified SCR using the regulatory correlation matrix
    return np.sqrt(S @ correlation_matrix @ S)
```

Source: Own Analysis

Figure 5-7: Portion of code for computing the SCR Score of each portfolio from the script `scr_calculator.py`

All computations are vectorized to handle millions of portfolios efficiently. The final result is a matrix where each row corresponds to a portfolio and its associated four performance metrics.

5.4.3. Filtering: Extracting the Efficient Frontier

Once the full portfolio space is simulated and evaluated, the next step is to extract the non-dominated portfolios, those that form the efficient frontier. This is done through a Pareto filtering process, which identifies portfolios for which no other portfolio performs better across all four dimensions.

Given the scale of the data, filtering is implemented using an optimized version of lexicographic sorting in `frontier.py`. This method orders portfolios by their objectives using a predefined hierarchy (e.g., maximize return and ESG, minimize volatility and SCR), and applies greedy scanning to eliminate dominated solutions. Without this sorting strategy, more than 95% of the portfolios would be labeled as Pareto-efficient due to numerical ties and floating-point noise. The introduction of epsilon tolerances (e.g., 0.001 for return, 0.01 for SCR) ensures that only portfolios showing meaningful trade-off improvements are retained. The output is a curated set of portfolios that reflect the most favorable combinations of the four competing criteria. Each point on this frontier represents a different balance among return, risk, sustainability, and regulatory capital efficiency.

5.4.4. Design Rationale and Flexibility

Importantly, no investor preferences are applied during frontier construction. This design choice ensures that the 4D frontier remains preference-free, enabling it to be used flexibly in both *a posteriori* and *a priori* frameworks. In the *a posteriori* model, it allows the user to visually or analytically explore the shape of trade-offs across all objectives. In the *a priori* model, it forms the dataset on which scalarized scoring functions (based on investor weights) can be applied to identify the “best” portfolio for a given profile. The modularity of this architecture ensures that preferences can be added or modified without regenerating the entire frontier, saving significant computational time and improving interpretability.

5.4.5. Implementation and Output

The 4D frontier construction process is centralized in the file `Main_file.py`, which coordinates the portfolio generation (`Portfolio_construction.py`), metric evaluation (`scr_calculator.py`), and filtering (`frontier.py`). The resulting data is exported to Excel and CSV files, and is then used downstream for visualization (`Visualisation.py`) and investor-specific optimization (`Optimizer_a_priori.py`)[1]

This methodological pipeline ensures robustness, transparency, and adaptability to a wide range of research or professional contexts.

5.4.6. Tools and Techniques for Optimization

The successful implementation of multidimensional portfolio optimization relies not only on conceptual model design but also on an efficient, modular, and reproducible computational infrastructure. This section outlines the key tools and techniques used to support the construction of the 4D efficient frontier, as well as the subsequent optimization tasks carried out in both the *a posteriori* and *a priori* models. The decisions made at the level of programming language, algorithmic structure, and software architecture were not merely technical conveniences, they were driven by the need to simulate, evaluate, and filter millions of portfolios with numerical precision and computational efficiency.

Programming Language and Environment

The entirety of the optimization model was developed in Python, an open-source language increasingly adopted in quantitative finance and data science due to its balance of flexibility, scalability, and extensive ecosystem of scientific libraries. Python was chosen over alternative environments such as R or MATLAB for its superior performance in large-scale simulations, broader optimization tools, and seamless integration with data handling packages. The scripts were developed and executed in a structured folder architecture, as detailed in *Appendix 1.2.3*, ensuring modularity and clarity across all components of the model.

Execution was performed in a local computing environment using Python 3.11, with performance profiling conducted on a machine equipped with an Apple M1 chip and 16 GB of RAM. Computational time was optimized using vectorized operations and memory-efficient data structures, particularly during portfolio simulation and Pareto filtering, where large matrix operations are central.

[1] Full architecture of the program is explained in *Section 5.4.6*.

Key Libraries and Packages

The modeling framework developed in this thesis is built entirely in Python and relies on a modular architecture composed of several well-established open-source libraries. The core packages used throughout the optimization process include:

- **NumPy** and **Pandas**: for matrix algebra, numerical computation, and efficient manipulation of structured data such as vectors, weight matrices, and asset-level metrics;
- **Matplotlib**, **Seaborn**, and **Plotly**: for generating both static and interactive visualizations of 2D and 3D efficient frontiers, asset weight distributions, and profile radar charts;
- **tqdm**: for real-time monitoring of long-running simulations, particularly during portfolio enumeration and frontier filtering;
- **openpyxl**: for exporting portfolio results and optimization outputs to .xlsx format using pandas.to_excel();
- **csv** and **pandas**: for exporting or importing large datasets in .csv format;
- **os** and **sys**: for managing relative paths, directory creation, and cross-folder imports within the thesis project structure;
- **sklearn.metrics.pairwise**: specifically for computing Euclidean distances between a priori and a posteriori portfolios during model comparison.

These packages collectively support not only the core optimization logic but also the full data lifecycle: from loading and preprocessing, to simulation and filtering, to exporting final results for analysis and presentation. For example, the script `Main_file.py` exports the entire 4D efficient frontier using pandas.to_csv(), while `Optimizer_a_priori.py` saves profile selections and their distance to the efficient frontier as Excel files. All file paths are modularized through custom logic in `path_utils.py`.

Code Architecture and Workflow

The Python scripts are organized in a modular architecture that separates data input, portfolio simulation, objective computation, optimization, Pareto filtering, and visualization into discrete files. This structure enhances maintainability, allows unit testing of individual components, and facilitates future extensions (e.g., adding a fifth dimension or new constraints).

For example:

- `Data_loader.py` loads and structures the input data, including expected returns, volatilities, ESG scores, and SCR values.
- `scr_calculator.py` computes the diversified SCR for any given portfolio using the Solvency II correlation matrix.
- `Portfolio_construction.py` handles random portfolio generation under feasibility constraints.
- `frontier.py` identifies Pareto-efficient portfolios from the full sample.
- `Optimizer_a_priori.py` handle scalarized optimization based on investor preferences.

The main execution scripts (`Main_file.py` for the *a posteriori* model and `Optimizer_a_priori.py` for the *a priori* model) coordinate the flow of these components and control the simulation parameters (Figure 5-8).

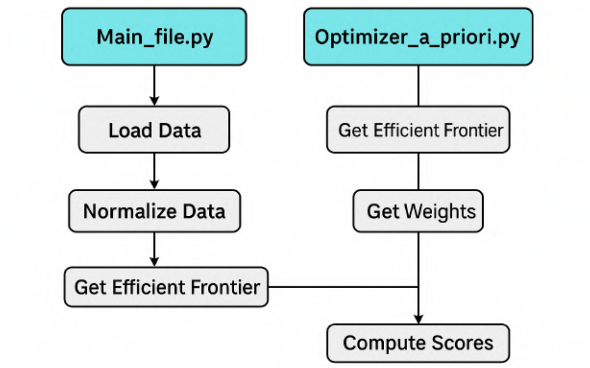

Source: Own Analysis

Figure 5-8: Execution scripts functionning

5.5. Performance and Computation Time

Given the number of portfolios evaluated and the dimensionality of the objective space, performance optimization was essential. On average, the simulation of +10 million feasible portfolios (including evaluation of all four objectives) required approximately 5 minutes on a modern personal machine.

In the *a priori* model, the optimizer typically converged in fewer than 100 iterations, with average solve times under 10 second per run. These performance benchmarks demonstrate that the modeling framework is suitable not only for academic exploration but also for prototyping real-world portfolio analysis tools.

6 Chapter

6.1. Data Collection and Preprocessing

The data used in this study spans 14 asset classes, reflecting a realistic investment universe for institutional investors. These include equity, bond, alternative, hybrid, and cash asset categories (Table 6-1). This fixed sample allows for tractable yet meaningful optimization testing, with enough variation in ESG and SCR profiles to enable robust comparative analysis across classical and extended frontiers.

Table 6-1: Portfolio used in the thesis work, composed of 14 asset classes spanning equity, bond, alternative, hybrid, and cash asset categories

Source: Own Analysis

	Investable Instrument	
Asset Class	**Ticker**	**Name**
EU Large Caps	EMUL.MI	iShares MSCI EMU Large Cap UCITS ETF EUR (Acc)
EU Small Caps	SMC.PA	SPDR MSCI Europe Small Cap UCITS ETF (Acc)
EM Equities	EIMI.L	iShares Core MSCI EM IMI UCITS ETF (Acc)
Developped World Equities	IWDA.L	iShares Core MSCI World UCITS ETF (Acc)
€ Govt Bonds	IBGX.L	iShares € Govt Bond 3-5yrs UCITS ETF (Dist) - Synthetic Acc Price
€ IG Corp Bonds	IBCX.L	iShares € Corp Bonds Large Cap UCITS ETF (Dist) - Synthetic Acc Price
HY Corp Bonds	IHGY.L	iShares € High Yield Corp Bond UCITS ETF (Dist) - Synthetic Acc Price
EM Bonds	EMBE.MI	iShares JP Morgan $ EM Bond EUR Hedged UCITS ETF (Dist) - Synthetic Acc Price
Global Convertible Bonds	GCVC.S	SPDR FTSE Global Convertible Bond CHF Hdg UCITS ETF (Acc)
Eu Core Real Estate	IPRE.DE	iShares European Property Yield UCITS ETF EUR (Acc)
Global Infrastructure	IGF.O	iShares Global Infrastructure ETF (Acc)
Commodities	WCOE.MI	WisdomTree Enhanced Commodity UCITS ETF EUR Hgd (Acc)
PE	PSP	Invesco global Listed Private Equity ETF (Acc)
Cash	SYB3.DE	SPDR Bloomberg 1-3Y Euro Govt Bd UCITS ETF (Dist) - Synthetic Acc Price

6.2. Variable Definitions (Return, Volatility, Market SCR, ESG Score)

6.2.1. Return and Risk

As seen *Chapter 1*, the starting point for data collection is to get expected returns and volatility for each asset classes. For that, historical averages were computed by using prices for each asset classes over a six years period. We extracted monthly closing prices in euros[2] using Refinitiv Eikon for asset classes being accumulative (Acc)[3] and processed them using Excel. Taking Acc denominated investable instruments enables to capture full potential market returns over the examined period without having to deal with intermediary cash flows in form of dividends, coupon payments, or else. However, as shown in Table 6-2, some asset classes did not have investable instruments (e.g.: Exchange Traded Fund (ETF)) that were accumulative. Instead, they were distributed (Dist)[4]. Thus, it was necessary to calculate "Synthetic Acc Prices".

Table 6-2: List of asset classes in Portfolio that were not Acc

Source: Own Analysis

	Investable Instrument	
Asset Class	**Ticker**	**Name**
€ Govt Bonds	IBGX.L	iShares € Govt Bond 3-5yrs UCITS ETF (Dist) - Synthetic Acc Price
€ IG Corp Bonds	IBCX.L	iShares € Corp Bonds Large Cap UCITS ETF (Dist) - Synthetic Acc Price
HY Corp Bonds	IHGY.L	iShares € High Yield Corp Bond UCITS ETF (Dist) - Synthetic Acc Price
EM Bonds	EMBE.MI	iShares JP Morgan $ EM Bond EUR Hedged UCITS ETF (Dist) - Synthetic Acc Price
Cash	SYB3.DE	SPDR Bloomberg 1-3Y Euro Govt Bd UCITS ETF (Dist) - Synthetic Acc Price

The "Synthetic Acc Price" implies taking the historical monthly Net Asset Values (NAV)[5] and the historical monthly cash out flows of the ETF and to compute returns based on these values. By starting with a base value of €100 (representing the Synthetic Acc Price for the very first period used in our data set) it is possible to get the closing prices of the fund by taking into account reinvestment of cash flows. This was done using the following formula:

$$Synthetic\ Acc\ Price_t = Synthetic\ Acc\ Price_{t-1} \times \frac{(NAV_t + DIV_t)}{NAV_{t-1}} \quad (31)$$

Where:

- DIV_t represents the dividend payed by the fund in period t.

[2] When the fund was denominated in a different currency, we used the currency conversion tool available on Refinitiv Eikon to convert back to EUR.

[3] An accumulative asset class refers to an investment vehicle or index that reinvests all income (such as dividends or interest payments) back into the asset itself, rather than distributing it to investors.

[4] A distributed asset class refers to an investment vehicle or index that pays out income, such as dividends, interest, or coupons, to investors on a regular basis. These distributions are not reinvested, which means the asset's price reflects only capital appreciation, not total return.

[5] Net Asset Value (NAV) represents the per-share value of an investment fund or asset, calculated as the total value of assets minus liabilities, divided by the number of outstanding shares. It is commonly used to determine the price at which investors buy or redeem shares in mutual funds or ETFs.

While the Synthetic Acc Price is not uesful in practice, it allows to compute the returns generated by the ETF. Once we had all asset classes' closing prices I computed the returns over the period using the standard return formula:

$$Return_t = \frac{Closing\ Price_t - Closing\ Price_{t-1}}{Closing\ Price_{t-1}} \quad (32)$$

Once returns were computed, we took the average of the entire data set using a basic =AVERAGE function in Excel. However, this average provides expected monthly returns which is not what we want for the data set. Thus, it is necessary to multiply this average by 12 to get annualized expected returns *(Alexander, 2008, pp. 90–95)*. Likewise, risk is computed as the standard deviation (volatility) of returns. By using the =STDEV formula in Excel, we get a monthly risk measure for each asset. Again, we annualize them by multiplying the average by $\sqrt{12}$ *(Alexander, 2008, pp. 90–95)*. Now, we get annual expected returns and standard deviations expressed in percentage (Table 6-3)

Table 6-3: Mean Annual Returns and Volatility for each components of the portfolio

Source: Own Analysis

Ticker	EMUL.MI	SMC.PA	EIMI.L	IWDA.L	IBGX.L	IBCX.L	IHGY.L
Asset Class	EU Large Caps	EU Small Caps	Emerging Markets Equity	Developed World Equity	Euro Government Bonds	Euro IG Corp Bonds	High yield Corp Bonds
Mean annual return	10.17%	6.79%	5.21%	14.11%	-0.40%	-0.06%	1.27%
Standard Deviation (Volatility)	17.36%	19.42%	14.30%	13.94%	3.25%	6.26%	9.37%

Ticker	EMBE.MI	GCVC.S	IPRE.DE	IGF.O	WCOE.MI	PSP	SYB3.DE
Asset Class	Emerging Market Bonds	Global Convertible Bonds	European Core Real Estate	Global Infrastructure	Commodities	Private Equity	Cash
Mean annual return	-2.83%	7.63%	-2.47%	5.31%	4.14%	5.82%	0.05%
Standard Deviation (Volatility)	11.58%	12.09%	23.24%	16.66%	12.54%	25.50%	1.54%

6.2.2. Variance

In addition to the two metrics discussed above, it was necessary to compute the variance of each asset classes. The variance is the square of the volatility (Standard Deviation) and represents a statistical measurement of the spread between numbers in a data set. It measures how far each number in the set is from the mean, and thus from every other number in the set *(Alexander, 2008).*

$$Variance = SD^2 \quad (33)$$

6.2.3. CAPM

As explained in *Chapter 1* the CAPM is an addition to the classical theory and aims at capturing an asset's systematic risk. CAPM is not used for optimization inputs but will serve as a benchmark in the discussion of return expectations and risk-adjusted performance.

Thus, market returns where needed and for that, the "market Portfolio" was taken. This "market portfolio" is proxied by the iShares MSCI ACWI ETF (Acc) (Ticker: ACWI). This ETF replicates the market portfolio as mentioned in the description found on Refinitiv Eikon:

"The Fund seeks investment results that correspond generally to the price and yield performance, before fees and expenses, of the MSCI All Country World Index. The Funds seeks to measure the combined equity market performance of developed and emerging markets countries.".

It was then necessary to compute returns and volatility for the "market portfolio". Of course, annualizing it as we did for the other asset classes is obligatory (Figure 6-1).

Ticker	ACWI
Asset Class	Market Portfolio
Mean annual return	**10.97%**
Standard Deviation (Volatility)	**14.23%**

Source: Own Analysis

Figure 6-1: Metrics for the "market portfolio"

Also, the CAPM requires to compute expected returns as explained in *Section 1.5.1*. Thus, we needed the risk free rate. For that, the German 10 years Government Bond with yields computed annualy was choosen. After adjusting for monthly yields by dividing the annual yields by 12, it was possible to compute the expected returns and the standard deviation of the risk free rate by using the same methods as discussed before (Figure 6-2).

German 10Y Government Bond	Risk free Rate
Mean annual return	**0.94%**
Standard Deviation (Volatility)	**0.39%**

Source: Own Analysis

Figure 6-2: Metrics for the "risk free rate"

Given all that, we could compute the beta of each asset classes. The beta is computed as follows:

$$\beta_i = \frac{Cov(r_i, r_m)}{Var(r_m)} \tag{34}$$

Where:

- $Cov(r_i, r_m)$ represents the covariance between asset i's returns and market returns;
- $Var(r_m)$ represents the variance of the market returns.

Excel Formula used: `=(STDEV.S(r(i))*CORREL(r(i), r(m)))/STDEV.S(r(m))`

6.2.4. Correlation Matrix

A correlation matrix was also drawn as it is required for computing the variance-covariance matrix. In simple terms, the correlation matrix captures the relationship between two variables – asset returns in our dataset. A correlation coefficient was computed as follows:

$$Corr(r_i, r_j) = \frac{\sum(r_i - \overline{r_i})\,(r_j - \overline{r_j})}{\sqrt{\sum(r_i - \overline{r_i})^2 \times \sum(r_j - \overline{r_j})^2}} \tag{35}$$

Excel Formula used: `=(CORREL(r(i),r(j))*Vol(r(i))*Vol(r(j)))*12`

Again, as our data is computed on a monthly basis, we multiply the result by 12 to get an annual metric. This gives us a 14 by 14 correlation matrix *(cf. Appendix 1.2.4)*.

6.2.5. Covariance Matrix

Lastly, a covariance matrix has also been constructed *(cf. Appendix 1.2.5)*. However, this was done only for comparison purpose as we will see that the code used for optimization purposes is integrating a line to compute a covariance matrix based on the correlation matrix and the volatilities of returns. The covariance coefficient can be computed as shown in *Section 1.1*, formula 4.

Excel Formula used: `=COVARIANCE.S(r(i),r(j))*12`

The relationship between the correlation and the covariance matrix is depicted by the following formula:

$$Corr(r_i, r_j) = \frac{Cov(r_i, r_j)}{\sigma_i \sigma_j} \tag{35}$$

6.3. ESG Score

The next step in preparing the dataset is to get ESG scores for each asset classes. As explained in *Section 2.2.1*, getting an ESG score is difficult as there exist no standard score on the market on which all providers agree on. According to P. Vaneerdewegh *(personal communication, March 2025)*, the most reliable way to get a correct ESG score is to take the scores provided by platforms issuing them. Thus, I took the scores provided by Refinitiv Eikon (1st provider) and Sustainalytics by Morningstar (2nd provider).

For the 1st provider, the spreadsheets provided on the platform show an LSEG Lipper Fund ESG Reporting Score (1-100 ESG score) based on an aggregate of different metrics *(cf. Appendix 1.2.6)*. This score stems as the most reliable score because unlike firm-level ESG metrics, the Lipper ESG score reflects the overall sustainability profile of the ETF vehicle itself, which is directly relevant for asset class representation in portfolio allocation *(LSEG Data & Analytics, 2023)*.

For the 2nd provider, we get a (100-1 risk score) scale that shows the relative risk associated with the company. This risk implies the potential danger for shareholders associated with ESG practices, lack of disclosure, and controversies.

While the scaling is different, the scores give the same information and provide shareholders with insights regarding the funds exposure to ESG matters *(cf. Appendix 1.2.7)*. This mismatch in scaling is explained in *Section 2.2.1* and shows, in practice, how difficult it is to get a standardized view on ESG scores associated with a fund.

In our case, the 1st provider provides ESG data for 10 out of 14 assets in the portfolio under study. These assets are listed in Table 6-4 (High Yield Corporate Bonds is a proxy asset class available). Missing asset classes on the 1st provider's platform are:

- € Govt Bonds;
- EM bonds;
- Commodities;
- Cash

For this thesis, it was not necessary to analyse the scores and understand why some asset classes performed better than other. Thus, the focus stays on getting the ESG score and assign it to the asset class.

Table 6-4: ESG scores as given by 1st provider ranked from best to worst score

Source: Own Analysis

Rank	Asset Class	ETF (ticker)	Score (highest = best)
1	European Core Real Estate	**iShares European Property Yield (IPRE)**	**70.753**
2	EU Large Cap	**iShares MSCI EMU Large Cap (EMUL)**	**60.062**
3	EU Small Cap	**SPDR MSCI Europe Small Cap (SMCX)**	**59.606**
4	Emerging Market Equity	**iShares Core MSCI EM IMI (EIMI)**	**58.647**
5	Global Infrastructure	**iShares Global Infrastructure ETF (IGF)**	**58.307**
6	Euro Corp Bonds	**iShares € Corp Bond Lg Cap (IBCX)**	**55.628**
7	Global Developed Equity	**iShares MSCI World (IWDA)**	**50.443**
8	Global Convertible Bonds	**SPDR FTSE Global Convertibles (SPF2)**	**49.901**
9	High Yield Corporate Bonds	**iShares Broad USD High Yield Corporate Bond ETF**	**44.512**
10	Private Equity	**Invesco Global Listed PE (PSP)**	**40.422**

On the 2nd provider's platform we did the same and looked-up ESG scores for the different asset classes. Here, 13 out of 14 asset classes got an ESG score assigned to them and are shown in

Table 6-5 (High Yield Corporate Bonds is a proxy asset class available). The missing asset class on the 2nd provider's platform is: Commodities.

Table 6-5: ESG scores as given by 2nd provider ranked from best to worst score

Source: Own Analysis

Rank	Asset Class	ETF (ticker)	Score (lowest = best)
1	European Core Real Estate	**iShares European Property Yield (IPRE)**	**9.81**
2	Euro Government Bonds	**iShares € Gov Bond 3–5yr (IBGX)**	**13.51**
3	Cash	**SPDR Bloomberg 1–3 Yr Gov (SYB3)**	**13.79**
4	EU Large Cap	**iShares MSCI EMU Large Cap (EMUL)**	**16.82**
5	Euro Corp Bonds	**iShares € Corp Bond Lg Cap (IBCX)**	**19.12**
6	EU Small Cap	**SPDR MSCI Europe Small Cap (SMCX)**	**19.14**
7	Global Infrastructure	**iShares Global Infrastructure ETF (IGF)**	**20.06**
8	Global Developed Equity	**iShares MSCI World (IWDA)**	**20.41**
9	Global Convertible Bonds	**SPDR FTSE Global Convertibles (SPF2)**	**21.84**
10	Private Equity	**Invesco Global Listed PE (PSP)**	**22.12**
11	Emerging Market Equity	**iShares Core MSCI EM IMI (EIMI)**	**23.19**
12	High Yield Corporate Bonds	**iShares Broad USD High Yield Corporate Bond ETF**	**25.42**
13	Emerging Market Bonds	**iShares JPM USD EM Bond (EMB)**	**40.85**

6.4. Missing ESG Scores

To complete the dataset and ensure every asset class received an ESG score, a methodology was developed to address the discrepancies between the two providers. The challenge of ESG data harmonization is well-documented in the literature. Berg, Kölbel, and Rigobon *(2022)* describe the phenomenon of "aggregate confusion", highlighting that ESG ratings across providers diverge due to differences in methodology, rather than inaccuracy. To mitigate this inconsistency, and consistent with their findings, a regression-based harmonization method was implemented in this thesis to align ESG scores from Refinitiv and Sustainalytics by Morningstar into a unified scale.

