

ESG-Aware Asset Selection and Portfolio Optimization: A Comparative Study of Multi-Criteria Decision-Making Methods

Ștefan GĂMAN*

*Marie Skłodowska-Curie Actions (MSCA) Digital Finance,
IDA Institute Digital Assets, Bucharest University of Economic Studies, Bucharest, Romania
Poznań University of Economics and Business, Poznań, Poland*

**Corresponding author, gamanstefan18@stud.ase.ro*

Rahul TAK

*Marie Skłodowska-Curie Actions (MSCA) Digital Finance,
IDA Institute Digital Assets, Bucharest University of Economic Studies, Bucharest, Romania
takrahul24@stud.ase.ro*

Siang-Li JHENG

*Marie Skłodowska-Curie Actions (MSCA) Digital Finance,
IDA Institute Digital Assets, Bucharest University of Economic Studies, Bucharest, Romania
david86888@gmail.com*

Daniel Traian PELE

*IDA Institute Digital Assets, Bucharest University of Economic Studies, Bucharest, Romania
Institute for Economic Forecasting, Romanian Academy, Bucharest, Romania
danpele@ase.ro*

Abstract. Besides the main goal of investment, which is to obtain profit, investors are also paying more attention to the broader impact of their choices. This awareness is reflected in portfolio construction, where environmental, social, and governance (ESG) scores are included alongside financial considerations. In this study, we propose a two-stage portfolio construction framework that uses ESG factors while separating asset selection from portfolio optimization, thereby allowing the effects of sustainability factors to be examined at both levels. Using a universe of 418 U.S. equities with ESG scores and sector classifications, assets are screened and ranked using three multi-criteria decision-making methods (AHP, TOPSIS, and VIKOR), based on ESG score, expected return, and volatility, while also ensuring sector balance. In the second stage, the selected asset universes are optimized in a mean-variance framework, producing classical, green (high ESG), and brown (low ESG) efficient frontiers. The findings indicate that ESG constraints have only a limited effect on portfolio efficiency. Green portfolios show similar risk-adjusted performance to classical portfolios. The asset selection methods mainly influence the composition and risk of the portfolios. VIKOR-based portfolios are more stable, TOPSIS-based portfolios are more volatile, and AHP portfolios are between these two cases. Overall, the results suggest that portfolio outcomes depend on both the asset selection method and the way sustainability constraints are introduced during optimization.

Keywords: ESG Investing, Mean-Variance Optimization, Efficient Frontier, Asset Selection Methods.

Introduction

Sustainability has become a central part of how financial decisions are made today. Institutional investors, regulators, and data providers increasingly agree that allocating capital should not be

based only on risk and return. Now, building a portfolio is also expected to consider the broader environmental and social impacts of companies. In this context, environmental, social, and governance (ESG) factors are seen not as barriers to profitability, but as additional drivers of long-term stability (Friede et al., 2015). Therefore, the main question is not whether sustainable finance is important, but how to apply its principles using solid quantitative methods.

Within this broader landscape, green investment means allocating funds to companies that adopt sustainable practices to protect people and the environment. The goal of these practices is to meet the needs of the present generation without limiting future generations' ability to meet their needs. Several global frameworks support these efforts. For example, the Paris Agreement (United Nations, 2015a) aims to reduce greenhouse gas emissions and fight climate change. The United Nations' 17 Sustainable Development Goals (United Nations, 2015b) address a wider range of challenges, including poverty, hunger, education, peace and justice, and climate action.

If we look at the global market for sustainable funds, Europe accounts for about 85% of all sustainable fund assets, while the United States holds around 10% (Morningstar, 2025). Even though fund flows fluctuate, the overall trend shows a consistent growth in sustainable assets, as shown in Figure 1. Morningstar includes only mutual funds and exchange-traded funds that officially claim to be sustainable.




Figure 1. Global sustainable fund assets (USD)

Source: Morningstar Direct. Data as of September 2025.

Despite current political and economic challenges, ESG adoption remains high. Global tensions, economic uncertainty, and changing regulations continue to influence investors, yet most still prefer sustainable investing. According to the ESG Global Study (Capital Group, 2025), ESG integration is now at 87% worldwide, only slightly below the record levels in 2023 and 2024. Over 90% of investors who already use ESG have either kept or increased their investments in sustainable strategies during the past year. Looking forward, most investors plan to maintain or grow their ESG investments, showing strong interest in sustainable finance even in an uncertain environment.