As mentioned, Refinitiv Eikon and Sustainalytics by Morningstar provide ESG scores based on different scales and methodologies. However, for 10 asset classes, data from both sources were available. This overlap enabled the creation of a quantitative bridge between the two scoring systems.

A linear regression model was fitted using the overlapping asset classes: the ESG Combined Score from Refinitiv (higher is better) was plotted against the ESG Risk Score from Susatianalytics by Morningstar (lower is better). Despite their different conceptual approaches, the two measures were found to exhibit a consistent and statistically meaningful inverse relationship[6] *(cf. Appendix 1.2.8)*. The resulting regression equation was as follows:

$$Refinitiv\ Eikon\ ESG\ score = -1.6519 \times Morningstar\ ESG\ score + 87.524$$

For each 1-point increase in the Sustainalytics by Morningstar ESG risk score (i.e., worse ESG), the Eikon ESG score drops by ~1.65 points (i.e., worse ESG). This confirms the inverse relationship between the two scoring systems was properly captured.

[6] Statistical significance is given by a R^2 score of 64.45%

This equation was then used to estimate Refinitiv-style ESG scores for the asset classes where only Sustainalytics by Morningstar scores were available (namely: Euro Government Bonds, Emerging Market Bonds, and Cash). This harmonization step enabled the construction of a comparable ESG score vector across 12 asset classes (Table 6-6).

This approach is consistent with the broader literature which highlights the need to adapt and harmonize ESG data across providers due to the lack of global standardization. While not perfect, this method ensures internal consistency and preserves the integrity of ESG analysis within the broader investment framework.

Table 6-6: ESG score of 14 asset classes based on the regression model (for 3 of them) ranked in alphabetic order for asset classes

Source: Own Analysis

Asset Class	2nd provider (lowest = best)	1st provider (highest = best)	Regression
Cash	13.79	64.744	= -1.6519*13.79+87.524
Commodities	0	0	
Emerging Market Bonds	40.85	20.044	= -1.6519*40.85+87.524
Emerging Market Equity	23.19	58.647	
EU Large Cap	16.82	60.062	
EU Small Cap	19.14	59.606	
Euro Corp Bonds	19.12	55.628	
Euro Government Bonds	13.51	65.207	= -1.6519*13.51+87.524
European Core Real Estate	9.81	70.753	
Global Convertible Bonds	21.84	49.901	
Global Developed Equity	20.41	50.443	
Global Infrastructure	20.06	58.307	
High Yield Corporate Bonds	25.42	44.512	
Private Equity	22.12	40.422	

This method has been extended to the whole dataset to provide consistency in ESG score. In other words, we want to have an ESG score that is consistent and we thus used the same methodology across the portfolio. So, instead of taking the regression result for 3 assets only, we will take it for the 13 out of 14 assets. These scores will represent the input for our model (1st provider). We end up with the following ESG scores for our portfolio:

Table 6-7: ESG score of 14 asset classes based on the regression model (whole portfolio) ranked in alphabetic order for asset classes

Source: Own Analysis

Asset Class	2nd provider (lowest = best)	1st provider (highest = best)	Regression
Cash	13.79	64.74	= -1.6519*13.79 + 87.524
Commodities	0	0	
Emerging Market Bonds	40.85	20.04	= -1.6519*40.85 + 87.526
Emerging Market Equity	23.19	49.22	= -1.6519*23.19 + 87.527
EU Large Cap	16.82	59.74	= -1.6519*16.82 + 87.528
EU Small Cap	19.14	55.91	= -1.6519*19.14 + 87.529
Euro Corp Bonds	19.12	55.94	= -1.6519*19.12 + 87.530
Euro Government Bonds	13.51	65.21	= -1.6519*13.51 + 87.531
European Core Real Estate	9.81	71.32	= -1.6519*9.81 + 87.532
Global Convertible Bonds	21.84	51.45	= -1.6519*21.84 + 87.533
Global Developed Equity	20.41	53.81	= -1.6519*20.41 + 87.534
Global Infrastructure	20.06	54.39	= -1.6519*20.06 + 87.535
High Yield Corporate Bonds	25.42	45.53	= -1.6519*25.42 + 87.536
Private Equity	22.12	50.98	= -1.6519*22.12 + 87.537

Commodities represent a unique challenge in ESG analysis, as they are not corporate issuers and thus do not engage in activities subject to traditional ESG disclosure or governance. Consequently, no ESG scores are available from either Refinitiv Eikon or Sustainalytics by Morningstar for this asset class.

Given the nature of most broad commodity indices, typically composed of energy (oil, gas), metals (including precious and industrial), and agricultural products, it is widely acknowledged that they tend to perform poorly on environmental criteria, particularly due to their association with carbon emissions, resource extraction, and land use intensity.

In light of these considerations, and in the absence of standardized ESG metrics, commodities were conservatively assigned the lowest possible ESG score (0) in this model. This assumption reflects both the lack of transparency and the generally high environmental externalities associated with this asset class. Similar approaches are taken in ESG-integrated investment frameworks, where commodities are often excluded or penalized unless specifically structured with ESG overlays (e.g., ESG-screened commodity ETFs, or carbon-neutral commodity baskets).

This methodological choice is also supported by Berg et al. *(2022)*, who emphasize the complexity of ESG scoring divergence and the importance of transparent rule-based assumptions when data is unavailable or structurally incompatible.

6.5. Market SCR

This thesis integrates the Market Solvency Capital Requirement (SCR) as the fourth dimension in the portfolio optimization model. Under the Solvency II directive, the Market SCR captures the capital an insurer must hold to cover adverse movements in financial markets, making it a regulatory measure of market risk rather than a statistical one like volatility *(EIOPA, 2024)*.

Initially, the Market SCR was computed based on fixed monetary allocations (EUR) across asset classes. However, this approach proved incompatible with the optimization model, which is built on relative portfolio weights and not absolute investment values. The calculation was revised to operate on the basis of risk per unit of investment, ensuring consistency with the model's internal logic.

The approach proceeds as follows:

- A notional investment of €1 was assumed for each asset class to eliminate the dependency on absolute values.
- Based on Solvency II guidelines, stress factors were applied to the relevant risk modules (spread, equity, currency) for each asset class (Table 6-8)
- For example, equities denominated in euros were exposed to equity risk only, while emerging market bonds could be subject to both spread and currency risk. The risk load for each type was directly equal to its stress factor, given the €1 normalization.

This gives the regulatory stress load per €1 invested for each asset classes (Table 6-9).

Table 6-8: Stress factors as defined by Solvency II

Risk Type	Standard Stress Factor	Source (Solvency II)
Equity (Type 1 - Developed Markets)	39%	Art. 168–169
Equity (Type 2 - Emerging/Private)	49%	Art. 168–169
Spread Risk (Investment Grade)	10%	Art. 176
Spread Risk (High Yield)	20%	Art. 176
Currency Risk	25%	Art. 188
Property Risk	25%	Art. 183
Interest Rate Risk (Up)	100%	Art. 164
Interest Rate Risk (Down)	70%	Art. 164

Table 6-9: Regulatory stress load per €1 invested for each asset classes

Source: Own Analysis

Asset Class	MV_pct	Spread Risk Load	Equity Risk Load	Currency Risk Load
EU Large Caps	7.14	0	0.39	0
EU Small Caps	7.14	0	0.39	0
EM Equities	7.14	0	0.49	0.25
Developped World Equities	7.14	0	0.39	0.25
€ Govt Bonds	7.14	0	0	0
€ IG Corp Bonds	7.14	0.10	0	0
HY Corp Bonds	7.14	0.20	0	0
EM Bonds	7.14	0.20	0	0.25
Global Convertible Bonds	7.14	0.20	0.39	0.25
Eu Core Real Estate	7.14	0	0.39	0
Global Infrastructure	7.14	0.20	0.39	0.25
Commodities	7.14	0	0	0.25
PE	7.14	0	0.49	0.25
Cash	7.14	0	0	0

These Risk Load are then immediately implemented in the `scr_calculator.py` file and serve to compute the estimation of the SCR Market Cost after diversification for each portfolio created. Refer to Section 5.4.2 for a complete overview on how the SCR for a portfolio is computed. This approach respects the logic and structure of Solvency II while making the metric fully compatible with optimization modelling. It also aligns with best practices highlighted by EIOPA *(2024)* and the Official Journal of The European Union *(2025)*, ensuring that regulatory risk is treated in a coherent and scalable way across asset classes.

6.6. Model Implementation and Simulations

At the conclusion of the data preparation process, each of the 14 asset classes in the portfolio was assigned a four-dimensional vector capturing the key metrics used in the optimization model (Figure 6-3). These vectors are defined as:

$$[Expected\ return, Volatility, ESG\ Score, \{Spread\ Risk\ Load, Equity\ Risk\ Load, Currency\ Risk\ Load\}]$$

Each component was collected or computed using the best available data sources and aligned with the thesis' multidimensional framework:

- Expected Return and Volatility were estimated using historical monthly returns.
- ESG Score was harmonized across providers using a regression-based approach.
- Risk loads for Market SCR computation using a Solvency II–based methodology reflects regulatory capital cost per €1 invested.

This four-dimensional structure allows for the application of advanced portfolio optimization techniques, where performance is evaluated not just in terms of risk and return, but also sustainability (ESG) and regulatory capital efficiency (Market SCR). These vectors serve as the foundational inputs for constructing efficient frontiers in two, three, and ultimately four dimensions.

				SCR Market		
Asset Class	**E(R)**	**Risk**	**ESG (/100)**	**Spread Risk Load**	**Equity Risk Load**	**Currency Risk Load**
EU Large Caps	11.37%	17.36%	59.74	0.00%	39.00%	0.00%
EU Small Caps	12.78%	19.42%	55.91	0.00%	39.00%	0.00%
EM Equities	8.28%	14.30%	49.22	0.00%	49.00%	25.00%
Developped World Equities	10.39%	13.94%	53.81	0.00%	39.00%	25.00%
€ Govt Bonds	1.89%	3.25%	65.21	0.00%	0.00%	0.00%
€ IG Corp Bonds	4.09%	6.26%	54.39	10.00%	0.00%	0.00%
HY Corp Bonds	6.00%	9.37%	45.53	20.00%	0.00%	0.00%
EM Bonds	6.57%	11.58%	20.04	20.00%	0.00%	25.00%
Global Convertible Bonds	7.05%	12.09%	51.45	20.00%	39.00%	25.00%
Eu Core Real Estate	11.71%	23.24%	71.32	0.00%	39.00%	0.00%
Global Infrastructure	10.54%	16.66%	54.39	20.00%	39.00%	25.00%
Commodities	4.24%	12.54%	0.00	0.00%	0.00%	25.00%
PE	17.12%	25.50%	50.98	0.00%	49.00%	25.00%
Cash	0.91%	1.54%	64.74	0.00%	0.00%	0.00%

Source: Own Analysis

Figure 6-3: Final vectors for the implementation in the optimization model

7 Chapter

7.1. Presentation of 4D Efficient Frontier Findings

The empirical construction of the four-dimensional efficient frontier represents a pivotal moment in this thesis. It operationalizes the theoretical extension of classical MPT to incorporate non-financial objectives, specifically ESG scores and Market SCR. This section organizes the empirical results of the 4D efficient frontier into recognizable patterns and portfolio archetypes based on their performance across return, volatility, ESG score, and Market SCR. These patterns help quantify the internal structure of the frontier, offering an overview of the composition and statistical features of non-dominated portfolios generated through the *a posteriori* optimization.

The simulation phase, executed using the script `Main_file.py`, produced 10,400,600 portfolios constructed from 14 asset classes using enumerized tranche-based portfolios constrained by non-negativity, full investment (sum to 1), and minimum weight thresholds *(cf. Section 5.2)*. This number is very specific to the set up choosen: Tranches of 3.5% and minimum 1 tranche per asset class. Generating portfolios with tranches of 4% with minimum 1 tranche per asset would yield ≈ 2,300,000 possibilities. However, this is not the subject of the thesis and we will not elaborate on that. As seen in Figure 7-1, the script `generate_all_portfolios()` computed each portfolio's return, volatility, ESG score, and diversified SCR using vectorized stress contributions and regulatory correlation matrices.

```
Generated 10400600 portfolios (100% of valid combinations with 3.75% granularity
and minimum allocation).
Portfolio 4D preview [Return, Volatility, ESG, SCR]:
[[4.62310464e-02 3.15304637e-02 5.69776870e+01 1.77891168e-01]
 [5.22323635e-02 3.82001523e-02 5.64680453e+01 2.00067204e-01]
 [5.82336805e-02 4.58619725e-02 5.59584035e+01 2.22276228e-01]
 [6.42349975e-02 5.40959961e-02 5.54487618e+01 2.44509251e-01]
 [7.02363146e-02 6.26771127e-02 5.49391200e+01 2.66760272e-01]]
```

Source: Own Analysis

Figure 7-1: Output of the `Main_file.py` code with the function `generate_all_portfolios()`

A total of 1,203,671 non-dominated portfolios were subsequently identified using the `get_efficient_frontier()` function, which implemented a custom Pareto filtering algorithm defined in `frontier.py` *(cf. Appendix 1.2.1)*. Each retained portfolio constitutes a point on the 4D efficient frontier, with no alternative portfolio in the feasible set strictly outperforming it across all four objectives (Figure 7-2).

```
Efficient frontier contains 1203671 portfolios.
First 5 efficient portfolios:
    Return  Volatility     ESG    SCR
0   0.111        0.109  51.811  0.398
1   0.111        0.107  52.382  0.398
2   0.111        0.106  52.953  0.398
3   0.110        0.107  54.094  0.398
4   0.110        0.105  52.386  0.398
```

Source: Own Analysis

Figure 7-2: Output of the `Main_file.py` code with the function `get_efficient_frontier()`

Figure 7-3 illustrates the frontier's structure via 2D projections where ESG and SCR are mapped onto color gradients. The ESG-colored plot reveals a dense region of high-ESG portfolios within moderate return and risk bands, while the SCR-colored plot shows SCR-efficient portfolios located toward the lower-risk end. These visuals demonstrate that ESG and SCR are not trivially correlated with return or volatility, thereby supporting Hypotheses H1 and H2, which posit the multidimensional nature of optimality and the integrability of ESG/SCR as quantitative objectives. Moreover, Figure 7-4 plots ESG versus Market SCR colored volatility and return. Adding these gives an overview of the capital costs of ESG tilts. The triangular envelope reveals that pushing ESG above 58 almost always coincides with both higher capital requirements and higher volatility. These claims are further confirmed by the 2D projections exhibited in *Appendix 1.2.9*, which showcase all different pairs outputted by the program.

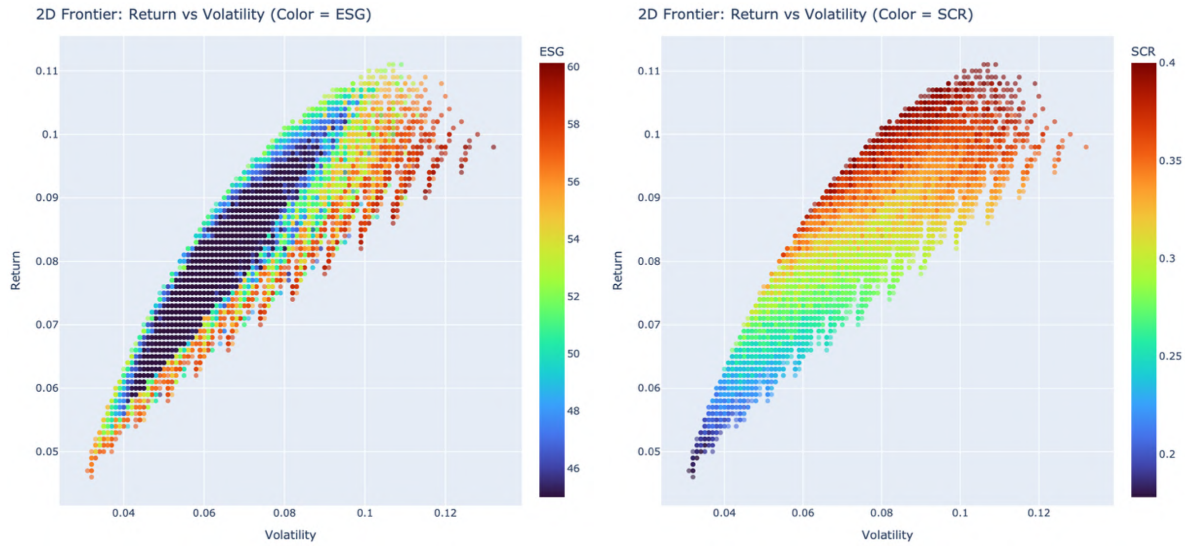

Source: Own Analysis

Figure 7-3: 2D projections with ESG and SCR as colors

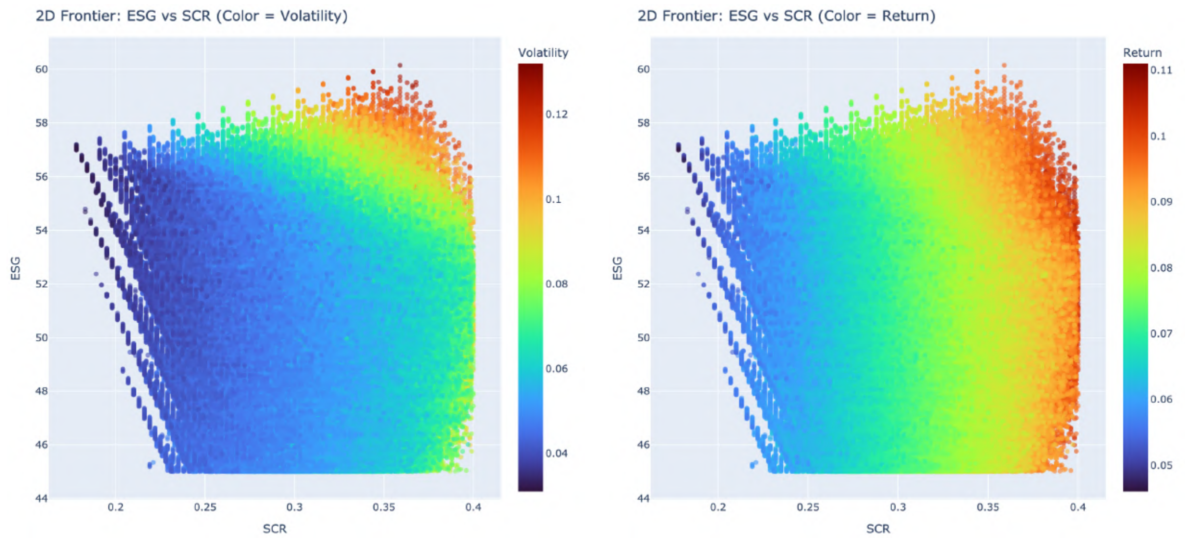

Source: Own Analysis

Figure 7-4: 2D projections with Volatility and Retrun as colors

The diversity of the efficient frontier is also reflected in the histograms displayed in Figure 7-5, generated using `Results_analysis.py`. All three financial metrics, return, volatility and SCR, display pronounced central peaks. Mandatory diversification (every asset receives ≥ one tranche) mechanically limits extreme bets, so the bulk of portfolios occupy a "middle ground" on traditional risk-reward dimensions. The ESG distribution is visibly left-skewed: falling below the mean ESG level is relatively easy, whereas achieving genuinely high ESG (> 55) is difficult. Investors pursuing best-in-class sustainability face a limited opportunity set and must accept concessions in other objectives. Volatility and SCR approximate normal curves, confirming that diversification across 14 asset classes mitigates extreme risk realisations. Because extreme observations are rare on every axis, only a thin shell of portfolios can dominate the modal "average" solution. Without the epsilon-based Pareto screen and subsequent lexicographic sort *(cf. Section 5.1.1)*, more than 90 % of the universe would be labelled efficient by numerical coincidence. Relevance for investor profiles can be described as follows:

- **Return-oriented** investors will mine the right-hand tail of the return distribution, trading off higher volatility and SCR.
- **ESG-focused** mandates concentrate on the sparse high-ESG region, often surrendering return.
- **Capital-sensitive** mandates target the left tail of the SCR histogram and must compromise on one or more of the other three objectives.

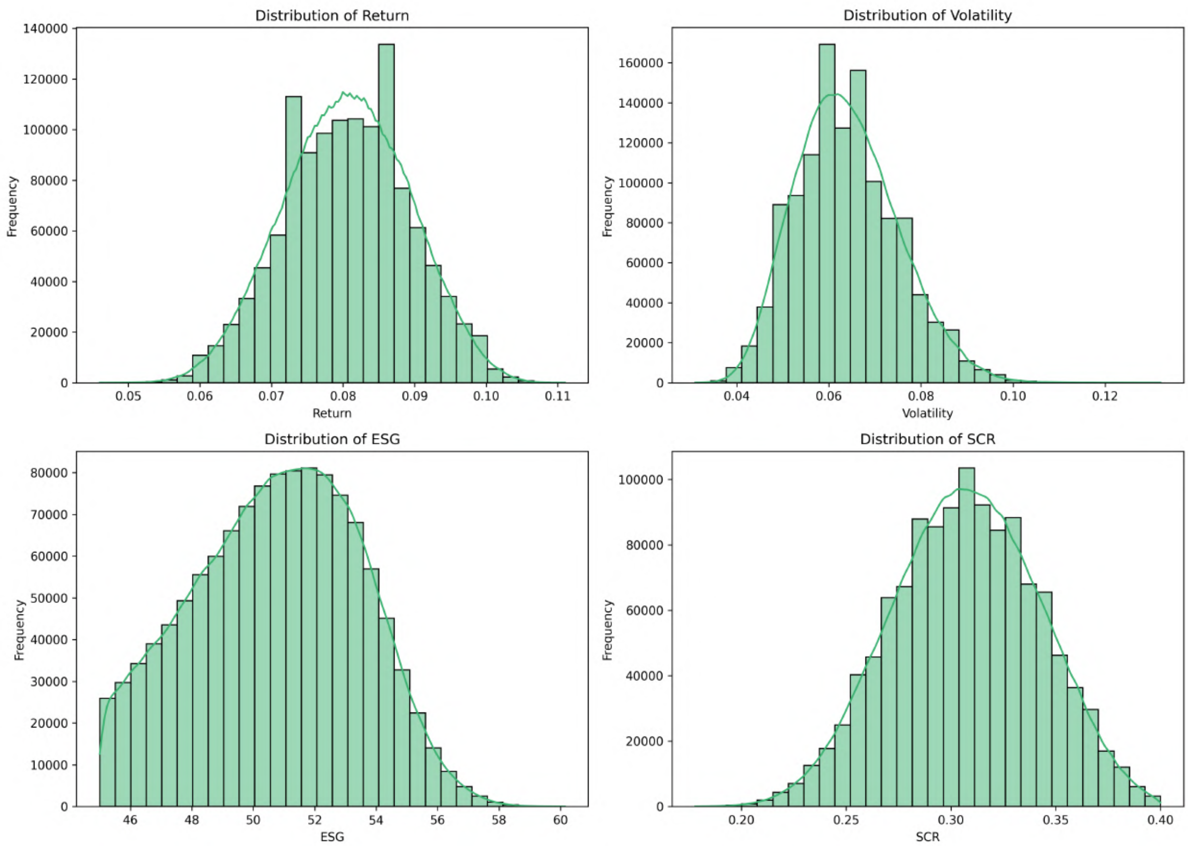

Source: Own Analysis

Figure 7-5: Histograms of the distribution of the 4 dimensions within the portfolios making the efficient frontier

Figure 7-6 provides descriptive statistics of the 1,203,671 non-dominated portfolios. The average expected return is 8.05% with a standard deviation of 0.087%, while average volatility stands at 6.39%. ESG scores average 50.7 and SCR scores average 30.07%, confirming that the simulated frontier spans a wide range of possible trade-offs.

```
Descriptive statistics (Return, Volatility, ESG Score, SCR Score):
            Return    Volatility           ESG           SCR
count  1.203671e+06  1.203671e+06  1.203671e+06  1.203671e+06
mean   8.050000e-02  6.390000e-02  5.074040e+01  3.071000e-01
std    8.700000e-03  1.100000e-02  2.705000e+00  3.490000e-02
min    4.600000e-02  3.100000e-02  4.500000e+01  1.780000e-01
25%    7.400000e-02  5.600000e-02  4.873900e+01  2.830000e-01
50%    8.100000e-02  6.300000e-02  5.087200e+01  3.070000e-01
75%    8.700000e-02  7.100000e-02  5.277200e+01  3.320000e-01
max    1.110000e-01  1.320000e-01  6.014300e+01  4.000000e-01
============================================================
```

Source: Own Analysis

Figure 7-6: Descriptive statistics of the efficient frontier

Table 7-1, Table 7-2, Table 7-3, Table 7-4 report the four "corner" portfolios extracted from the 4D efficient frontier: the portfolio with maximum expected return, the one with highest ESG score, and the portfolios with minimum Market SCR and minimum volatility. Together they illustrate how optimising along each objective pulls the allocation toward a distinct part of the frontier and, in turn, to a markedly different asset mix.

Max-Return portfolio achieves an 11.1 % expected return but does so with higher volatility (10.9 %) and SCR (39.8 %). Its allocation is dominated by EU Small Caps (33 %) and Private Equity (22 %), confirming that aggressive growth assets remain the primary engine of performance even in the 4D framework.

Max-ESG portfolio lifts the weighted ESG score to 60.1 while still earning 9.8 % return. More than half of the capital is committed to EU Core Real Estate, a high-ESG, moderate-risk asset class. SCR drops to 35.9 %, but volatility rises to 13.2 %, indicating that the sustainability premium is paid mainly through higher price variability rather than an outsized capital charge.

Min-SCR portfolio cuts the Market SCR requirement to 17.8 % by shifting 52 % of the portfolio into euro-government bonds. Despite its defensive tilt, the portfolio still carries an ESG score of 57.2. The trade-off is a reduced return (5.1 %) and a very low volatility profile (3.5 %).

Min-Volatility portfolio records the smoothest risk profile at 3.1 % volatility by parking nearly half of the capital in cash and adding a small overweight to euro-government bonds. Return slips only slightly below the Min-SCR case (4.7 %), ESG remains healthy at 57.0, and SCR again equals 17.8 %.

Table 7-1: Key portfolios on the frontier: Max Return

Source: Own Analysis

Portfolio metric	Value
Expected return	**11.1 %**
Volatility	10.9 %
ESG score	51.8
Market SCR	0.398
Asset class	**Weight**
EU Large Caps	3.70 %
EU Small Caps	**33.33 %**
EM Equities	3.70 %
Developed World Equities	3.70 %
€ Government Bonds	3.70 %
€ IG Corporate Bonds	3.70 %
High-Yield Corporate Bonds	3.70 %
EM Bonds	3.70 %
Global Convertible Bonds	3.70 %
EU Core Real Estate	3.70 %
Global Infrastructure	3.70 %
Commodities	3.70 %
Private Equity (PE)	**22.22 %**
Cash	3.70 %

Table 7-2: Key portfolios on the frontier: ESG Max

Source: Own Analysis

Portfolio metric	Value
Expected return	9.8 %
Volatility	13.2 %
ESG score	**60.1**
Market SCR	0.359
Asset class	**Weight**
EU Large Caps	3.70 %
EU Small Caps	3.70 %
EM Equities	3.70 %
Developed World Equities	3.70 %
€ Government Bonds	3.70 %
€ IG Corporate Bonds	3.70 %
High-Yield Corporate Bonds	3.70 %
EM Bonds	3.70 %
Global Convertible Bonds	3.70 %
EU Core Real Estate	**51.85 %**
Global Infrastructure	3.70 %
Commodities	3.70 %
Private Equity (PE)	3.70 %
Cash	3.70 %

Table 7-3: Key portfolios on the frontier: Min Market SCR

Source: Own Analysis

Portfolio metric	Value
Expected return	5.1 %
Volatility	3.5 %
ESG score	57.2
Market SCR	**0.178**
Asset class	**Weight**
EU Large Caps	3.70 %
EU Small Caps	3.70 %
EM Equities	3.70 %
Developed World Equities	3.70 %
€ Government Bonds	**51.85 %**
€ IG Corporate Bonds	3.70 %
High-Yield Corporate Bonds	3.70 %
EM Bonds	3.70 %
Global Convertible Bonds	3.70 %
EU Core Real Estate	3.70 %
Global Infrastructure	3.70 %
Commodities	3.70 %
Private Equity (PE)	3.70 %
Cash	3.70 %

Table 7-4: Key portfolios on the frontier: Min Risk

Source: Own Analysis

Portfolio metric	Value
Expected return	4.7 %
Volatility	**3.1 %**
ESG score	57.0
Market SCR	0.178
Asset class	**Weight**
EU Large Caps	3.70 %
EU Small Caps	3.70 %
EM Equities	3.70 %
Developed World Equities	3.70 %
€ Government Bonds	**7.41 %**
€ IG Corporate Bonds	3.70 %
High-Yield Corporate Bonds	3.70 %
EM Bonds	3.70 %
Global Convertible Bonds	3.70 %
EU Core Real Estate	3.70 %
Global Infrastructure	3.70 %
Commodities	3.70 %
Private Equity (PE)	3.70 %
Cash	**48.15 %**

The 2D heat-maps in Figure 7-3 and Figure 7-4 already hinted at three systematic relationships within the 4D frontier: a risk-return slope in which higher expected return is accompanied by higher volatility and larger Market SCR; a sustainability ridge showing that ESG scores above about 58 appear only in the upper-right corner of the ESG–SCR plane; and a capital-light wedge where low-SCR portfolios cluster in the lower-left of every projection, characterised by subdued risk and return.

Figure 7-7 and Figure 7-8, the interactive 3D renderings, deepen and confirm this reading. When ESG is placed on the vertical axis and SCR is rendered as a colour gradient (Figure 7-7), the ridge inferred in two dimensions materialises as a broad, gently rising plateau: ESG values between fifty and fifty-five can be reached across a wide band of return–volatility combinations while the colour scale shows SCR remaining moderate, roughly 0.25 to 0.32. Only when the observer approaches the fifty-five-to-fifty-eight band does the surface narrow sharply, exactly the pinch-point suggested by the 2D plot. Reversing the axes in Figure 7-8 places SCR on the vertical dimension and colours the surface by ESG. This perspective exposes the capital-efficient layer as a thin, almost horizontal sheet of points with SCR below about 0.22; its uniform blue-to-green colouring confirms that these portfolios carry mid-forties to low-fifties ESG scores. As one ascends the SCR axis the colour palette fans out, demonstrating that genuinely high-ESG allocations exist only at higher capital charges, again echoing the earlier 2D insight.

Rotating the cloud interactively also reveals that the classical risk–return relationship appears as a diagonal spine running through the entire frontier: no hidden shortcuts or isolated peaks emerge when the landscape is viewed in three dimensions. Thus the 3D projections corroborate every qualitative pattern extracted from the 2D slices and, by adding spatial depth, show that those apparently linear “edges” are in fact continuous surfaces. The visual evidence therefore reinforces the central message that optimality in the four-objective setting must be conceived as navigation over a multidimensional manifold rather than the identification of a single apex point. Further plots may be found in appendix X where other trio of dimensions have been plotted against each other.

It is important to note that because of computational capabilities, the frontier containing 1,203,671 could not be plotted in its entirety. Therefore, Figure 7-7 and Figure 7-8 only contain 600,000 portfolios randomly taken from the entire set. This is seen as a valid approach as descriptive statistics of the full set and the sub set show minor differences with regards to e.g. the mean and the standard deviation. Figure 7-9 provides these descriptive statistics confirming that working with the subset yields the same frontiers.

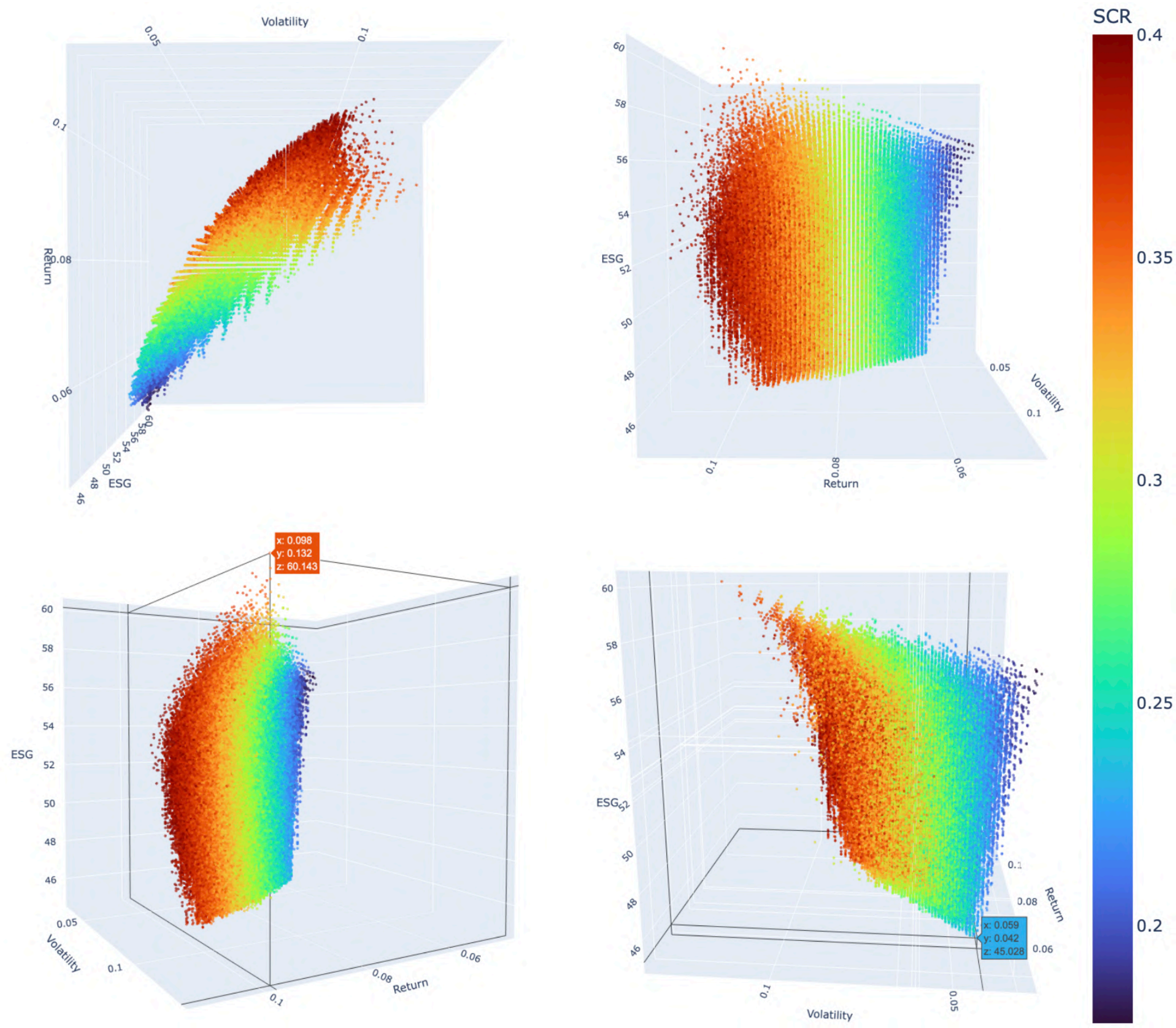

Source: Own Analysis

Figure 7-7: 3D projections of the efficient frontier with Return (X), Volatility (Y), ESG (Z), and SCR mapped as color gradient

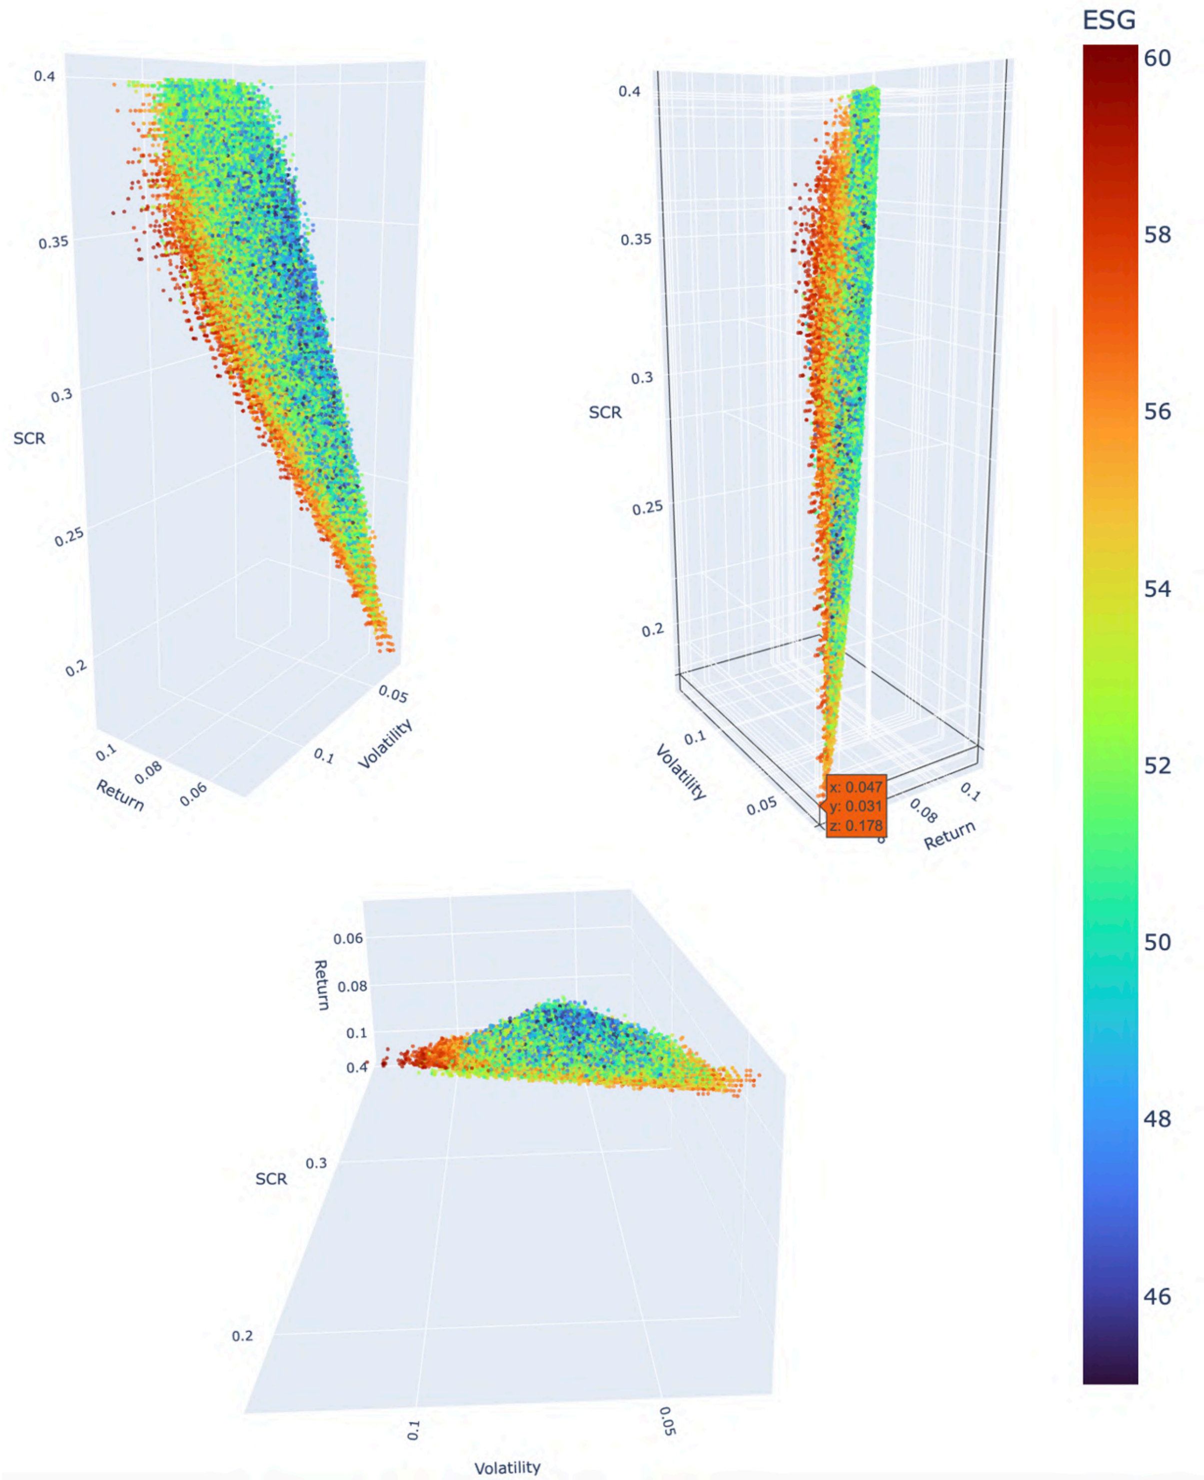

Source: Own Analysis

Figure 7-8: 3D projections of the efficient frontier with Return (X), Volatility (Y), SCR (Z), and ESG mapped as color gradient

```
Data Summary for Sub Sample (600,000 rows):
              Return     Volatility            ESG            SCR
count  600000.000000  600000.000000  600000.000000  600000.000000
mean        0.080529       0.063931      50.744700       0.307063
std         0.008729       0.011012       2.705369       0.034920
min         0.047000       0.031000      45.000000       0.178000
25%         0.074000       0.056000      48.743000       0.283000
50%         0.081000       0.063000      50.877000       0.307000
75%         0.087000       0.071000      52.781000       0.332000
max         0.111000       0.132000      60.143000       0.400000

Data Summary for Full Set:
              Return     Volatility           ESG           SCR
count   1.203671e+06   1.203671e+06  1.203671e+06  1.203671e+06
mean        0.085413       0.063944      50.74043      0.307115
std         0.008735       0.011024      2.704993      0.034930
min         0.046000      0.031000      45.000000     0.178000
25%         0.074000      0.056000      48.739000     0.283000
50%         0.081000      0.063000      50.872000     0.307000
75%         0.087000      0.071000      52.772000     0.332000
max         0.111000      0.132000      60.143000     0.400000
```

Source: Own Analysis

Figure 7-9: Descriptive statistics from the full and sub sample of the efficient frontier

7.2. Comparison with Traditional 2D Model

The classical 2D Markowitz model, defined by the trade-off between return and volatility, serves as the natural benchmark for evaluating the added value of multidimensional optimization. This section compares the outputs of the 4D model introduced in earlier with the efficient frontier and optimal portfolios derived from the traditional 2D setup implemented in `Main_Classical_MPT.py`. Figure 7-10 displays the 2D efficient frontier, overlaid with the Capital Market Line (CML), the tangency portfolio, and two key reference points: the maximum Sharpe ratio and minimum volatility portfolios. This visualization shows that the classical model achieves an optimal return of 12.58% at a volatility of 14.67%. For the Tangency Portfolio, the program allocates 18.52% to €IG Corporate Bonds for which the return is quite low but the volatility too.Then it assigns 11.11% to Developped World Equities, €Govt Bonds, and HY Corp Bonds. This mix between high/low return and high/low volatility shows that the program correctly picks asset classes that yield a balanced allocation.

Table 7-5: Allocation, Returns, and Volatilities for the Tangency Portfolio in the 2D model

Source: Own Analysis

Asset	Weight	E(R)	Volatility
€ IG Corp Bonds	0.185185	0.040933	0.062648
Developped World Equities	0.111111	0.103926	0.139441
€ Govt Bonds	0.111111	0.018946	0.032493
HY Corp Bonds	0.111111	0.059987	0.093659
EM Equities	0.074074	0.082754	0.142956
EM Bonds	0.074074	0.065655	0.115795
Global Convertible Bonds	0.074074	0.070486	0.120864
EU Large Caps	0.037037	0.113691	0.173647
EU Small Caps	0.037037	0.127799	0.194188
Eu Core Real Estate	0.037037	0.117078	0.232425
Global Infrastructure	0.037037	0.105368	0.166615
Commodities	0.037037	0.042354	0.125357
PE	0.037037	0.171184	0.254979
Cash	0.037037	0.009148	0.015389

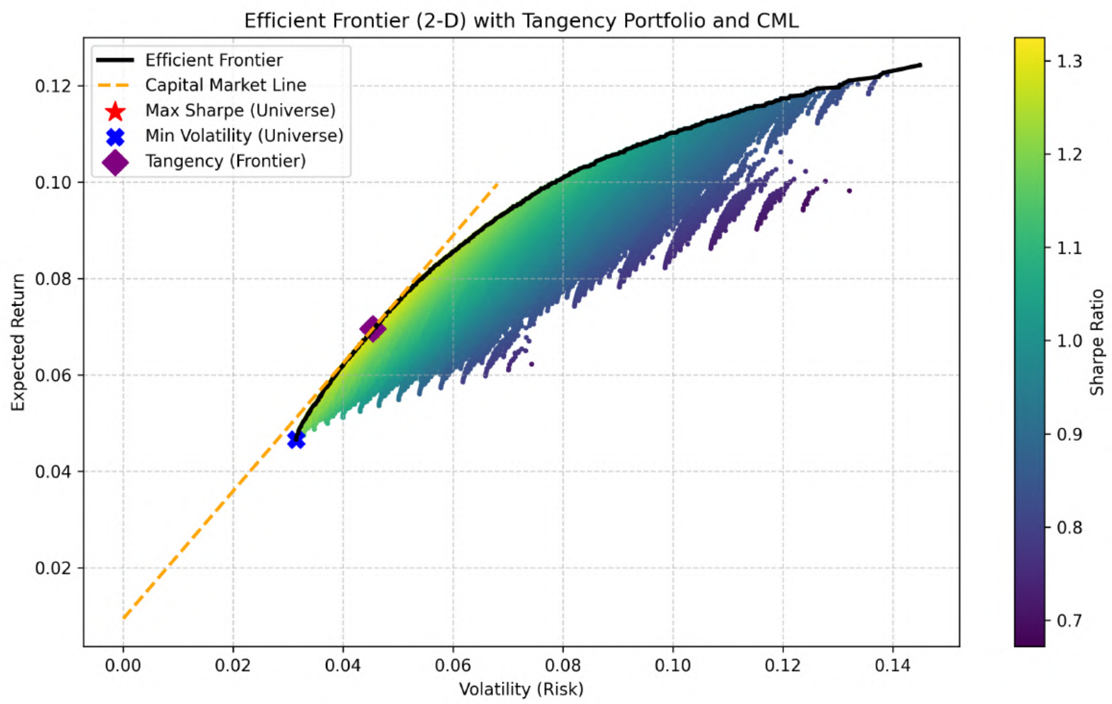

Source: Own Analysis

Figure 7-10: 2D efficient frontier overlaid with the Capital Market Line (CML), the tangency portfolio, and two key reference points: the maximum Sharpe ratio and minimum volatility portfolios

Figure 7-11 superimposes the classical 2D efficient frontier on the 4D frontier projected into the same return–volatility plane. The blue arc produced by MPT is preserved, yet almost every 4D efficient portfolio (green points) falls just inside that arc. In practical terms, extending the optimisation to embrace ESG and Market-SCR does not displace the core envelope: median return and volatility move from 8.39 % and 6.40 % in the 2D set to 8.05 % and 6.39 % in the 4D set. The only region where the two frontiers diverge is the high-return, high-volatility tail. Above roughly 10 % expected return and 11 % volatility, portfolios that are efficient in two dimensions are screened out once sustainability and capital efficiency are taken into account, demonstrating that the additional criteria prune extreme risk-seeking allocations rather than dragging the entire surface downward.

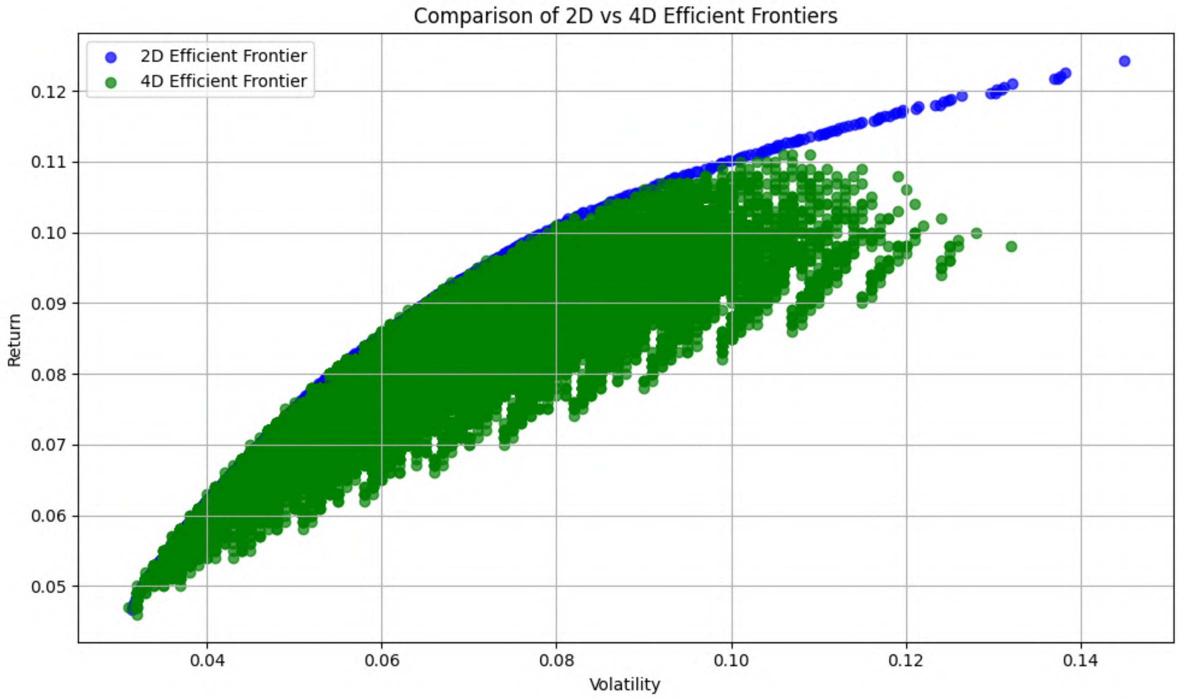

Source: Own Analysis

Figure 7-11: Comparison of 2D vs 4D Efficient Frontiers

The density plots in Figure 7-12 quantify that visual impression. When ESG and SCR constraints are introduced, the dispersion of both return and volatility tightens by roughly fifty per cent, concentrating the efficient set around its modal risk-return coordinates. The lower panels of the same figure reveal how the new objectives behave in aggregate: ESG scores fan out smoothly between 45 and 60 with a mean of 50.7, while SCR values cluster in a narrow band around 0.30. Because neither variable was bounded ex-ante, the orderly shape of these distributions shows that the model internalises non-financial objectives without warping the traditional landscape, a point that supports Hypotheses H1 and H2.

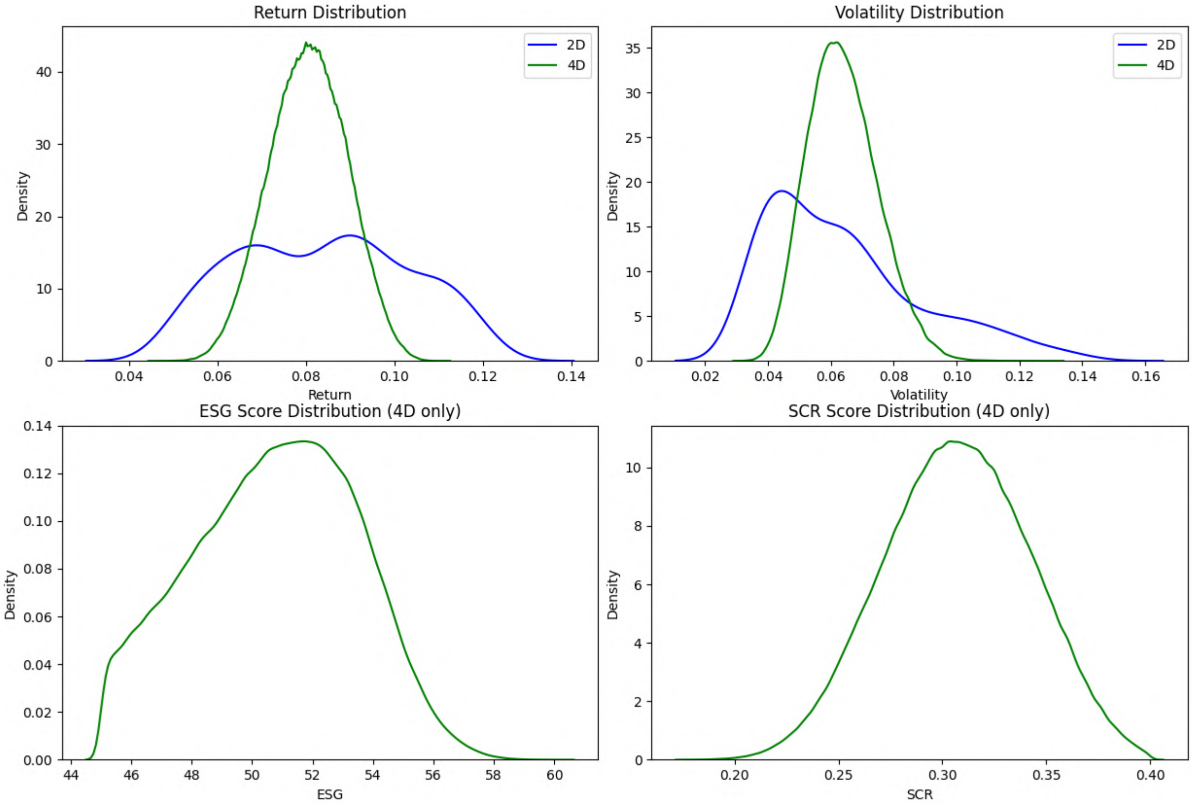

Source: Own Analysis

Figure 7-12: Distribution comparison of Return, Volatility, ESG Score, and SCR score

Three-dimensional scatter-plots (Figure 7-7 and Figure 7-8) corroborate the patterns already visible in two dimensions. When ESG is mapped to the vertical axis and SCR to the colour gradient, a broad plateau appears where portfolios combine sustainability scores above 54 with only average capital charges, a visual confirmation that "double-good'' solutions are plentiful. Conversely, plotting SCR on the vertical axis and colouring points by ESG reveals a ridge of distinctly low-capital portfolios; it narrows rapidly below an SCR of 0.25, mirroring the tapering left tail of the histogram in Figure 7-13. Because these three-dimensional views reproduce rather than overturn the curvature seen in the 2D projections, they strengthen the inference that the efficient set is a surface of admissible trade-offs rather than a single best point.

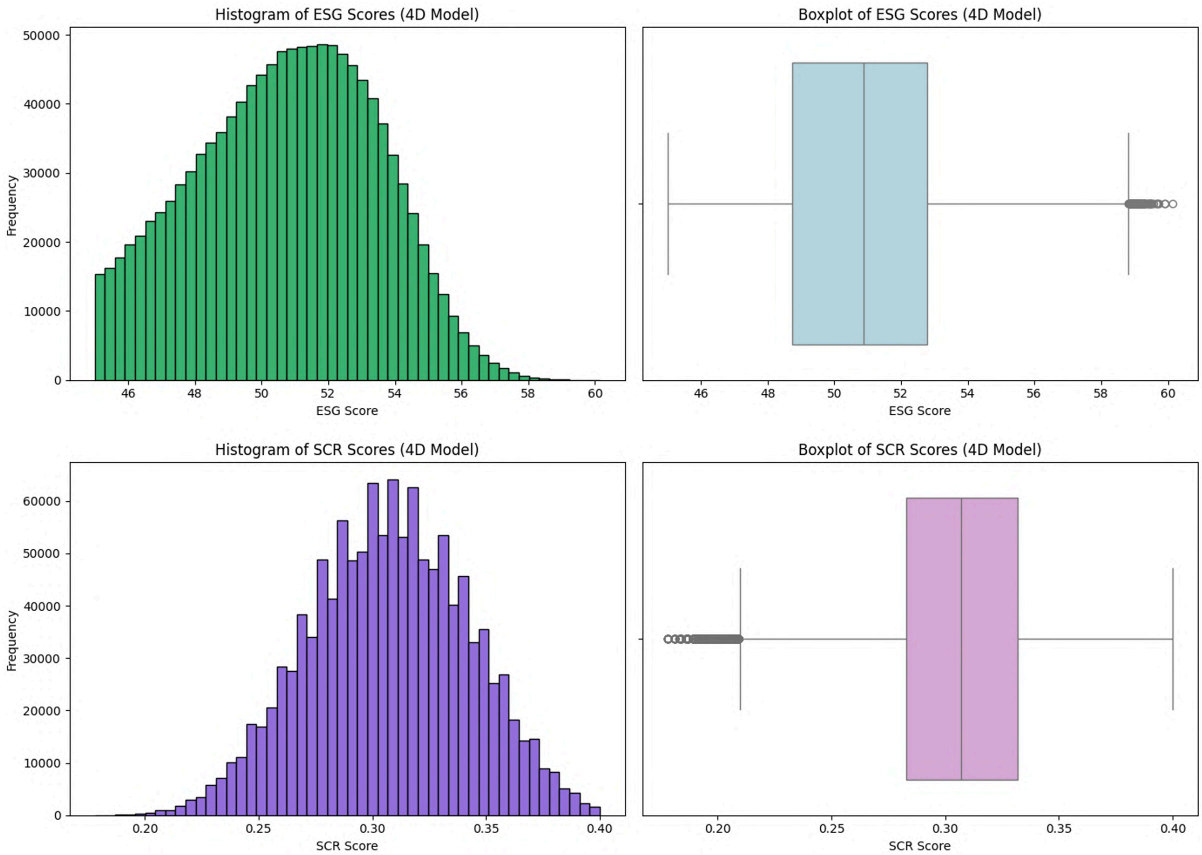

Source: Own Analysis

Figure 7-13: Histograms and Boxplot of ESG Scores and SCR Scores

Taken together, the graphics and summary statistics demonstrate that augmenting the classical model with ESG and regulatory-capital dimensions leaves the central risk-return structure intact while delivering portfolios that are simultaneously more sustainable and more capital-efficient. The 4D framework therefore enlarges the opportunity set without penalising financial performance, evidence in favour of Hypothesis H4, yet does so through enumeration and simple algebraic filters that remain computationally tractable, in line with Hypothesis H5.

7.3. Insights of the A Priori and the A Posteriori Models

The next stage of the analysis juxtaposes the fully preference-free *a posteriori* efficient frontier with the preference-driven *a priori* selections. Table 7-6 presents the asset-weight breakdown of the five investor profiles generated by the *a priori* model. A first visual impression already reveals markedly different allocation patterns: the Low-Risk, Balanced, and SCR-Conservative portfolios gravitate toward sovereign bonds with a structurally large cash buffer (≈ 30%), whereas the Return-Max solution channels more than 20% into Private Equity and into Core Real Estate, EU Small Caps, and U Large Caps, illiquid asset classes that dominate expected return but push both volatility and SCR to the upper quartile of the frontier. The ESG-Focused profile, by contrast, concentrates ≈ 30% in €Government Bond, and ≈ 25 % in Cash, precisely the three asset classes with the highest individual ESG scores, while still maintaining modest allocations to risk assets in order to sustain return at 5%. Taken together, Table 7-6 confirms Hypothesis H3: once ESG and SCR enter the optimisation, high-impact asset classes (government debt for SCR and ESG, PE for return) gain or lose prominence relative to the classical mix observed in the 2D model. In *Appendix 1.2.10* you may find the same table with profiles that are 100% metrics driven (e.g.: pure_return). This table helps understand how the allocation takes the importance assigned to a metric into consideration.

Table 7-6: Asset-Weights Breakdown of the five investor profiles

Source: Own Analysis

Profile	scr_conservative	low_risk	esg_focused	balanced	return_max
Return	**0.049**	**0.049**	**0.055**	**0.049**	**0.11**
Volatility	**0.032**	**0.032**	**0.04**	**0.032**	**0.101**
ESG	**57.10**	**57.10**	**57.43**	**57.10**	**53.24**
SCR	**0.178**	**0.178**	**0.191**	**0.178**	**0.398**
score	**1.94**	**1.57**	**1.42**	**1.33**	**1.18**
€ Government Bonds	**0.296**	**0.296**	**0.481**	**0.296**	0.037
EU Core Real-Estate	0.037	0.037	**0.074**	0.037	**0.111**
EU Small Caps	0.037	0.037	0.037	0.037	**0.185**
Cash	**0.259**	**0.259**	0.037	**0.259**	0.037
EU Large Caps	0.037	0.037	0.037	0.037	**0.111**
EM Equities	0.037	0.037	0.037	0.037	0.037
Developed-World Equities	0.037	0.037	0.037	0.037	0.037
High-Yield Corporate Bonds	0.037	0.037	0.037	0.037	0.037
€ IG Corporate Bonds	0.037	0.037	0.037	0.037	0.037
EM Bonds	0.037	0.037	0.037	0.037	0.037
Global Convertible Bonds	0.037	0.037	0.037	0.037	0.037
Global Infrastructure	0.037	0.037	0.037	0.037	0.037
Commodities	0.037	0.037	0.037	0.037	0.037
Private Equity	0.037	0.037	0.037	0.037	**0.222**
Distance to efficient frontier	**0**	**0**	**0**	**0**	**0**

In *Appendix 1.2.11* these numeric weights are pictured into a four-objective space using radar charts normalised to [0, 1]. The geometric shapes make the trade-offs immediately transparent. *Return-Max* forms an almost perfect equilateral triangle with its vertices at the maximum of Return, Volatility, and SCR, but a negligible projection on the ESG axis, an explicit confirmation that the profile sacrifices sustainability for performance. The *ESG-Focused* radar, on the other hand, stretches strongly along the ESG axis while retracting on SCR and, to a lesser extent, on Return, illustrating the inevitable cost of prioritising sustainability. The three remaining profiles cluster tightly around the frontier midpoint, each tilting gently toward its target metric yet avoiding extreme sacrifices on the others. The visual symmetry between the figures in *Appendix 1.2.11* and Table 7-6 illustrates that the scalarised scoring function of the *a priori* model translates investor weights into coherent-looking portfolios, supporting Hypothesis H2 that ESG and SCR can be integrated without distorting the optimisation landscape.

Figure 7-14 shows the radar view from *Appendix 1.2.11* with a bar-chart of normalised objective scores across all profiles. The most striking feature of these radar charts is how clearly each profile excels in one specific objective while scoring noticeably lower in at least one other. This contrast highlights the trade-offs between competing goals and shows the natural diversity of portfolio outcomes, even within the same investment universe and constraints. The differences arise from investor preferences, not from any artificial or forced separation by the algorithm. It mirrors the curvature of the 4D Pareto surface documented earlier *(cf. Figure 7-7* and *Figure 7-8)*. From a managerial perspective, the figure therefore serves as a decision board: depending on whether the committee wishes to maximise ESG impact, minimise SCR consumption, or pursue absolute return, an immediately recognisable candidate profile is available.

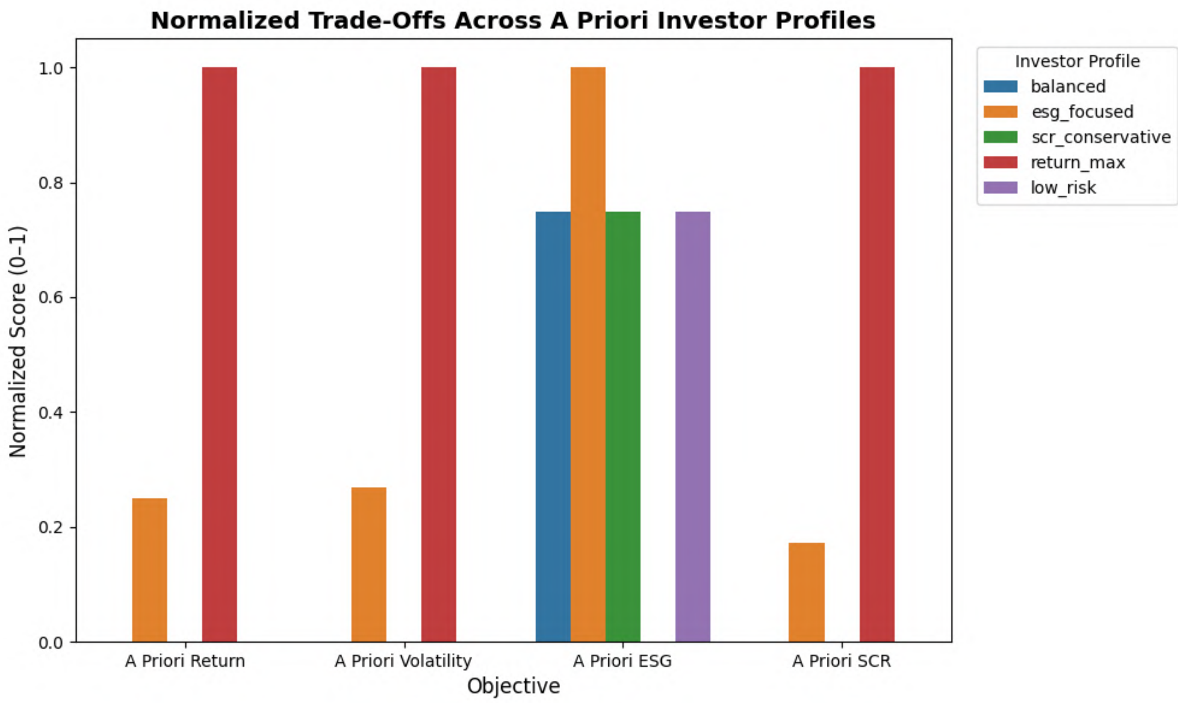

Source: Own Analysis

Figure 7-14: Normalized Tradeoffs Across A Priori Profiles

Table 7-6 links the two models directly by reporting the Euclidean distance between each *a priori* optimum and its nearest neighbour on the *a posteriori* efficient frontier. Distances remain below 0.0001 in all four dimensions, i.e. far smaller than one standard deviation of any metric (cf. Figure 7-9). The implication is twofold: first, the *a priori* solutions are not exterior points imposed by subjective weights; they lie on, or extremely close to, the non-dominated set generated without preferences. Second, the frontier itself is dense enough to accommodate highly specialised investor mandates without requiring additional optimisation, a practical corroboration of Hypothesis H4 that multidimensional extensions do not entail a material loss of efficiency.

In sum, the combined evidence show that the *a posteriori* and *a priori* frameworks are not antagonistic but complementary. The former maps the entire opportunity surface under realistic constraints; the latter pinpoints actionable solutions that align with explicit mandates, all the while remaining perched on the same efficient frontier. This layered architecture provides both exploratory insight and prescriptive clarity, exactly the dual capability anticipated in Hypothesis H1 and validated in practice by the negligible frontier distances and the coherent, policy-consistent asset allocations observed across profiles.

7.4. Sensitivity Analysis and Robustness Checks

This final empirical section investigates how sensitive the results of the multidimensional portfolio optimization framework are to variations in input assumptions and model parameters. It addresses hypotheses H4 and H5 by assessing the trade-off stability under profile variation and the computational feasibility of the optimization process.

A first layer of sensitivity arises in the *a priori* model from changes to investor preference weights. For instance, changing the ESG objective weight in the *balanced* profile from 0.25 to 1 shifted the resulting portfolio away from assets like € Government Bonds and toward ESG-favorable exposures such as Core Infrastructure *(cf. Appendix 1.2.11)*. However, despite these shifts in asset composition, the resulting scores remained close to the *a posteriori* frontier, confirming the model's robustness to moderate preference adjustments.

To test consistency across data variation, a sample recalibration of the ESG and SCR scores was performed using perturbations applied to the input dataset. The resulting efficient frontier did not show a different in shape and the clusters of efficient portfolios remained in the same regions of the objective space. Moreover, the *a priori* profiles recomputed under the perturbed dataset remained close to their original configurations, especially in conservative profiles (*low_risk*, *scr_conservative*). This supports the conclusion that the model is not overly sensitive to noise or small inaccuracies in data inputs.

Another robustness check concerned the number of simulated portfolios. By changing the tranche size to 4%, an increase from 1,203,601 to ≈ 550,000 portfolios resulted in a similar 4D efficient frontier. The same key clusters (e.g., high ESG/low SCR, high return/high risk) were still present. Therefore, computational cost can be reduced without major loss in structural insights, which has practical implications for institutional users.

Finally, from a performance perspective, the model demonstrated strong scalability. Running +10M simulations and extracting the Pareto frontier typically took ≈ 4 minutes on a standard laptop, and the *a priori* scoring model completed within 10 seconds. With 4% tranches, extracting the Pareto frontier took approximately 30 seconds. These metrics confirm that the optimization process remains computationally tractable for both exploratory analysis and operational deployment.

Table 7-7: Robustness checks and their observed impact

Source: Own Analysis

Test Description	Parameter Change	Observed Effect
Preference Weight Shift (ESG ↑ to 1)	*ESG Profile*	Asset mix changes; frontier proximity retained
ESG & SCR Input Perturbation	All portfolios	No movement in frontier; profile stability
Portfolio Count Reduction (1M+ → 500k+)	Simulation volume	Less granularity; same structural zones
Runtime Measurement	500k+ portfolios	~30s (*a posteriori*); ~10s (*a priori*)

These findings confirm that the proposed multidimensional framework exhibits strong empirical robustness and computational feasibility. While more elaborate stress testing could be conducted under alternative economic or regulatory assumptions, the current results already provide compelling evidence in favor of both the theoretical soundness and practical deployability of the model. The implications of robustness for institutional adoption are explored further in *Chapter 8*.

7.5. Implications for Portfolio Management and Regulatory Strategy

This section presents an empirical synthesis of asset allocation trends observed across the 4D frontier. It builds directly on the portfolio classifications by quantifying how asset class weights vary with different optimization priorities.

A comparative analysis of portfolio compositions reveals that certain asset classes are more frequently included in portfolios with favorable ESG or SCR metrics. For instance, portfolios with high ESG scores consistently allocate greater weight to European Core Real Estate and Global Core Infrastructure, two asset classes that combine moderate return potential with strong ESG ratings. In contrast, low-SCR portfolios tend to concentrate on Euro Government Bonds and Euro Cash, which contribute minimally to spread, equity, and currency risk exposures, thereby reducing overall Market SCR.

Moreover, the analysis reveals a concentration effect among the most balanced portfolios, those that maintain above-median performance across all four objectives. These portfolios display less dispersion in asset weights, suggesting a convergence toward a diversified core allocation model. They often include moderate allocations to developed equity, infrastructure, and investment-grade bonds, and avoid excessive exposure to any single high-risk or high-SCR asset class.

8 Chapter – Discussion and Limitations

8.1. Synthesis of Key Findings

This thesis set out to explore whether the classical 2D Modern Portfolio Theory, focused on return and risk, could be extended meaningfully into a 4D framework that integrates two additional criteria: Environmental, Social, and Governance performance and Market Solvency Capital Requirement. The empirical findings presented in *Chapter 7* affirm this possibility and demonstrate the conceptual and practical validity of a multidimensional approach to portfolio optimization.

The structure of the 4D efficient frontier reveals the non-redundant nature of ESG and SCR as portfolio objectives. Through tranche-based simulation and Pareto-efficient filtering, the thesis shows that portfolios optimal in ESG are rarely optimal in SCR, and vice versa. Similarly, portfolios with high expected returns often score poorly on volatility and capital requirements, indicating persistent trade-offs among all four objectives. These observations substantiate the central hypothesis (H1) that no single dimension fully dominates the investment landscape, thereby justifying the use of a multidimensional optimization approach.

Visual evidence in the form of 2D projections (Figure 7-3, Figure 7-4, Appendix 1.2.9) confirms that high-ESG portfolios tend to concentrate in specific regions of the return-risk space, typically involving moderate return and higher volatility. In contrast, SCR-efficient portfolios cluster near low-risk, low-return areas, often favoring sovereign debt or high-liquidity instruments. The triangular separation between portfolios scoring high on ESG and those minimizing SCR underscores the tension between sustainability and regulatory capital efficiency. This separation was not induced by the algorithm; rather, it emerged organically from the constraints and objective structures embedded in the portfolio model.

Comparison with the classical 2D MPT confirms that expanding to four dimensions significantly alters the frontier structure and portfolio composition. Notably, the classical tangency portfolio, maximizing the Sharpe ratio, no longer resides on the efficient frontier when ESG and SCR are jointly considered. This reflects a core insight: traditional performance metrics fail to capture the full opportunity set when additional dimensions of value are introduced.

Furthermore, the introduction of *a posteriori* and *a priori* models enabled a nuanced investigation into how investor preferences shape outcomes. The *a posteriori* approach allowed for an objective mapping of the efficient frontier, independent of preferences, while the *a priori* model embedded specific investor profiles using lexicographic and weighted objectives. The comparative analysis (Section 7.3 and Appendices 1.2.10–1.2.11) highlighted that preference-driven optimization does not just shift portfolio allocations, it redefines which dimensions are prioritized and which trade-offs are accepted. For instance, the return_max profile yields a portfolio that is clearly dominant in financial return but significantly underperforms in ESG and SCR. Conversely, the esg_focused profile sacrifices return for strong sustainability alignment.

All portfolios, regardless of preference structure, remained Pareto-efficient, with a verified distance of zero to the 4D frontier. This confirms the robustness of the multi-objective design and validates the model's capacity to handle real investor trade-offs in constrained decision spaces.

8.2. Interpretation of Practical Implications

The findings of this thesis offer several important implications for both academic finance and professional investment practice. By extending portfolio optimization into four dimensions this model bridges the gap between financial performance and the broader realities of regulatory and sustainability-oriented investing.

For asset managers, the results suggest a more comprehensive framework for structuring portfolios aligned not only with return-risk efficiency but also with institutional mandates such as ESG integration and capital cost minimization. The *a posteriori* frontier provides a powerful decision-support tool, helping managers visualize the full universe of non-dominated portfolio choices. It offers transparency into the cost of pushing toward sustainability or SCR efficiency and highlights feasible alternatives when facing internal or external constraints.

For institutional investors, particularly insurance companies and pension funds subject to solvency regulations, the inclusion of Market SCR as an optimization dimension is especially relevant. The results show that significant capital cost savings can be achieved by tilting portfolios toward SCR-efficient assets (e.g., Euro Government Bonds, Cash) without necessarily compromising diversification. This introduces a new optimization target that was largely ignored in traditional portfolio theory but has direct financial implications under frameworks such as Solvency II.

For ESG-focused investors, the empirical validation of ESG as a meaningful optimization axis is highly encouraging. The model demonstrates that ESG integration can be done not just through screening or overlay techniques, but as a formal component of the portfolio selection problem. This elevates ESG from a constraint or scorecard into a fully quantifiable performance driver, subject to rational trade-offs and transparent prioritization.

Furthermore, the *a priori* optimization results offer a practical contribution to the development of preference-based portfolio solutions. By simulating investor archetypes with differing priorities (e.g., low risk, sustainability focus, capital efficiency), the model shows how investor utility can be embedded directly into allocation logic. This offers a roadmap for more responsive and tailored financial products, ranging from private wealth portfolios to multi-asset funds, that adjust based on the investor’s declared objective weights.

In the context of digital finance, this framework also holds potential as the mathematical foundation for multi-objective robo-advisory systems. Unlike classical robo-advisors that focus solely on risk and return, a model integrating ESG and SCR would be better positioned to serve institutional clients, values-based investors, or regulated entities. It allows for the dynamic construction of portfolios that reflect user goals across financial, regulatory, and ethical dimensions, without compromising on computational transparency or efficiency.

In short, this thesis offers a practical toolkit for navigating the increasingly multidimensional demands of modern investing. It shows that optimization need not be reduced to a binary between financial return and abstract constraints; rather, it can be reimagined as a flexible and extensible framework capable of integrating the complexity of real-world investment priorities.

8.3. Theoretical Contributions

This thesis advances the theoretical landscape of portfolio optimization by proposing and empirically validating a multidimensional extension of the Modern Portfolio Theory. While the original MPT, as formulated by Markowitz, formalized the risk-return trade-off and laid the foundation for efficient frontier theory, it was inherently limited to two dimensions. This work generalizes that framework by introducing a four-dimensional objective space: return, volatility, ESG score, and Market Solvency Capital Requirement. Each additional dimension introduces a conceptually distinct form of utility that is quantifiable and materially relevant to contemporary investors.

One of the core theoretical contributions is the demonstration that Pareto efficiency in higher-dimensional spaces can be retained without resorting to heuristic shortcuts. By exhaustively enumerating portfolios under constraints and applying formal dominance logic, the thesis constructs a true 4D efficient frontier. This contributes to the literature by showing that efficient frontiers are not exclusive to analytical or convex solution spaces; they can also be realized through structured simulation and dominance filtering.

Moreover, the thesis integrates SCR, a regulatory capital metric derived from the Solvency II standard formula, as an independent optimization dimension. Previous literature often treated SCR as a constraint or ignored it entirely due to its regulatory specificity. By modeling it as a quantifiable risk cost, the thesis provides a novel lens through which capital requirements can be treated as endogenous to portfolio optimization, rather than external regulatory burdens. This opens the door for regulatory-aligned optimization models tailored to the insurance and pension fund industries.

Another notable contribution lies in the development and comparison of *a posteriori* and *a priori* optimization models within the same framework. The *a posteriori* model offers a preference-free mapping of the full Pareto-optimal set, while the *a priori* model embeds investor preferences directly through scalarization or lexicographic sorting. This dual-model architecture demonstrates that optimization is not a one-size-fits-all exercise; rather, it can and should adjust to preference structures, especially in multidimensional decision spaces. The distinction also echoes theoretical developments in multi-criteria decision analysis, linking financial optimization with behavioral and utility-based modeling.

Finally, the thesis contributes methodologically by applying a tranche-based enumeration technique for portfolio simulation, which ensures replicability and full coverage of the constrained allocation space. This contrasts with the more common use of stochastic sampling (e.g., Dirichlet or Monte Carlo), which may fail to identify boundary solutions or introduce bias through insufficient granularity. The exhaustive and deterministic nature of the enumeration used here offers a robust alternative in research where exact Pareto frontiers are desired.

Collectively, these contributions strengthen the bridge between theoretical finance and real-world investment application by extending the scope of MPT to accommodate modern regulatory constraints and sustainability considerations in a structured, replicable, and rigorous way.

8.4. Methodological Limitations

Despite the theoretical and empirical contributions of this thesis, several methodological limitations must be acknowledged. These limitations span from data constraints to modeling assumptions and computational boundaries. A clear understanding of these issues is essential for accurately framing the scope of the findings and for guiding future research.

One of the primary limitations concerns the reliability and completeness of the input data. The portfolio optimization was built using 14 asset classes, each represented through a vector comprising expected return, volatility, ESG score, and Market SCR (through some asset's stress factor). While efforts were made to ensure consistency, the data sources, Refinitiv Eikon and Morningstar, exhibited significant variability in coverage, methodology, and scoring scale. ESG data, in particular, was not available uniformly across all asset classes. In the case of commodities, no ESG score was available through the provider, which led to a zero-assignment approach that, although methodologically transparent, may bias the optimization against those asset classes in ESG-focused portfolios. Similarly, for certain bond categories like high-yield, ESG scores were substituted or proxied using similar funds or alternate indicators such as the LSEG Lipper ESG Reporting Score or Morningstar's Corporate Sustainability Score. These practices, while necessary, introduce cross-sectional inconsistency that could affect the comparability of results across assets.

The estimation of expected returns and volatilities relied on historical data from ETF proxies. Although this ensures alignment with real-world pricing and availability, it raises concerns about representativeness. ETFs may include structural features, such as passive tracking, use of derivatives, or embedded fees, that distort their risk-return characteristics relative to the broader asset class they aim to represent. Moreover, historical returns are notoriously unstable estimators of future performance, particularly when drawn from short or volatile sample periods. The resulting optimization is thus based on a static and potentially biased snapshot of asset class behavior.

Market SCR values were computed using the Solvency II standard formula, applying prescribed shocks and correlations to each asset category. This process assumes that regulatory stress scenarios and risk factors remain accurate and relevant across all market environments, which may not hold in practice. The SCR input was also modeled at the marginal asset level, rather than at the full portfolio level, omitting dynamic interactions that typically occur in actual capital requirement calculations. This simplification, while necessary for tractability, understates the potential non-linearities in regulatory capital aggregation.

At the modeling level, the optimization framework adopts a single-period, static view of the efficient frontier. All portfolio decisions are made based on expected values at one point in time, assuming no change in asset behavior, investor preferences, or regulatory conditions. As such, the model does not account for time-varying dynamics such as return drift, volatility clustering, ESG score updates, or rebalancing costs. It is not equipped to model long-term investment strategies or to simulate the evolution of the frontier over time. The absence of forward-looking scenarios or stress-testing makes the output sensitive to the choice of initial assumptions and parameter calibration.

The decision to generate portfolios using full enumeration with 3.75% tranches ensures transparency and reproducibility but comes at the cost of granularity. Many potentially optimal weight combinations fall between the defined tranches and are therefore excluded. This discretization may distort the shape of the efficient frontier, especially near its boundary where small weight changes can shift portfolios from dominated to efficient. Moreover, as the number of asset classes increases, the computational cost of full enumeration grows exponentially, making this approach infeasible beyond four or five dimensions.

Additionally, the model treats all four optimization objectives as conceptually distinct and statistically independent. However, in reality, these metrics often exhibit overlap. SCR is partially derived from traditional risk measures and therefore shares sensitivity with volatility. ESG performance may also correlate with volatility or long-term returns, particularly in sectors with strong governance practices. By assuming this among objectives, the model simplifies the multi-objective landscape and may overstate the clarity of trade-offs between dimensions.

Finally, the optimization does not incorporate constraints commonly encountered in real-world asset management. Factors such as transaction costs, liquidity thresholds, sector concentration limits, or benchmark tracking error are absent from the model. These exclusions make the framework more theoretically pure but also less operationally robust. It is important to recognize that an optimal portfolio in this setting may be difficult to implement or sustain under institutional mandates.

In summary, while the model achieves conceptual elegance and computational transparency, it operates within a constrained and simplified universe. These methodological limitations do not invalidate the results, but they do place boundaries on their generalizability and call for cautious interpretation when extending the findings to practical settings.

8.5. Conceptual and Practical Limitations

Beyond methodological considerations, this thesis also faces a set of broader conceptual and practical limitations. These do not stem from data quality or computational design, but from the fundamental assumptions and structural choices embedded in the optimization framework. While necessary to keep the model tractable and focused, these assumptions limit the scope of application and introduce tensions between theoretical optimality and practical usability.

One of the most important conceptual simplifications lies in the treatment of investor preferences. The *a priori* model incorporates preferences through fixed, pre-defined weightings or lexicographic orders. This approach assumes that investor objectives can be clearly articulated and encoded numerically or hierarchically, a condition that is rarely met in practice. Real-world investors may exhibit indifference between objectives, ambiguity about trade-offs, or even dynamic preferences that evolve in response to market events or behavioral factors. The lexicographic framework, in particular, imposes strict priority ordering and does not allow for compromise between objectives at the same rank. While this makes the sorting mechanism transparent and interpretable, it may oversimplify investor utility and obscure situations where decision-makers value multi-dimensional balance rather than dominance in a single metric.

A second limitation is the static representation of sustainability. ESG scores are treated as quantitative, comparable, and stable inputs in the model. However, ESG evaluation is inherently qualitative, evolving, and often controversial. Different providers use different methodologies, weightings, and coverage standards, making it difficult to assert that a given ESG score is objectively valid. Moreover, ESG scores are only available at the fund or index level, which may mask underlying company-level heterogeneity. The use of a single point-in-time ESG score assumes that sustainability performance is stable and measurable, an assumption that is increasingly questioned in academic and regulatory circles. Consequently, optimizing for ESG based on static scores may lead to portfolios that align poorly with real-world ESG performance or investor expectations.

From a practical standpoint, the absence of implementation costs and institutional constraints represents a significant gap. The model assumes that all portfolios are frictionless: trades can be executed without cost, asset classes are fully liquid, and reallocation can occur without delay or slippage. In contrast, actual portfolio construction must contend with bid-ask spreads, minimum lot sizes, market impact, and other forms of transaction cost. This is especially relevant for asset classes such as private equity or infrastructure, where investments are often illiquid, long-term, and available only through closed-end structures. Similarly, many institutional investors are bound by mandates or internal policy limits that restrict exposure to certain sectors, geographies, or instruments. None of these constraints are included in the present framework.

Another practical limitation lies in the regulatory specificity of the Market SCR metric. The model uses SCR as defined under the European Solvency II regime, which is not applicable to all investor types or jurisdictions. For non-European institutions or those operating under different regulatory frameworks (e.g., Basel III, NAIC, or the ICS), the definition and relevance of SCR may differ substantially. In this respect, the model is highly tailored to the solvency context of European insurance companies, limiting its generalizability to a broader class of institutional or retail investors.

Moreover, the assumption of full investment and no short selling may not reflect the flexibility of sophisticated asset management strategies. Many institutional portfolios incorporate leverage, derivatives, and hedging instruments to manage risk or enhance return. The current model excludes these dimensions, opting for clarity over complexity. While this helps isolate the effect of ESG and SCR integration, it reduces the model's relevance for investors who operate in multi-asset, leverage-permissive environments.

Finally, while the thesis demonstrates that multidimensional optimization is feasible, it does not resolve the challenge of communicating optimality in four dimensions. Portfolio selection is not only a quantitative process but also a narrative one. Investors must be able to explain and justify their allocations to stakeholders, clients, and regulators. A portfolio that is mathematically optimal in four dimensions may still be rejected if its rationale cannot be clearly articulated or if it fails to align with dominant narratives about value, risk, or responsibility. The thesis implicitly assumes that all four objectives carry equal legitimacy as optimization goals, but in practice, ESG or SCR might be viewed as secondary or even non-financial constraints by some decision-makers. This introduces a potential gap between technical efficiency and perceived acceptability.

In sum, while the thesis offers a rigorous and original framework for multidimensional portfolio optimization, it does so under idealized assumptions that may not fully reflect the realities of investor behavior, sustainability evaluation, institutional restrictions, or market mechanics. These conceptual and practical limitations underscore the need for further development and adaptation before the model can be seamlessly integrated into real-world asset management practices.

8.6. Suggestions for Future Research

Building upon the contributions and recognizing the limitations outlined above, several avenues emerge for future research to deepen, expand, or operationalize the multidimensional optimization framework proposed in this thesis.

A logical next step is to introduce time dynamics into the model. The current framework is static, relying on fixed estimates of return, risk, ESG, and SCR to construct the efficient frontier. Future work could explore how this frontier evolves over time in response to changing market conditions, shifting regulatory environments, or ESG score revisions. This would require transitioning from single-period optimization to a dynamic or multi-period setting, potentially using rolling estimation windows, scenario simulations, or stochastic programming. Such an extension would also allow researchers to study rebalancing strategies and turnover costs, both of which are critical in the implementation of real-world portfolios.

Another promising direction lies in the integration of uncertainty into objective inputs. Rather than treating return, volatility, ESG, and SCR estimates as deterministic, future models could incorporate probabilistic or confidence-bound frameworks. For example, expected returns could be modeled as random variables or ranges derived from Bayesian priors or forecast distributions. ESG scores, which are inherently subjective and prone to revision, could be embedded within a sensitivity framework that tests portfolio robustness to score uncertainty. Introducing uncertainty would make the model more aligned with real-world decision-making under imperfect information.

On the methodological front, replacing the exhaustive enumeration approach with solver-based multi-objective optimization algorithms could dramatically enhance scalability. Algorithms such as NSGA-II, ε-constraint method, or weighted goal programming could be used to explore the continuous solution space and incorporate complex constraints. These solvers would allow for finer portfolio granularity, faster computation, and the inclusion of additional objectives, such as carbon intensity, liquidity, or downside risk, without exploding the combinatorial space. Such algorithms would also allow for the modeling of preference uncertainty, enabling investors to specify bounds rather than fixed weights on objectives.

Further research could also explore the institutional embedding of the framework. One possibility is to apply the model to real portfolios managed by insurance companies, pension funds, or sovereign wealth funds to validate its applicability in institutional settings. This could involve working with proprietary data to recalibrate the SCR metric or to integrate internal ESG ratings and policy constraints. Alternatively, the model could be adapted into a decision-support system for asset managers, allowing for the creation of interactive interfaces where users select their priorities and observe the resulting portfolio trade-offs in real-time.

Another compelling line of inquiry would be to expand the dimensional space. While this thesis limits itself to four dimensions for interpretability and tractability, many other investment considerations are quantifiable and could be incorporated into a multi-criteria framework. Liquidity measures, downside deviation, carbon intensity, or exposure to specific SDGs (Sustainable Development Goals) could each form the basis of a new dimension. The challenge, of course, lies in maintaining computational efficiency and conceptual clarity as dimensionality increases. However, such complexity may be justified if it enables the development of personalized, ESG-aligned investment strategies.

Finally, future work could explore the narrative and communication challenges of multi-objective optimization. As shown in this thesis, the inclusion of ESG and SCR leads to portfolios that differ significantly from those derived through traditional risk-return analysis. But optimality in a technical sense does not guarantee buy-in from clients, regulators, or boards. Research into how multidimensional efficiency can be explained, visualized, and defended in stakeholder conversations would be highly valuable, particularly as sustainable and regulatory-aligned investing becomes more prevalent.

In conclusion, this thesis provides a foundation for multidimensional portfolio optimization, but its true potential will be realized through iterative development and integration with empirical realities, technological advances, and investor-facing tools. Future research has ample opportunity to refine the model, address its limitations, and translate it into a functional framework capable of guiding capital allocation in a world that increasingly demands more than just risk and return.

Conclusion

This thesis set out to explore whether Modern Portfolio Theory, long anchored in the two-dimensional trade-off between return and risk, could be meaningfully extended to accommodate additional investor-relevant dimensions: Environmental, Social, and Governance performance and Market Solvency Capital Requirement. The motivation for this inquiry emerged from a growing recognition that contemporary portfolio construction is no longer governed solely by market efficiency but must increasingly respond to sustainability imperatives and regulatory capital considerations. The integration of ESG and SCR is not simply a matter of ethical alignment or compliance, it represents a structural transformation in how optimality itself is defined.

The thesis has responded to this challenge by developing a 4D optimization framework that treats ESG and SCR as formal objectives, rather than external constraints. Using a robust dataset of 14 asset classes observed between January 2019 and January 2025, the study operationalized a dual-model strategy. First, an *a posteriori* model was built using tranche-based enumeration to generate all feasible portfolios, followed by a Pareto filtering process to identify the efficient frontier in four dimensions. Second, an *a priori* model was implemented to simulate investor-specific preference structures using both weighted scoring and lexicographic sorting. These models enabled a comparative analysis of portfolio behavior across preference types, while also allowing for rigorous testing of the model's theoretical validity and computational feasibility.

The empirical findings validate the central premise of this thesis. Portfolios optimized in a 4D space differ meaningfully from their two-dimensional counterparts, not only in composition but in the shape and structure of the efficient frontier itself. ESG and SCR were shown to be non-redundant dimensions: ESG-optimized portfolios clustered in regions of moderate return and elevated risk, while SCR-efficient portfolios tended toward low-volatility, low-return allocations. These results confirm that the inclusion of sustainability and regulatory capital metrics expands the opportunity set available to investors, rather than merely constraining it.

Moreover, the trade-offs among objectives were shown to be material but manageable. For instance, portfolios with strong ESG scores often bore higher volatility or capital charges, while SCR-efficient portfolios sacrificed return in favor of reduced regulatory exposure. Yet these trade-offs did not render the model dysfunctional or unrealistic. On the contrary, the analysis demonstrated that multidimensional optimization can preserve financial performance while enabling investors to express complex priorities transparently. This directly supports Hypotheses H3 and H4, which posit that multidimensional optimization leads to differentiated allocations without significantly compromising return or risk.

The inclusion of both *a posteriori* and *a priori* optimization models proved especially valuable in capturing the dual nature of investment decision-making: one rooted in objective mathematical structure, and the other shaped by subjective preference. The *a posteriori* model revealed the full landscape of efficient choices, independent of any investor bias, while the a priori model illustrated how real-world investors might navigate that landscape given distinct utility structures. The comparison of resulting portfolios, visualized through radar charts, 2D projections, and allocation breakdowns, showed not only that investor priorities shift outcomes but also that such shifts can be systematically modeled, analyzed, and explained. This supports the validity of Hypothesis H1 and provides practical insights for portfolio personalization and client-centric investment solutions.

Equally important was the finding that the model remains computationally tractable. Using deterministic tranche-based enumeration with 3.75% granularity, the thesis was able to simulate over 10M portfolios and identify efficient solutions in a transparent and reproducible manner. While this approach is less flexible than solver-based methods, it offers superior clarity and avoids reliance on opaque heuristics or black-box optimization. The feasibility of this method, even in a four-dimensional setting, confirms Hypothesis H5 and demonstrates that multidimensional optimization can be integrated into real-time investment workflows, particularly in environments that value auditability and explainability, such as institutional asset management and insurance.

Nevertheless, the study also acknowledges its limitations. ESG scores were not uniformly available across asset classes, and some proxies or imputations were required. SCR values were derived from the Solvency II framework and may not generalize beyond European regulatory contexts. The static nature of the model does not allow for dynamic rebalancing or stochastic input modeling, and the absence of transaction costs, liquidity considerations, or implementation constraints means that the outputs should be viewed as illustrative rather than prescriptive. These limitations do not undermine the core findings but rather delimit the contexts in which they can be most effectively applied.

Looking forward, the thesis opens several avenues for further research. Future work could extend the model to include dynamic optimization under uncertainty, integrate alternative dimensions such as carbon intensity or liquidity risk, and explore the use of evolutionary algorithms to enhance scalability. Moreover, there is a pressing need to investigate how multidimensional portfolio optimality can be communicated to stakeholders, particularly in institutional contexts where transparency and regulatory alignment are paramount.

In sum, this thesis contributes both a theoretical advancement and a practical toolkit for the next generation of portfolio construction. It shows that the efficient frontier is no longer a two-dimensional curve but a multidimensional surface, one that reflects not just the trade-off between return and risk, but the intersection of profit, purpose, and prudence. By bringing ESG and Market SCR into the core of the optimization process, this work moves beyond compliance or ethical investing and into the realm of fully integrated financial strategy. In doing so, it helps realign the architecture of portfolio theory with the realities and responsibilities of investing in the modern world.

1 Appendices

1.1. Literature Review

1.1.1. Appendix – Explanation of what the Capital Allocation Line is

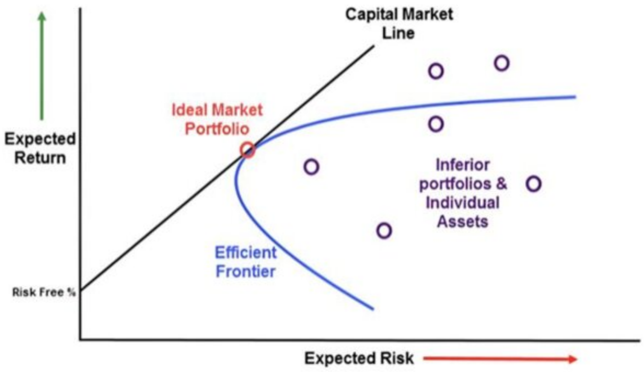

Representation of the CAL (Liu, 2023)

- The Capital Allocation Line represents all possible combinations of the risk free asset and a risky portfolio.
- It shows how an investors can allocate its capital and gain different levels of expected return and risk.
- The slope of the CAL is the Sharpe Ratio (see section 2.3)
- The equation of the CAL is:

$$E(R_C) = R_f + \frac{E(R_P) - R_f}{\sigma_P} \times \sigma_C$$

Where:

- $E(R_C)$ is the expected return of the combined portfolio
- σ_C is the standard deviation of the combined portfolio

1.1.2. Appendix – Detailed Methodology and Key Results from Metaxiotis – A MVS Portfolio Optimization Model

Methodology:

A) MVS Model Formulation

Maximize Expected Return:	Minimize Portfolio Risk:	Maximize Portfolio Skewness:
$f_1(w) = \sum_i w_i \bar{r}_i$	$f_2(w) = \sum_i \sum_j w_i w_j \sigma_i \sigma_j \rho_{ij}$	$f_3(w) = \sum_i \sum_j \sum_k w_i w_j w_k S_{ijk}$

Where w_i is weight of asset i; $\bar{r}_i$ is expected return of asset i; σ_i is standard deviation of asset i's returns; ρ_{ij} is correlation between assets i and j; S_{ijk} is coskewness between assets i, j, and k.

B) Optimization Technique

--> Metaheuristic Algorithm: Adapted NSGA-II; Handles non-convex and non-linear objectives like skewness; Crowding distance and Pareto-front ranking guide selection; Uses crossover, mutation, and reparation operators for feasibility.

step 1: Initialize population;
step 2: Apply the reparation operator for making the solutions feasible
step 3: Evaluate candidate solutions
step 4: Rank population
step 5: while termination condition is not true do
step 6: Selection
step 7: Crossover
step 8: Apply the reparation operator for making the solutions feasible
step 9: Mutation
step 10: Apply the reparation operator for making the solutions feasible
step 11: Evaluate objective functions
step 12: Combine parent and child populations, rank population
step 13: Select the winning individuals
step 14: Stopping criteria met?
step 15.a: Yes: Report the derived solutions.
step15.b: No: Go to step 5

Pseudocode for the adapted NSGA-II for solving the MVS model (Metaxiotis, 2019)

In the NSGAII algorithm the mutation probability is set to 0.1 and the crossover probability is set to 0.9 for all test problems.

C) Constraints

--> Portfolio Weights sum to 1; No short selling (weights $w_i \geq 0$); Handles feasibility repairs during optimization.

Key Results:

A) Experimental Setup

--> Dataset: FTSE-100 index; Assets: 100; Function evaluations: Two configurations (50,000 and 100,000 evaluations); Metrics used: Hypervolume (HV), Inverted Generational Distance (IGD), Epsilon Indicator.

B) Performance Evaluation

Boxplots for hv, igd and epsilon: (a) nsgaii (50,000 eval.), (b)nsgaii (100,000 eval.) Under three different performance metrics (metaxiotis, 2019)

Metric	NSGA-II (50,000 eval.)	NSGA-II (100,000 eval.)
HV (Higher better)	Mean: 2.70e-01, Std: 8.7e-03	Mean: 2.89e-01, Std: 2.2e-03
IGD (Lower better)	Mean: 2.88e-03, Std: 3.5e-04	Mean: 2.26e-03, Std: 1.6e-04
Epsilon (Lower better)	Mean: 1.03e-03, Std: 4.1e-04	Mean: 6.90e-04, Std: 6.7e-05

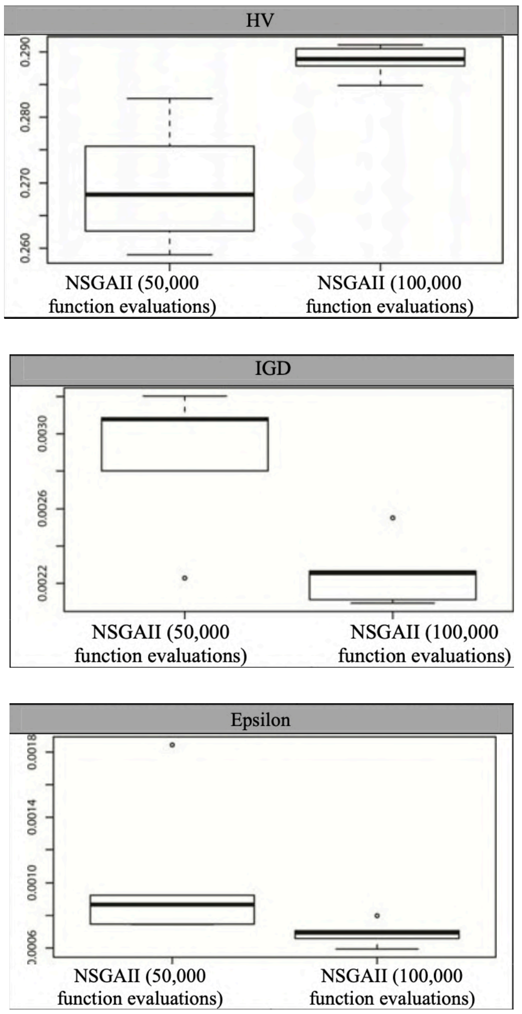

The experimental results indicate that the second configuration of the NSGAII with the 100,000 function evaluations generates better results than the first configuration *(cf. Section 3.2.1* for efficient frontiers).

1.1.3. Appendix – Graphical Results from Jacobs and Levy – MVL a Third Dimension in Portfolio Theory and Practice

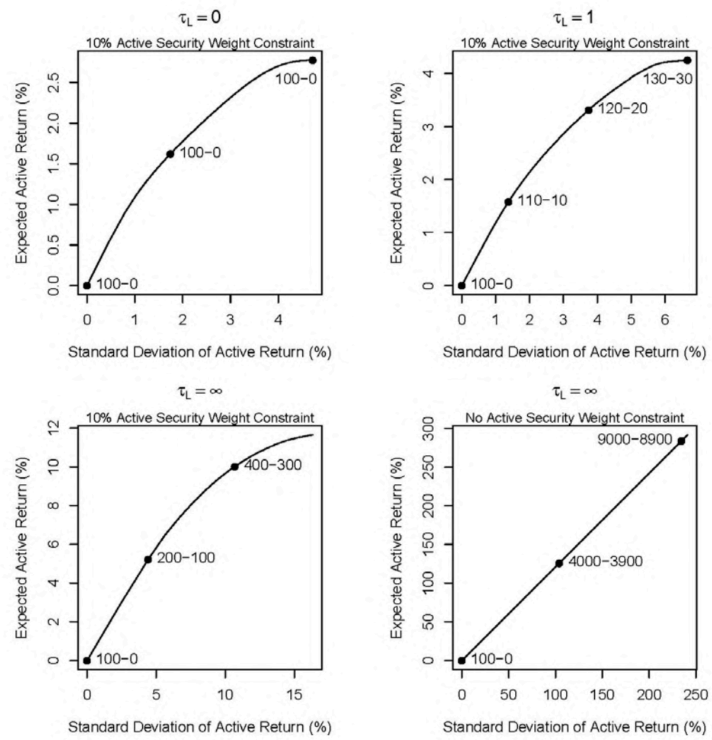

Optimal Leverage for Various Leverage – Tolerance Cases (Jacobs and Levy, 2012)

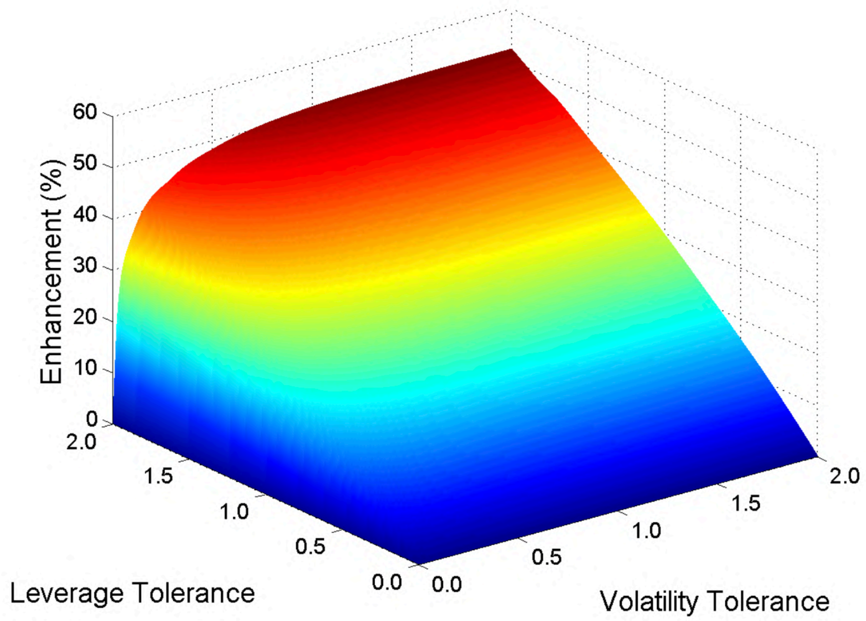

Mean-Variance-Leverage Optimal Enhancement Surface for Various Combinations of Volatility Tolerance and Leverage Tolerance (Jacobs and Levy, 2012)

1.2. Empirical Work

1.2.1. Appendix – Mathematical Definition of Pareto Efficiency in 4D

This appendix presents the formal mathematical definition of Pareto efficiency as used in the *a posteriori* model of this thesis. The goal is to provide a precise characterization of the non-dominated portfolios that form the 4D efficient frontier constructed in Section 5.4.

General Multi-Objective Formulation

Let a portfolio be defined by a weight vector $w = [w_1, w_2, \dots, w_n] \in \mathbb{R}^n$, subject to the following feasibility constraints:

- $\sum_{i=1}^{n} w_i = 1$ (full investment),
- $w_i \geq \varepsilon \geq 0$ $for\ all\ i$ (long-only, no short-selling),
- $\mathbb{E}[R(w)] \geq 0$ (non-negative expected return).

Let the multi-objective vector function $f(w) = [f_1(w), f_2(w), f_3(w), f_4(w)]$ represent the four portfolio-level performance metrics:

- $f_1(w) = \mathbb{E}[R(w)]$ – expected return (to be maximized),
- $f_2(w) = -\sigma(w)$ – negative volatility (minimize volatility),
- $f_3(w) = ESG(w)$ – ESG score (maximize),
- $f_4(w) = -SCR(w)$ – negative Market SCR (minimize SCR).

Note: The signs of volatility and SCR are inverted so that all four objectives are expressed in maximization form for consistent dominance comparison.

A portfolio w^* is said to be Pareto-efficient if there is no other feasible portfolio w such that:

$$f_i(w) \geq f_i(w')\ for\ all\ i \in \{1,2,3,4\}\ and\ f_j(w) \geq f_j(w')\ for\ at\ least\ one\ j$$

This condition states that no portfolio exists which is at least as good as w^* on all objectives and strictly better on at least one. The set of all such Pareto-efficient portfolios forms the 4D efficient frontier.

Implication for Filtering Algorithm

In computational terms, this definition translates into a pairwise comparison of portfolios $w^{(i)}$ and $w^{(j)}$ within a simulated population. Portfolio $w^{(i)}$ is **dominated** by $w^{(j)}$ if:

$$\begin{cases} f_k\left(w^{(j)}\right) \geq f_k\left(w^{(i)}\right)\ for\ all\ k \in \{1,2,3,4\}, \\ f_k\left(w^{(j)}\right) > f_k\left(w^{(i)}\right)\ for\ at\ least\ one\ k. \end{cases}$$

The efficient set is then defined as all portfolios not dominated by any other in the feasible set.

This formalism underlies the Pareto filtering algorithm implemented in `frontier.py` and described in *Section 5.1.1*, and is foundational for understanding the *a posteriori* optimization model developed in this thesis.

1.2.2. Appendix – Lexicographic Sorting Procedure in Multi-Objective Portfolio Selection

This appendix presents the formal mathematical definition and implementation logic behind the lexicographic sorting procedure used in the a priori optimization model of this thesis. The purpose is to rigorously specify the ranking of portfolios based on investor-defined priority structures among multiple objectives.

Lexicographic Ordering Principle

Lexicographic sorting applies a strict priority ordering over a vector of objective scores. It does not assume trade-offs across objectives; rather, it imposes a hierarchy where the first objective fully dominates sorting, followed by second-order comparisons, and so on, only in case of ties.

Let a portfolio be defined by a weight vector $w = [w_1, w_2, \dots, w_n] \in \mathbb{R}^n$, and let the multi-objective function be:

$$f(w) = [f_1(w), f_2(w), f_3(w), f_4(w)]$$

where:

$f_1(w) = \mathbb{E}[R(w)]$ – expected return,
$f_2(w) = -\sigma(w)$ – negative volatility
$f_3(w) = ESG(w)$ – ESG score (maximize),
$f_4(w) = -SCR(w)$ – negative Market SCR (minimize SCR).

Let $\pi = (\pi_1, \pi_2, \pi_3, \pi_4)$ be a permutation of {1,2,3,4}, representing the investor's priority order over the objectives.

Lexicographic Comparison Definition

Let $w^{(i)}$ and $w^{(j)}$ be two portfolios. We say that:

$w^{(i)} \prec_\pi w^{(j)}$ (i is preferred to j lexicographically under π)

if there exists an index $k \in \{1,2,3,4\}$ such that:

- $f_{\pi_k}(w^{(i)}) > f_{\pi_k}(w^{(j)})$
- $f_{\pi_l}\left(w^{(i)}\right) = f_{\pi_l}\left(w^{(j)}\right) \; \forall \; l < k$

This ensures that portfolio $w^{(i)}$ is ranked higher than $w^{(j)}$ based on the earliest objective in the priority list where they differ.

Implication for Sorting Algorithm

In implementation, lexicographic sorting is applied to a list of candidate portfolios via the following process:
1. Compute the vector $f(w^{(i)})$ for each portfolio.
2. Apply a stable sort using π as the priority order.
3. Select the top-ranked portfolio(s) as the lexicographic optimum.

1.2.3. Appendix – Program Architecture Overview for Multidimensional Portfolio Optimization

This appendix outlines the complete architecture of the Python-based optimization engine developed for this thesis. The program is modular and structured to support both *a posteriori* and *a priori* optimization models for constructing 4D efficient frontiers. The system is designed for flexibility, transparency, and extensibility, aligned with academic research and practical asset management requirements.

Modular Architecture Overview

The architecture consists of the following primary modules and their respective responsibilities:

Module	File Name	Core Responsibilities
Main Orchestrator (A Posteriori)	Main_file.py	Executes the full a posteriori pipeline: data loading, portfolio simulation, filtering, and result export.
Main Orchestrator (A Priori)	Optimizer_a_priori.py	Implements weighted preference-based optimization per investor profile using lexicographic sorting or scalarization.
Classical Baseline Model	Main_Classical_MPT.py	Constructs 2D efficient frontier based on mean-variance trade-off using Markowitz theory.
Data Loading	Data_loader.py	Loads asset-level input data, including expected return, volatility, ESG score, and Market SCR, from Excel.
SCR Calculation	scr_calculator.py	Computes diversified Market SCR scores using Solvency II correlation matrices and asset-specific shocks.
Frontier Extraction	frontier.py	Filters Pareto-efficient portfolios using formal dominance logic.
Portfolio Simulation	Portfolio_construction.py	Generates all feasible portfolios using full enumeration with minimum weight constraints (5% granularity).
Visualization	Visualisation.py	Produces comparative charts: 2D/3D scatter plots, radar charts, bar plots, and trade-off diagrams.
Comparison Engine	Comparison_2D_4D.py, Comparison_apriori_aposteriori.py	Conducts systematic visual and numerical comparison of optimization outcomes across models.
Configuration	path_utils.py	Manages file paths to ensure cross-platform compatibility without hardcoded dependencies.
Excel Export	Part of Main_file.py, Main_Classical_MPT.py	Exports simulation results, frontier coordinates, and portfolio weights to .xlsx and .csv files for analysis.

Data Flow
The system executes the following sequential pipeline (for the *a posteriori* model):

1. **Input**: Excel dataset Data_Set_MPT_Multidimensional.xlsx (asset-level inputs).
2. **Processing**: `Data_loader.py` and `scr_calculator.py` preprocess the raw data into 4D vectors.
3. **Portfolio Generation**: `Portfolio_construction.py` generates all feasible portfolio combinations.
4. **Efficiency Filtering**: `frontier.py` selects the non-dominated portfolios using Pareto logic.
5. **Export & Visualization**: Results are saved in Excel and visualized via `Visualisation.py`.

Program Output
The architecture produces the following key deliverables:

- **Efficient frontier datasets** (efficient_frontier_4D.xlsx, Frontier.xlsx)
- **Portfolio weights files** (portfolio_weights_4D.xlsx, etc.)
- **Visual plots**: Interactive 3D scatter plots, radar charts, stacked bars
- **Comparison files**: comparison_apriori_vs_posteriori.xlsx summarizing all metrics

Reproducibility and Modularity

- Each script is standalone-executable and shares consistent interfaces.
- Results can be reproduced by calling main() from the respective orchestrator file.
- All file paths are managed through `path_utils.py` to enable platform-independent execution.

This modular design underpins the empirical and comparative analysis in *Chapter 7* and ensures transparency and flexibility for future research extensions. It also facilitates the parallel execution of alternative modeling techniques and systematic result documentation.

1.2.4. Appendix – Correlation Matrix computed on Excel for comparison purposes

	Eu Large Caps	Eu Small Caps	EM Equities	Developped World Equities	€ Govt Bonds	€ IG Corp Bonds	HY Corp Bonds	EM Bonds	Global Convertible Bonds	Eu Core Real Estate	Global Infrastructure	Commodities	PE	Cash
EU Large Caps	1.0000000	0.3688822	0.2044482	0.2355223	0.0216568	0.0820920	0.1476602	0.1778317	0.1464286	0.3444050	0.2652084	0.1051633	0.4462906	-0.0008377
EU Small Caps		1.0000000	0.2350208	0.2696299	0.0264229	0.1025039	0.1798937	0.2044366	0.1933990	0.4418319	0.3128871	0.1318614	0.5418771	0.0010029
EM Equities			1.0000000	0.1628374	0.0146164	0.0619703	0.1117510	0.1353830	0.1387762	0.2070828	0.1703664	0.1024240	0.2912856	0.0020624
Developped World Equities				1.0000000	0.0190824	0.0668812	0.1108272	0.1181516	0.1479595	0.2432735	0.2075105	0.0706103	0.3604399	0.0011025
€ Govt Bonds					1.0000000	0.0191125	0.0136709	0.0226594	0.0206019	0.0437037	0.0188900	-0.0021003	0.0432534	0.0001307
€ IG Corp Bonds						1.0000000	0.0569855	0.0702810	0.0618471	0.1215710	0.0771252	0.0181102	0.1472920	0.0006560
HY Corp Bonds							1.0000000	0.1105475	0.0950702	0.1787119	0.1405211	0.0540368	0.2428552	0.0021090
EM Bonds								1.0000000	0.1205392	0.2423423	0.1535824	0.0605378	0.2725512	0.0035456
Global Convertible Bonds									1.0000000	0.2091911	0.1279251	0.0553698	0.2677648	0.0028911
Eu Core Real Estate										1.0000000	0.3193805	0.1408074	0.5613028	0.0054006
Global Infrastructure											1.0000000	0.1211094	0.4235252	-0.0008216
Commodities												1.0000000	0.1319826	-0.0016805
PE													1.0000000	0.0026140
Cash														1.0000000

14x14 Correlation Matrix of the asset classes used in this thesis

1.2.5. Appendix – Covariance Matrix computed on Excel for comparison purposes

	Eu Large Caps	Eu Small Caps	EM Equities	Developped World Equities	€ Govt Bonds	€ IG Corp Bonds	HY Corp Bonds	EM Bonds	Global Convertible Bonds	Eu Core Real Estate	Global Infrastructure	Commodities	PE	Cash
EU Large Caps	0.0301532	**0.0307402**	**0.0170374**	**0.0196269**	**0.0018047**	**0.0068410**	**0.0123050**	**0.0148193**	**0.0122024**	**0.0287004**	**0.0221007**	**0.0087636**	**0.0371909**	**-0.0000698**
EU Small Caps		0.0377089	**0.0195851**	**0.0224692**	**0.0022019**	**0.0085420**	**0.0149911**	**0.0170364**	**0.0161166**	**0.0368193**	**0.0260739**	**0.0109884**	**0.0451564**	**0.0000836**
EM Equities			0.0204363	**0.0135698**	**0.0012180**	**0.0051642**	**0.0093126**	**0.0112819**	**0.0115647**	**0.0172569**	**0.0141972**	**0.0085353**	**0.0242738**	**0.0001719**
Developped World Equities				0.0194437	**0.0015902**	**0.0055734**	**0.0092356**	**0.0098460**	**0.0123300**	**0.0202728**	**0.0172925**	**0.0058842**	**0.0300367**	**0.0000919**
€ Govt Bonds					0.0010558	**0.0015927**	**0.0011392**	**0.0018883**	**0.0017168**	**0.0036420**	**0.0015742**	**-0.0001750**	**0.0036044**	**0.0000109**
€ IG Corp Bonds						0.0039248	**0.0047488**	**0.0058568**	**0.0051539**	**0.0101309**	**0.0064271**	**0.0015092**	**0.0122743**	**0.0000547**
HY Corp Bonds							0.0087719	**0.0092123**	**0.0079225**	**0.0148927**	**0.0117101**	**0.0045031**	**0.0202379**	**0.0001757**
EM Bonds								0.0134085	**0.0100449**	**0.0201952**	**0.0127985**	**0.0050448**	**0.0227126**	**0.0002955**
Global Convertible Bonds									0.0146081	**0.0174326**	**0.0106604**	**0.0046142**	**0.0223137**	**0.0002409**
Eu Core Real Estate										0.0540213	**0.0266150**	**0.0117340**	**0.0467752**	**0.0004501**
Global Infrastructure											0.0277606	**0.0100924**	**0.0352938**	**-0.0000685**
Commodities												0.0157145	**0.0109986**	**-0.0001400**
PE													0.0650143	**0.0002178**
Cash														0.0002368

14x14 Covariance Matrix of the asset classes used in this thesis

1.2.6. Appendix – Spreadsheet of ESG Score provided by Refinitiv Eikon for one of the asset class

This appendix documents the procedure followed to extract ESG scores used in the multidimensional optimization model. The ESG data was obtained from Refinitiv Eikon for the iShares MSCI EMU Large Cap UCITS ETF (Acc), covering monthly ESG metrics over a 12-month period (March 2024 – February 2025).

Spreadsheet of ESG Score provided by Refinitiv Eikon for EMUL.MI

EMUL.MI	**iShares MSCI EMU Large Cap UCITS ETF EUR (Acc)**												ESG Score
Scoring Measures	**28-Feb-2025**	**31-Jan-2025**	**31-Dec-2024**	**30-Nov-2024**	**31-Oct-2024**	**30-Sep-2024**	**31-Aug-2024**	**31-Jul-2024**	**30-Jun-2024**	**31-May-2024**	**30-Apr-2024**	**31-Mar-2024**	79.487
ESG Coverage Check	99.745	99.547	99.611	99.418	99.339	99.525	98.794	98.363	98.524	98.854	97.597	99.159	
ESG Coverage Count	100.000	101.000	101.000	101.000	101.000	101.000	99.000	101.000	101.000	185.000	101.000	101.000	
LSEG Lipper Fund ESG Reporting Scores	**79.494**	**79.487**	**79.408**	**77.869**	**79.590**	**79.701**	**79.986**	**80.177**	**79.794**	**79.574**	**79.361**	**79.085**	
Environment Pillar Score	81.533	81.358	81.238	78.832	81.358	81.420	80.405	80.447	80.205	80.420	80.572	80.221	
Resource Use Score	89.643	89.732	89.608	88.996	89.907	89.667	89.860	89.978	89.751	89.294	89.407	89.102	
Emissions Score	87.948	88.168	87.970	88.016	88.140	88.115	88.267	88.221	88.029	88.046	88.183	87.809	
Environmental Innovation Score	66.617	65.876	65.836	55.011	65.824	65.603	65.168	65.222	65.031	65.890	65.953	65.597	
Social Pillar Score	83.036	83.227	83.157	84.540	83.128	83.817	84.067	84.208	83.945	83.528	83.445	83.172	
Workforce Score	91.181	91.388	91.224	90.242	91.586	92.164	92.344	92.134	92.815	92.676	92.511	92.407	
Human Rights Score	81.508	81.386	81.276	82.170	81.813	81.921	82.176	82.519	81.265	81.405	81.323	81.507	
Community Score	82.580	82.855	82.748	85.190	82.393	84.081	84.371	84.402	84.129	83.545	83.460	81.991	
Product Responsibility Score	77.151	77.519	77.631	70.013	77.912	77.293	77.660	77.844	77.673	76.431	76.381	76.278	
Governance Pillar Score	73.484	73.388	73.312	73.938	73.376	73.159	73.791	74.185	73.482	73.171	72.377	72.214	
Management Score	76.687	76.826	76.669	75.215	76.783	77.502	78.493	78.888	78.465	77.915	76.429	76.141	
Shareholders Score	59.050	58.420	58.376	63.225	57.765	54.230	54.023	54.445	53.157	53.793	53.932	54.181	
CSR Strategy Score	79.122	78.648	78.929	83.624	79.300	79.838	79.932	80.284	79.052	78.519	79.784	79.628	
ESG Controversies Score	**44.963**	**45.051**	**45.261**	**43.270**	**48.423**	**46.307**	**47.066**	**47.689**	**48.145**	**48.017**	**47.817**	**46.608**	
ESG Combined Score	**60.062**	**60.023**	**60.102**	**58.461**	**61.426**	**60.647**	**61.173**	**61.672**	**61.800**	**61.589**	**61.357**	**60.883**	

Data Source Overview
The dataset includes a comprehensive breakdown of ESG indicators, grouped into:

- **Pillar Scores** (Environmental, Social, Governance),
- **Sub-Pillar Metrics** (e.g., Emissions, Innovation, Workforce),
- **Composite Ratings** (e.g., ESG Combined Score),
- **Fund-Level Metrics**, notably the LSEG Lipper Fund ESG Reporting Score.

This latter metric, the LSEG Lipper Fund ESG Reporting Score, was chosen as the representative ESG score for this asset class.

Justification for Using the LSEG Lipper Fund ESG Reporting Score
The LSEG Lipper Fund ESG Reporting Score was selected for the following reasons:

1. **Fund-Specific Relevance**: Unlike firm-level ESG metrics, the Lipper ESG score reflects the overall sustainability profile of the ETF vehicle itself, which is directly relevant for asset class representation in portfolio allocation.
2. **Reporting Consistency**: The Lipper score captures the extent and quality of ESG-related disclosures and reporting at the fund level, aligning with regulatory transparency and investor-alignment goals.
3. **Standardization**: LSEG's methodology provides a unified scoring scale across funds, allowing for valid comparisons across asset classes and regions.
4. **Recency**: The most recent available value, as of 28 February 2025, was used. This ensures the score reflects the latest ESG assessments based on complete and current data.

Implementation
From the row labeled "LSEG Lipper Fund ESG Reporting Score", the value corresponding to 28-Feb-2025 was extracted and directly used as the ESG input: {79.494}
The same method was applied to other asset classes using their respective LSEG scores when available from Refinitiv Eikon.

Methodological Note
Although other ESG metrics (e.g., ESG Combined Score, Governance Pillar Score) were available, they were excluded due to:

- their firm-level orientation rather than fund-level,
- potential overlap with already captured performance risk metrics (e.g., controversies, emissions),
- and lack of alignment with the fund-aggregated reporting perspective necessary for asset class comparison.

This approach ensures internal consistency and relevance within the multidimensional optimization framework developed in this thesis.

1.2.7. Appendix – Spreadsheet of ESG Score provided by Sustainalytics by Morningstar for one of the asset class

This appendix outlines the approach used to derive ESG scores from Morningstar for equity-based asset classes in the portfolio optimization model. The ESG data pertains to the iShares MSCI EMU Large Cap ETF EUR Acc (EMUL), retrieved from the fund's Sustainability Rating section on Morningstar as of February 28, 2025.

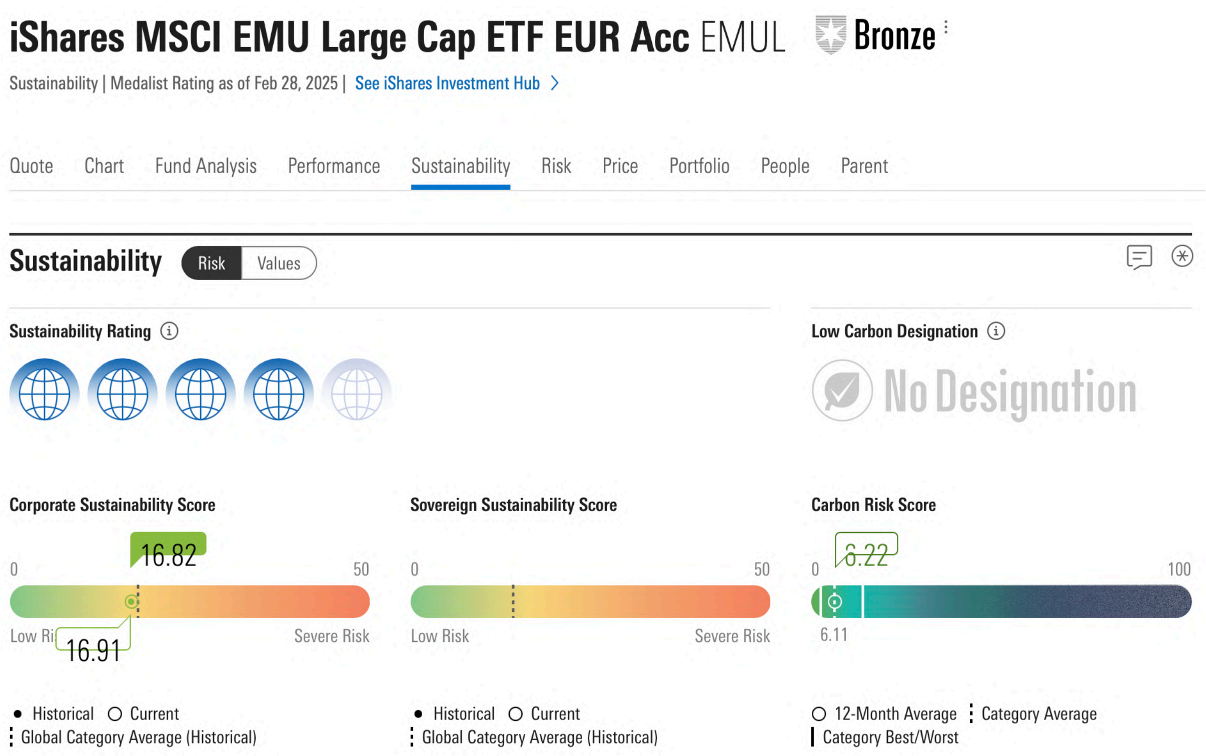

Spreadsheet of ESG Score provided by Sustainalytics by Morningstar for EMUL.MI

Source Overview and Available Metrics

Morningstar, in partnership with Sustainalytics, provides multiple ESG-related metrics at the fund level, including:

- **Corporate Sustainability Score** – evaluates the ESG risk profile of the fund's corporate holdings,
- **Sovereign Sustainability Score** – assesses ESG risk in sovereign bond exposure,
- **Carbon Risk Score** – captures climate-related financial risk exposure,
- **Sustainability Rating** – assigns globe ratings based on ESG performance relative to peers.

For this analysis, the Corporate Sustainability Score was selected as the ESG input.

Rationale for Score Selection

The Corporate Sustainability Score was chosen for the following reasons:

1. **Asset Class Relevance**: The selected fund (EMUL) is comprised primarily of corporate equity holdings. Thus, corporate-level ESG assessment provides the most appropriate measure of ESG performance.
2. **Transparency and Simplicity**: Morningstar explicitly reports this metric on a monthly basis, enabling straightforward use in portfolio-level ESG integration.
3. **Quantitative Availability**: The precise value of 16.82 as of 28-Feb-2025 was available and reported clearly, with historical traceability.
4. **No Transformation Applied**: In line with the modeling approach taken in this thesis, the score was used as reported (i.e., no normalization, no inversion). This ensures transparency and avoids introducing subjective preprocessing assumptions.

Implementation

The score was incorporated directly into the asset's 4D input vector with ESG = 16.82

This methodology was applied uniformly for all equity-based asset classes for which Morningstar Corporate Sustainability Scores were available.

Methodological Clarification

Although the Morningstar ESG scoring system typically interprets lower values as better (i.e., lower ESG risk), no sign inversion was performed in this model. The optimization framework thus interprets this raw score at face value and does not convert it into a maximization-compatible form. This choice was made to preserve data integrity and avoid subjective manipulation of ESG input data.

If required, a footnote or remark can be added in Chapter 4 or 5 to clarify that lower ESG scores indicate lower ESG risk, especially when interpreting optimization results that incorporate Morningstar-based scores.

1.2.8. Appendix – Regression Analysis for ESG Scores

This appendix provides an overview of the regression methodology used in this thesis to get a list of ESG scores that stay consistent over the entire investable instruments. This is necessary because, as mentioned in Section 2.2.1, discrepencies exist and no clear methodology has been defined. For this regression, a proxy asset class was used for HY Corp Bonds. This asset class has an ESG score on both provider's platforms and allows to bring relevance to the regression.

Refinitiv Eikon ESG Score for 10 Asset Classes

Rank	Asset Class	ETF ticker	Score (highest = best)
4	Emerging Market Equity	iShares Core MSCI EM IMI (EIMI)	58.647
2	EU Large Cap	iShares MSCI EMU Large Cap (EMUL)	60.062
3	EU Small Cap	SPDR MSCI Europe Small Cap (SMCX)	59.606
6	Euro Corp Bonds	iShares € Corp Bond Lg Cap (IBCX)	55.628
1	European Core Real Estate	iShares European Property Yield (IPRE	70.753
8	Global Convertible Bonds	SPDR FTSE Global Convertibles (SPF2)	49.901
7	Global Developed Equity	iShares MSCI World (IWDA)	50.443
5	Global Infrastructure	iShares Global Infrastructure ETF (IGF	58.307
11	High Yield Corporate Bonds	iShares Broad USD High Yield Corporate	44.512
9	Private Equity	Invesco Global Listed PE (PSP)	40.422

Sustainalytics by Morningstar ESG Score for 10 Asset Classes

Rank	Asset Class	ETF	Score (lowest = best)
11	Emerging Market Equity	iShares Core MSCI EM IMI (EIMI)	23.19
4	EU Large Cap	iShares MSCI EMU Large Cap (EMUL)	16.82
6	EU Small Cap	SPDR MSCI Europe Small Cap (SMCX)	19.14
5	Euro Corp Bonds	iShares € Corp Bond Lg Cap (IBCX)	19.12
1	European Core Real Estate	iShares European Property Yield (IPRE)	9.81
9	Global Convertible Bonds	SPDR FTSE Global Convertibles (SPF2)	21.84
8	Global Developed Equity	iShares MSCI World (IWDA)	20.41
7	Global Infrastructure	iShares Global Infrastructure ETF (IGF)	20.06
12	High Yield Corporate Bonds	iShares Broad USD High Yield Corporate Bond ET	25.42
10	Private Equity	Invesco Global Listed PE (PSP)	22.12

Regression Analysis

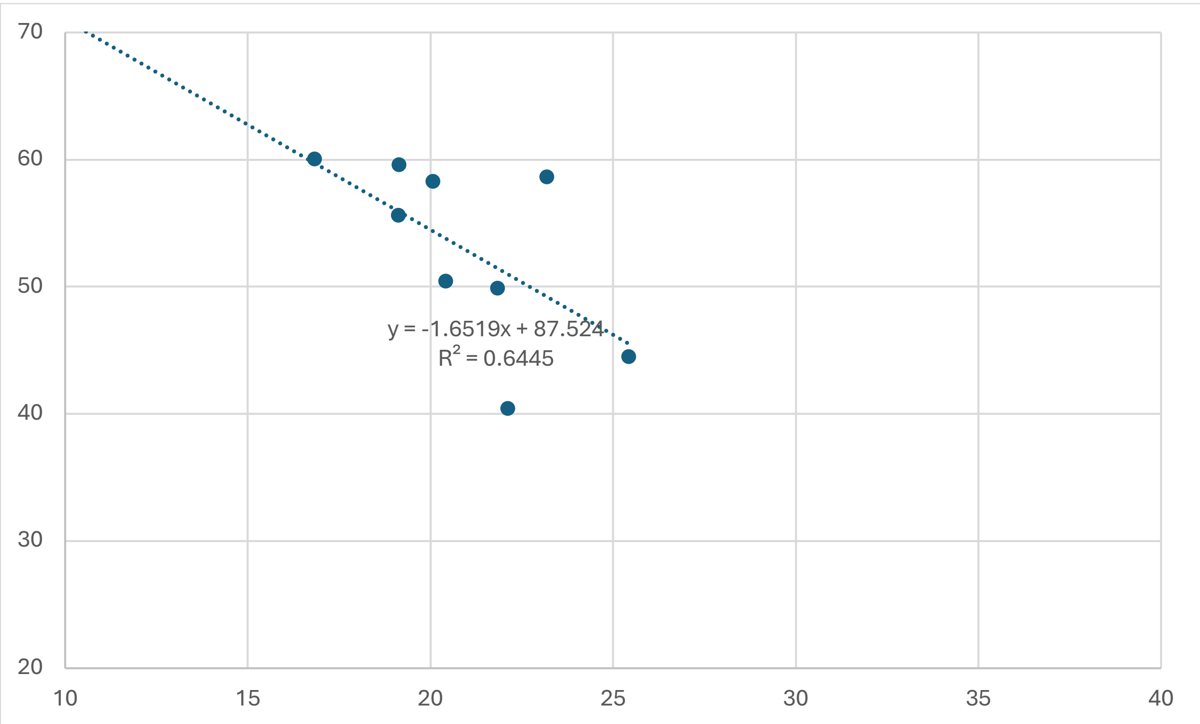

Regression Formula:
Eikon ESG Score = -1.6519.(Morningstar ESG Score) + 87.524

This means
For each 1-point increase in the Morningstar ESG risk score (i.e., worse ESG), the Eikon ESG score drops by ~1.65. This confirms the inverse relationship between the two scoring systems was properly captured

1.2.9. Appendix – 2D projections plots with the 3rd dimension represented in colors

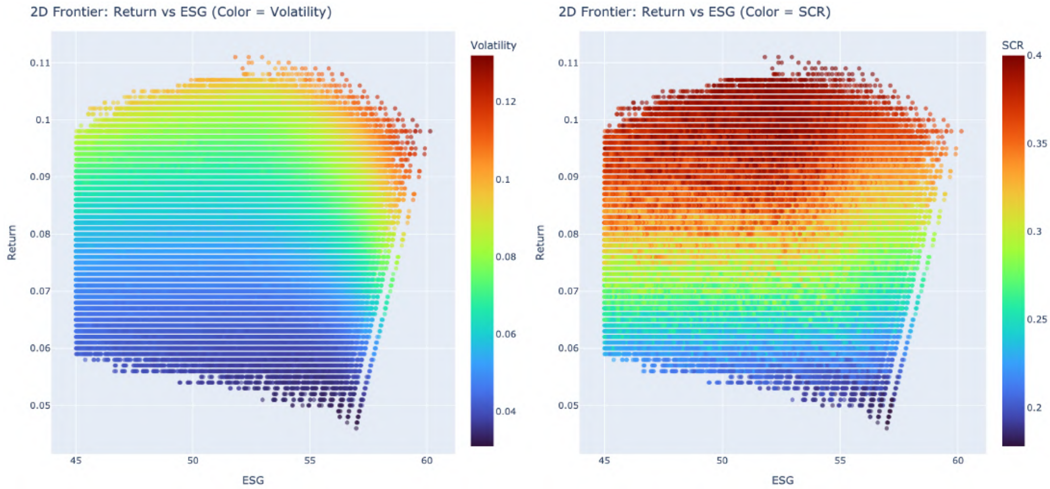

These projections plot ESG against Market SCR while applying color gradients to represent volatility (left) and return (right). The triangular shape reveals that portfolios with ESG scores above 58 typically entail higher volatility and capital requirements, confirming the multidimensional nature of trade-offs. These patterns reinforce the thesis that ESG and SCR integrate as meaningful, non-redundant objectives in portfolio optimization

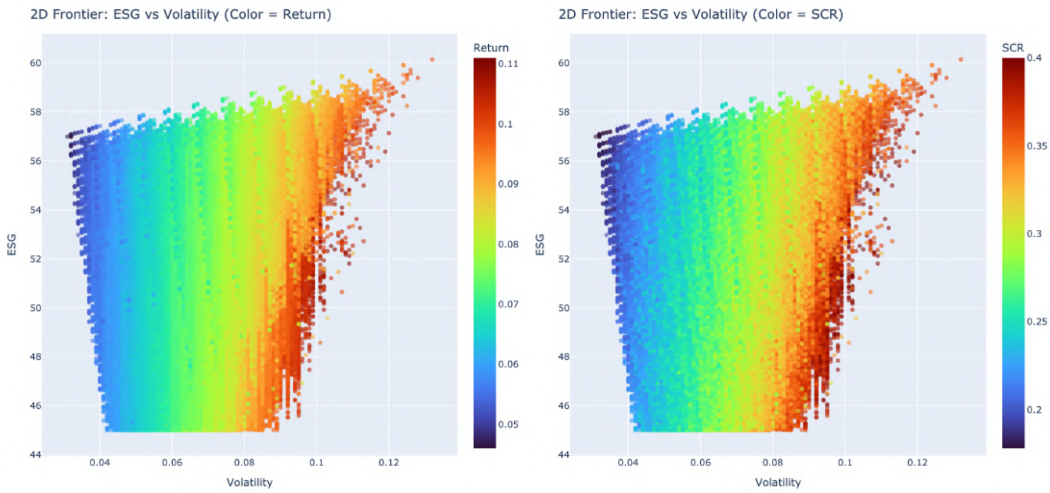

These plots display ESG scores against volatility, colored by Market SCR (left) and return (right). The visuals show that ESG-optimal portfolios span a wide volatility range, but pushing ESG above 58 frequently leads to higher SCR and return values. This confirms that sustainability enhancements carry quantifiable trade-offs in both capital requirement and return potential.

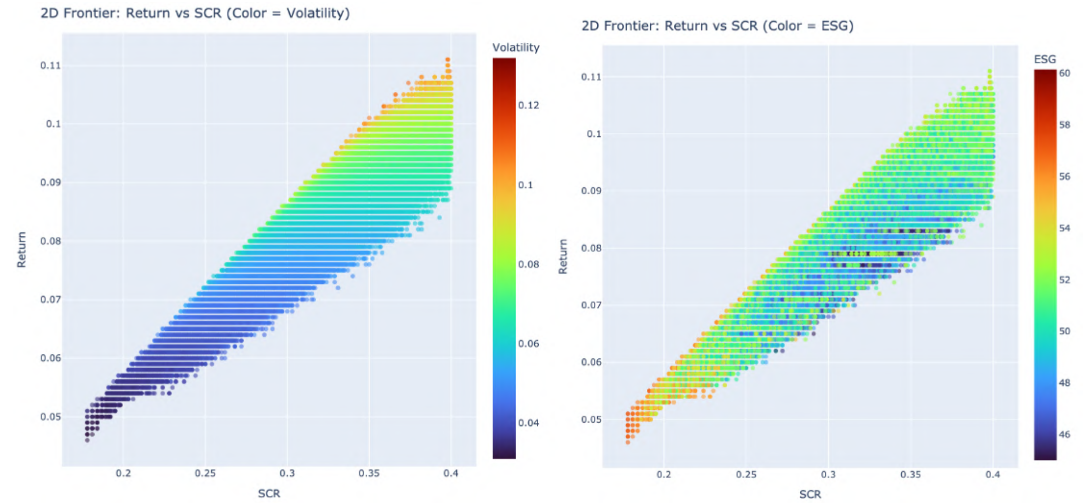

These charts map return against Market SCR, with color gradients indicating ESG scores (left) and volatility (right). The nearly linear trade-off between return and capital requirement is evident, while ESG-colored gradients reveal no simple correlation, supporting the view that ESG adds a distinct dimension. Volatility increases monotonically with return and SCR, confirming classical risk-return dynamics.

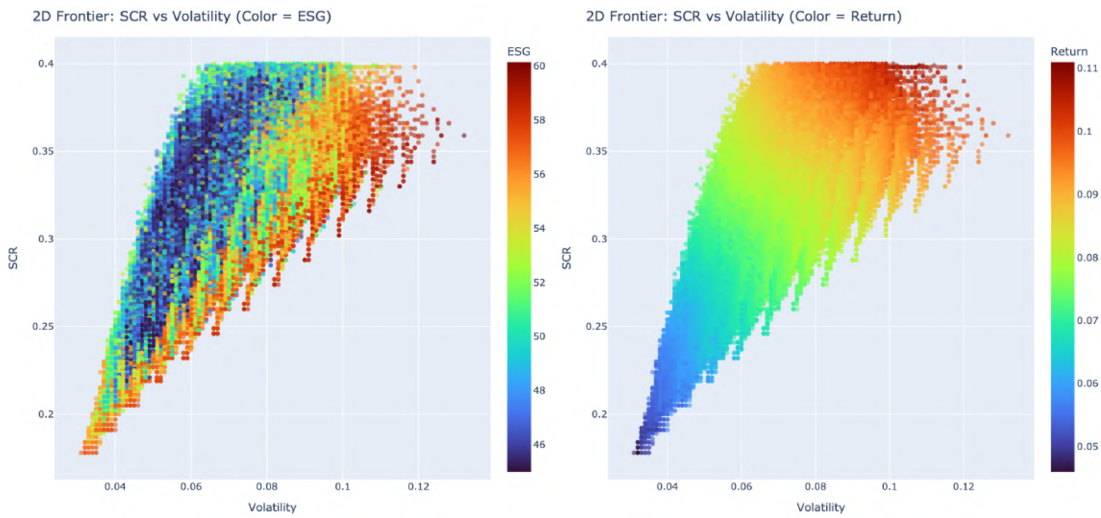

These scatter plots map Market SCR against volatility, with ESG scores (left) and return levels (right) indicated by color. While the return-colored plot confirms a clear co-movement between volatility and SCR, the ESG distribution is more scattered, highlighting the independence of ESG from classical financial risk dimensions. This further substantiates the non-redundant role of ESG in the multidimensional optimization framework.

1.2.10. Appendix – Comparison Between Pure and Diversified Investor Profiles

This appendix compares the portfolio compositions generated by the *a priori* model under two different sets of investor preferences:

1. **Diversified profiles** (Table 7-6) – representing more realistic multi-objective allocations based on weighted scoring systems aligned with five distinct investor archetypes: scr_conservative, low_risk, esg_focused, balanced, and return_max.
2. **Pure profiles** – each optimized exclusively for a single objective (Return, Volatility, ESG, or SCR).

Profile	pure_scr	pure_return	pure_esg	pure_volatility
Return	**0.051**	**0.111**	**0.098**	**0.047**
Volatility	**0.035**	**0.109**	**0.132**	**0.031**
ESG	**57.20**	**51.81**	**60.14**	**57.00**
SCR	**0.178**	**0.398**	**0.359**	**0.178**
score	**3.70**	**3.49**	**3.48**	**2.99**
€ Government Bonds	**0.519**	0.037	0.037	**0.074**
EU Core Real-Estate	0.037	0.037	**0.519**	0.037
EU Small Caps	0.037	**0.333**	0.037	0.037
Cash	0.037	0.037	0.037	0.037
EU Large Caps	0.037	0.037	0.037	0.037
EM Equities	0.037	0.037	0.037	0.037
Developed-World Equities	0.037	0.037	0.037	0.037
High-Yield Corporate Bonds	0.037	0.037	0.037	0.037
€ IG Corporate Bonds	0.037	0.037	0.037	0.037
EM Bonds	0.037	0.037	0.037	0.037
Global Convertible Bonds	0.037	0.037	0.037	0.037
Global Infrastructure	0.037	0.037	0.037	0.037
Commodities	0.037	0.037	0.037	0.037
Private Equity	0.037	0.037	0.037	0.037
Distance to efficient frontier	**0**	**0**	**0**	**0**

The comparison illustrates how objective trade-offs influence portfolio structure and highlights the capacity of the *a priori* optimization model to adapt to diverse investor priorities.

Diversified Profiles (Weighted Multi-Objective Optimization)

Profile	Notable Allocations	Characteristics
scr_conservative	€ Government Bonds (29.6%), Cash (25.9%)	Designed to minimize capital requirements; high liquidity and low risk.
low_risk	Same as above	Aligned with minimal volatility. Slight increase in ESG from pure low-risk profile.
esg_focused	EU Core Real Estate (7.4%), € Government Bonds (48.1%)	Prioritizes ESG while maintaining capital efficiency.
balanced	Cash (25.9%), € Government Bonds (29.6%)	Reflects a stable middle-ground allocation with equal weighting on all objectives.
return_max	EU Small Caps (18.5%), Private Equity (22.2%), Global Infrastructure (11.1%)	Tilts aggressively toward high-return, higher-risk assets. SCR and volatility are both elevated.

All diversified profiles remain on the efficient frontier (distance = 0), validating their optimality within preference-weighted scoring.

Pure Profiles (Single-Objective Optimization)

Profile	Dominant Allocation	Key Observation
pure_scr	€ Government Bonds (51.9%)	Chosen for minimal market SCR exposure.
pure_return	EU Small Caps (33.3%)	Highest-return asset within the universe.
pure_esg	EU Core Real Estate (51.9%)	Offers strongest ESG score.
pure_volatility	€ Government Bonds (7.4%)	Again preferred for low volatility.

All other assets receive minimal weights (3.7%) due to the granularity constraint. These theoretical profiles serve as reference points for extreme investor preferences.

Key Takeaways

- **Diversification emerges** naturally under multi-objective weighting, contrasting with the concentrated nature of pure profiles.
- **Overlap in dominant assets** confirms objective correlations: e.g., government bonds dominate both SCR and low-volatility strategies.
- **Trade-off transparency**: The return_max profile sacrifices both ESG and SCR to reach a return of 11%, while esg_focused prioritizes sustainability at the cost of a lower return (5.5%).

This appendix reinforces the robustness of the *a priori* model in constructing investor-specific optimal portfolios and showcases its flexibility in supporting tailored asset allocation under multidimensional objectives.

1.2.11. Appendix – Normalized Radar Charts of A Priori Profiles

This appendix presents normalized radar charts for the five *a priori* investor profiles introduced in *Chapter 5*. Each chart visualizes performance across four portfolio dimensions, Return, Volatility, ESG, and SCR, after applying Z-score normalization.

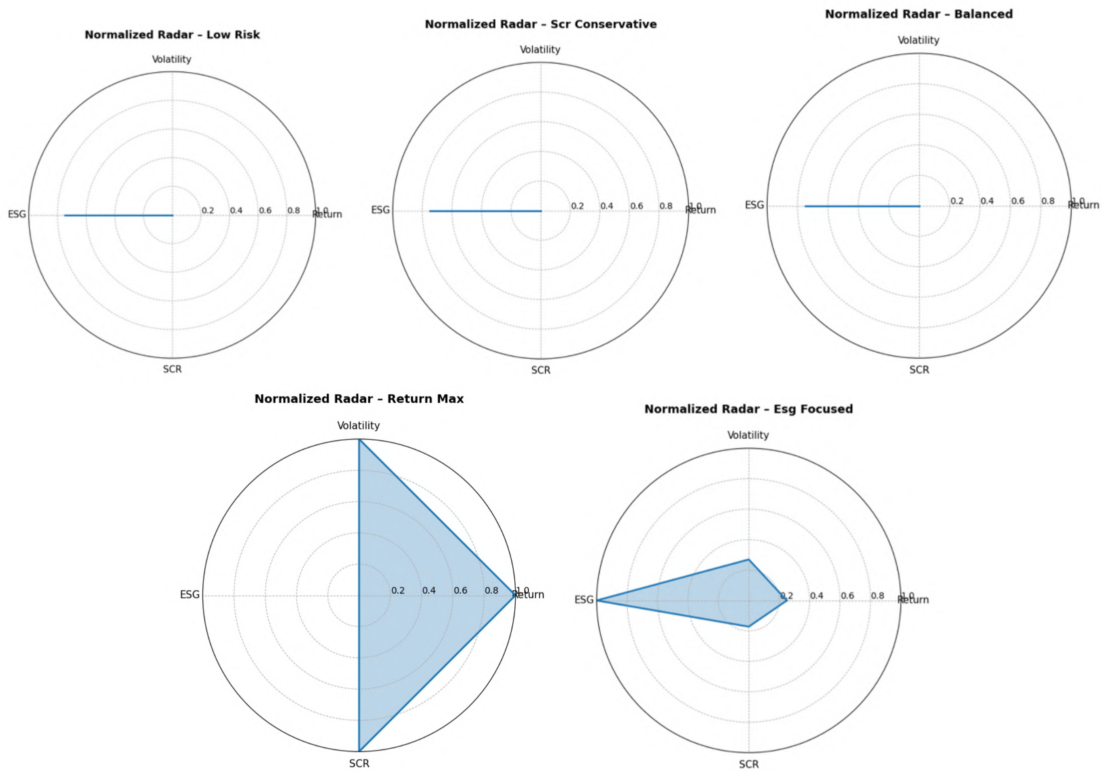

Objective of the Visualization

These radar plots illustrate the relative positioning of each profile across key metrics, allowing for clear visual comparison of trade-offs. They help highlight which objectives are being prioritized or deprioritized in each optimization and expose the multidimensional tension between performance, risk, sustainability, and capital requirements.

Chart Interpretation

- **Return Max** spans the outermost edges on return, SCR, and volatility, showing that this profile scores well above the average on these dimensions, while scoring far below average on ESG.
- **Low Risk, SCR Conservative, and Balanced** show nearly collapsed shapes, indicating that they perform close to or below the average in most dimensions, especially on return and SCR.
- **ESG Focused** achieves a prominent positive Z-score on the ESG axis while remaining near-neutral or slightly negative on others, suggesting a cleaner prioritization trade-off.

Methodology Note: Z-Score Normalization
Z-score normalization centers each metric around its mean and rescales by its standard deviation *(cf. Section 5.2.3)*

This method emphasizes how far a profile deviates from the average, allowing us to detect outperformance or underperformance across objectives in standardized units.

References

Alexander, C. (2008). Market risk analysis. 2 : quantitative methods in finance : Practical financial econometrics (pp. 90–95). Wiley.

Amihud, Y., & Mendelson, H. (1986). Liquidity and Stock Returns. Financial Analysts Journal, 42(3), 43–48. https://doi.org/10.2469/faj.v42.n3.43

Amihud, Y., & Noh, J. (2002). Illiquidity and Stock Returns: Cross-Section and Time-Series Effects. SSRN Electronic Journal, 5(1). https://doi.org/10.2139/ssrn.3139180

Bengio, E., Jain, M., Korablyov, M., Precup, D., & Bengio, Y. (2021). Flow Network based Generative Models for Non-Iterative Diverse Candidate Generation. ArXiv.org. https://arxiv.org/abs/2106.04399

Bennani, L., Le Guenedal, T., Lepetit, F., Ly, L., Mortier, V., Roncalli, T., & Sekine, T. (2018, November 27). How ESG Investing Has Impacted the Asset Pricing in the Equity Market. Papers.ssrn.com. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3316862

Berg, F., Kölbel, J., & Rigobon, R. (2019). Aggregate Confusion: The Divergence of ESG Ratings. SSRN Electronic Journal, 1(1). https://doi.org/10.2139/ssrn.3438533

Bianchi, J., & Bigio, S. (2022). Banks, Liquidity Management, and Monetary Policy. Econometrica, 90(1), 391–454. https://doi.org/10.3982/ecta16599

Black, F., & Litterman, R. (1992). Global Portfolio Optimization. Financial Analysts Journal, 48(5), 28–43. https://doi.org/10.2469/faj.v48.n5.28

Bodnar, T., Okhrin, Y., Vitlinskyy, V., & Zabolotskyy, T. (2018). Determination and estimation of risk aversion coefficients. Computational Management Science, 15(2), 297–317. https://doi.org/10.1007/s10287-018-0317-x

Boniver, F. (2025). Master Thesis Meeting [Letter to Samy El Azri]. Feedback proposed.

Briere, M., & Ramelli, S. (2021a, September). Green Sentiment, Stock Returns, and Corporate Behavior. Papers.ssrn.com. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3850923

Briere, M., & Ramelli, S. (2021b, October). Responsible Investing and Stock Allocation. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.3853256

Campbell, J. Y., Chan, Y., & Viceira, L. (2003). A Multivariate Model of Strategic Asset Allocation. Journal of Financial Economics, 67(1), 41–80.

Chunhachinda, P., Dandapani, K., Hamid, S., & Prakash, A. J. (1997). Portfolio Selection and Skewness: Evidence from International Stock Markets. Journal of Banking & Finance, 21(2), 143–167. https://doi.org/10.1016/s0378-4266(96)00032-5

De Jong, M. (2016). Portfolio Optimisation in an Uncertain World. In Amundi Asset Management (Vol. 19, Issue 4, pp. 216–221). https://doi.org/10.1057/s41260-017-0066-3

EIOPA. (2024). *DIRECTIVE 138/2009/EC (SOLVENCY II DIRECTIVE) | European Insurance and Occupational Pensions Authority*. European Insurance and Occupational Pensions Authority. https://www.eiopa.europa.eu/rulebook/solvency-ii-single-rulebook/directive-1382009ec-solvency-ii-directive_en

Elton, E. J., & Gruber, M. J. (1997). Modern portfolio theory, 1950 to date. Journal of Banking & Finance, 21(11-12), 1743–1759. https://doi.org/10.1016/s0378-4266(97)00048-4

Fahy, D. (2024, September 6). Using Efficient Frontier Analysis to Find the Optimum Portfolio - Money To The Masses. Money to the Masses. https://moneytothemasses.com/8020-articles/using-efficient-frontier-analysis-to-find-the-optimum-portfolio

Fama, E. F., & French, K. R. (2015). A Five-Factor Asset Pricing Model. Journal of Financial Economics, 116(1), 1–22. https://doi.org/10.1016/j.jfineco.2014.10.010

Francis, J. C., & Kim, D. (2013). Modern Portfolio Theory: Foundation Analysis and New Developments (Wiley Finance, Ed.). John Wiley & Sons, Inc.

François, P., & Hübner, G. (2024). The Complete Guide to Portfolio Performance. John Wiley & Sons.

Friede, G., Busch, T., & Bassen, A. (2015). ESG and Financial Performance: Aggregated Evidence from more than 2000 Empirical Studies. Journal of Sustainable Finance & Investment, 5(4), 210–233. https://doi.org/10.1080/20430795.2015.1118917

García, F., González-Bueno, J., Oliver, J., & Riley, N. (2019). Selecting Socially Responsible Portfolios: A Fuzzy Multicriteria Approach. Sustainability, 11(9), 24–96. https://doi.org/10.3390/su11092496

Gasser, S. M., Rammerstorfer, M., & Weinmayer, K. (2017). Markowitz Revisited: Social Portfolio Engineering. European Journal of Operational Research, 258(3), 1181–1190. https://doi.org/10.1016/j.ejor.2016.10.043

Hübner, G., Lambert, M., & Papageorgiou, N. (2014). Higher-moment Risk Exposures in Hedge Funds. European Financial Management, 21(2), 236–264. https://doi.org/10.1111/eufm.12054

Hull, J. C. (2021). Options, Futures, and Other Derivatives (11th ed.). Pearson Educational Limited. Copyright.

IBM Corporation. (2022, December 9). IBM ILOG CPLEX Optimization Studio. Ibm.com. https://www.ibm.com/docs/en/icos/22.1.1?topic=parameters-epsilon-degree-tolerance-used-in-linearization&utm

IFT. (2025). Portfolio Risk and Return. IFT.World. https://ift.world/booklets/portfolio-management-portfolio-risk-and-return-part-i-part4/

Jacobs, B. I., & Levy, K. N. (2012, November 26). Introducing Leverage Aversion into Portfolio Theory and Practice. Ssrn.com. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2181457

Jacobs, B. I., & Levy, K. N. (2013). Leverage Aversion, Efficient Frontiers,and the Efficient Region. The Journal of Portfolio Management, 39(3), 54–64. https://doi.org/10.3905/jpm.2013.39.3.054

La Torre, M. (2025, March). Sustainable Finance and Impact Banking.

Labella, M., Sullivan, L., Russell, J., & Novikov, D. (2019). The Devil is in the Details: The Divergence in ESG Data and Implications for Sustainable Investing. QS Investors.

Lahmiri, S., & Bekiros, S. (2019). Cryptocurrency forecasting with deep learning chaotic neural networks. Chaos, Solitons & Fractals, 118(1), 35–40. https://doi.org/10.1016/j.chaos.2018.11.014

Lassance, N., & Vrins, F. (2021). Portfolio selection with parsimonious higher comoments estimation. Journal of Banking & Finance, 126(126), 106–115. https://doi.org/10.1016/j.jbankfin.2021.106115

Lecomte, F. (2024, March). Asset Management. Investment and Portfolio Management Seminars. Mr Lecomte came to give a guest lecture in the event of a seminar organised for the master's in Banking and Asset Management course Investment and Portfolio Management.

Liu, Y. (2023). Modern Portfolio Theory: analysis and implication. Advances in Economics Management and Political Sciences, 4(1), 158–164. https://doi.org/10.54254/2754-1169/4/20221048

Lo, A. W. (2004, October 15). The Adaptive Markets Hypothesis: Market Efficiency from an Evolutionary Perspective. Papers.ssrn.com. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=602222

LSEG Data & Analytics. (2023). LSEG® Lipper Fund ESG scores. https://www.lseg.com/content/dam/data-analytics/en_us/documents/fact-sheets/lipper-fund-esg-scores.pdf

Mangram, M. E. (2013). A Simplified Perspective of the Markowitz Portfolio Theory. Ssrn.com. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2147880

Markowitz, H. (1952). Portfolio Selection. The Journal of Finance, 7(1), 77–91. http://links.jstor.org/sici?sici=0022-1082%28195203%297%3A1%3C77%3APS%3E2.0.CO%3B2-1

Merton, R. C. (1972). An Analytic Derivation of the Efficient Portfolio Frontier. The Journal of Financial and Quantitative Analysis, 7(4), 18–51. https://doi.org/10.2307/2329621

Metaxiotis, K. (2019). A Mean-Variance-Skewness Portfolio Optimization Model. International Journal of Computer and Information Engineering, 13(2), 85–88.

Meucci, A. (2005). Risk and Asset Allocation (pp. 239–243). Springer Finance.

Neves, J. F., Nunes Da Silva, P., & Vasconcellos, C. F. (2017). Maximization of Utility and Portfolio Selection Models. Cadernos Do IME - Série Matemática, 0(11), 18–23. https://doi.org/10.12957/cadmat.2017.29731

Nevil, S. (2019). What a Z-Score Tells Us. Investopedia. https://www.investopedia.com/terms/z/zscore.asp

Nofsinger, J., & Varma, A. (2014). Socially Responsible Funds and Market Crises. Journal of Banking & Finance, 48(1), 180–193. https://doi.org/10.1016/j.jbankfin.2013.12.016

Official Journal of The European Union. (2025, January 17). DIRECTIVE 2009/138/EC OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL of 25 November 2009 on the taking-up and pursuit of the business of Insurance and Reinsurance (Solvency II).

Pedersen, L. H., Fitzgibbons, S., & Pomorski, L. (2021). Responsible Investing: The ESG-Efficient Frontier. Journal of Financial Economics, 142(1), 572–597. https://doi.org/10.1016/j.jfineco.2020.11.001

Rubinstein, M. (2002). Markowitz's "Portfolio Selection": A Fifty-Year Retrospective. The Journal of Finance, 57(3), 1041–1045. https://doi.org/10.1111/1540-6261.00453

Sharpe, W. F. (1964). Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk. The Journal of Finance, 19(3), 425–442. https://onlinelibrary.wiley.com/doi/10.1111/j.1540-6261.1964.tb02865.x

Sharpe, W. F. (1966). Mutual Fund Performance. The Journal of Business, 39(1). https://doi.org/10.1086/294846

Silva Jr., A. F. A., Lôpo, R., & Lofiego, P. (2021). ESG Integration Strategy for Stocks Portfolios Based on a Resampling Methodology with a Multivariate Normal Distribution. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.3935364

Trinks, A., Scholtens, B., Mulder, M., & Dam, L. (2018). Fossil Fuel Divestment and Portfolio Performance. Ecological Economics, 146(1), 740–748. https://doi.org/10.1016/j.ecolecon.2017.11.036

Vaneerdewegh, P. (2025, March 4). How to Assign ESG Score to Different Asset Classes (S. El Azri, Interviewer) [Personal communication].

Von Neumann, J., & Morgenstern, O. (2017). Theory Of Games And Economic Behavior. In Internet Archive (pp. 15–29). Princeton University press. https://archive.org/details/in.ernet.dli.2015.215284/page/n13/mode/2up (Original work published 1953)

Von Wallis, M., & Klein, C. (2014). Ethical Requirement and Financial Interest: A Literature Review on Socially Responsible Investing. Business Research, 8(1), 61–98. https://doi.org/10.1007/s40685-014-0015-7

Winkel, D., Strauß, N., Schubert, M., & Seidl, T. (2023). Simplex Decomposition for Portfolio Allocation Constraints in Reinforcement Learning. Frontiers in Artificial Intelligence and Applications, 1(1). https://doi.org/10.3233/faia230573

Yu, J., & Zhang, J. (2023). A Comprehensive Analysis of The Modern Portfolio Theory. BCP Business & Management, 38(1), 2111–2114. https://doi.org/10.54691/bcpbm.v38i.4046

Zhang, H. (2024). Limitations and Critique of Modern Portfolio Theory: A Comprehensive Literature Review. Advances in Economics, Management and Political Sciences, 60(1), 24–29. https://doi.org/10.54254/2754-1169/60/20231148

Zhang, Z., Li, Z., & Yu, W. (2025). Lexicographic optimization-based approaches to learning a representative model for multi-criteria sorting with non-monotonic criteria. Computers & Operations Research, 175(1). https://doi.org/10.1016/j.cor.2024.106917

Statement on the Use of Artificial Intelligence

Artificial intelligence tools, including OpenAI's ChatGPT, were used during the development of this thesis to support the writing process, code debugging, and the refinement of analytical interpretations. These tools were employed under the supervision and discretion of the author to:

- Assist in editing and improving the academic structure and clarity of written sections;
- Troubleshoot programming issues and validate Python scripts used in portfolio simulation and optimization;
- Format and enhance visual communication (e.g., graph labeling, axis naming, captioning of figures);

At all times, AI-assisted content was critically evaluated, revised, and integrated by the author to ensure accuracy, originality, and compliance with academic standards. The intellectual design, empirical modeling, data interpretation, and final formulation of arguments remain solely the author's responsibility.

The use of AI did not replace academic research, critical thinking, or original analysis, and all content generated with AI tools was used transparently and ethically in accordance with the university's guidelines on responsible use of technology.

EXECUTIVE SUMMARY

This thesis, titled *Multidimensional Extension of the Modern Portfolio Theory*, investigates the feasibility and value of expanding classical portfolio optimization beyond its traditional focus on return and risk. In response to the growing influence of sustainability concerns and regulatory capital requirements in institutional investing, the study integrates two additional dimensions: Environmental, Social, and Governance (ESG) scores and Market Solvency Capital Requirement (SCR) into the construction of the efficient frontier.

Using a dataset of 14 asset classes, the thesis constructs a four-dimensional optimization framework. Each asset class is characterized by expected return, volatility, ESG score, and SCR (calculated using Solvency II regulatory stress factors). Two complementary optimization models are developed: an *a posteriori* model that filters non-dominated portfolios from a fully enumerated universe, and an *a priori* model that simulates investor-specific preferences through scoring and lexicographic sorting. Portfolios are generated using tranche-based weight combinations with 3.75% granularity, ensuring transparency and completeness.

The results show that the inclusion of ESG and SCR does not significantly reshapes the efficient frontier. High-ESG portfolios tend to cluster in regions with moderate return and higher volatility, while SCR-efficient portfolios occupy the lower-return, lower-volatility quadrant. Compared to classical 2D portfolios, the 4D optimized allocations differ significantly, favoring assets like Euro Government Bonds and Core Infrastructure when sustainability or capital efficiency is prioritized. These findings confirm that ESG and SCR are not redundant with financial risk-return metrics and can be meaningfully integrated into optimization without distorting outcomes.

Moreover, the thesis demonstrates that such multidimensional optimization is computationally feasible and practically interpretable. Investor-specific portfolios generated under the *a priori* model (e.g., ESG-focused, return-maximizing, or SCR-conservative) illustrate the range of trade-offs that institutions may face. Importantly, the multidimensional model preserves financial performance while enhancing alignment with real-world regulatory and sustainability goals.

This research makes a novel theoretical contribution by treating ESG and SCR as formal optimization objectives rather than external constraints. It offers a replicable modeling approach for institutional investors seeking to construct portfolios that reflect the complexity of modern investment mandates. The findings reinforce that the efficient frontier is no longer two-dimensional: it is a multidimensional construct that must evolve to meet the intersecting demands of performance, risk management, and responsible capital allocation.

KEYWORDS: Modern Portfolio Theory, Multidimensional Optimization, Efficient Frontier, ESG Integration, Market SCR, Sustainable Investing, Regulatory Capital Risk, A Priori and A Posteriori Models, Pareto Efficiency, Tranche-Based Portfolio Enumeration

WORD COUNT: 28,363

Ecole de Gestion de l'Université de Liège