Motivated by the ongoing growth of green finance and the increasing need for clear quantitative tools for ESG integration, this study proposes a two-stage portfolio construction framework. The framework combines asset selection with constrained portfolio optimization. In the first stage, assets are ranked by their ESG score, expected return, and volatility. This ranking is carried out using three multi-criteria decision-making methods (AHP, TOPSIS, and VIKOR), resulting in sector-balanced sets of investable assets. In the second stage, portfolios are constructed using a standard long-only mean–variance optimization approach. Sustainability constraints are applied only to the green (high-ESG) and brown (low-ESG) portfolios, while the classical

portfolios remain unconstrained and serve as a financial benchmark. For the ESG-constrained portfolios, additional diversification rules are imposed, including minimum sector exposure requirements and upper limits on investment in each asset.

For the empirical analysis, we use a large sample of 418 stocks from 11 sectors over a period of ten years to test how the proposed framework performs in realistic situations. The results show that adding ESG constraints changes the efficient frontier only slightly. Green portfolios remain close to classical portfolios in terms of both risk and return. Instead of causing big efficiency losses, ESG constraints mostly affect which assets are included in the portfolio and how risk is spread among them. These effects depend on how the investable universe is chosen, because different asset selection methods lead to small but visible differences in portfolio outcomes. Overall, the findings suggest that responsible investing shifts the efficient frontier in ways that depend on both the asset selection and optimization steps, but does not lead to a significant decrease in portfolio quality. With this framework, the study makes a data-driven link between sustainability considerations and quantitative asset management, helping to better understand how ESG principles can be included in formal portfolio optimization.

The rest of the paper is structured as follows. We begin with a review of the scientific literature in the field. We then introduce the theoretical framework and describe the model design. After that, we present the empirical analysis and discuss the key findings. The paper concludes with a summary of the main insights. All data, code, and ongoing updates are available through Quantlet  and Quantinar .

Literature review

The foundation of modern portfolio theory (MPT), developed by Markowitz (1952), is based on the maximization of expected return for a given level of risk. However, the investment world now moves more toward a "triple-bottom-line" approach. Pedersen et al. (2021) challenged the classic two-dimensional MPT by introducing the "ESG efficient frontier". Their framework shows that ESG scores have two main roles: (1) they provide information about company fundamentals, and (2) they reflect what investors prefer. By including these aspects, they show that adding ESG preferences alters the traditional Sharpe-ratio-maximizing portfolio. In this case, the "cost" of green is offset by the information gains or the utility derived from ethical alignment.

One important debate in finance is whether adding ESG factors creates "alpha" or reduces performance. Friede et al. (2015) did a large meta-analysis of more than 2,000 empirical studies and found that about 90% of them show a non-negative relationship between ESG and company financial performance. The study also points out that the "G" (Governance) factor usually has the strongest link to operational performance, while "E" (Environmental) and "S" (Social) factors are more related to reducing tail risk and meeting regulatory requirements.

Recent empirical work by Ceccarelli et al. (2023) looks at the "low carbon" premium and finds that mutual funds with better environmental ratings attract much higher inflows, especially after climate-related news. On the other hand, Bolton and Kacperczyk (2021) argue that investors want higher returns for stocks with high carbon emissions. This suggests that the "carbon premium" is a reward for the risk of future carbon taxes or stranded assets. These mixed findings highlight the need for strong asset selection methods that can balance these different signals.

Including ESG factors changes portfolio selection from a single-objective (risk/return) optimization to a more complex decision with multiple attributes. Traditional ranking methods often fail to capture the trade-offs between sustainability and financial metrics. Consequently, the

literature now uses multi-criteria decision-making (MCDM) methods to bring more mathematical structure to the process.

The Analytic Hierarchy Process (AHP), developed by Saaty (1980), helps determine the weights of different ESG and financial criteria. It uses pairwise comparisons, so investors can incorporate their expert opinions quantitatively. This method makes sure the importance of certain factors (like giving a high weight to "E" in the energy sector) is clearly shown.

The Technique for Order of Preference by Similarity to Ideal Solution (TOPSIS), developed by Hwang and Yoon (1981), ranks assets by measuring how close they are to a "positive ideal" solution (which means highest ESG and return, and lowest risk) and how far they are from a "negative ideal" solution. This approach is known for its ability to handle large datasets well.

VIKOR, introduced by Opricovic (1998), is designed to find a compromise solution. Unlike TOPSIS, which ranks based on distance to the ideal, VIKOR tries to maximize "group utility" for most options and minimize "individual regret" for the least satisfied. In finance, this method helps select assets that achieve a balanced performance across ESG and financial criteria, avoiding very poor results in any one area.

Moving from asset selection to portfolio optimization is the second step in modern ESG frameworks. Markowitz's mean-variance method is still the standard, but it is often criticized for being sensitive to estimation errors and for assuming returns are normally distributed. Rockafellar and Uryasev (2000) introduced Conditional Value-at-Risk (CVaR) as a better way to measure risk, especially for ESG-aware portfolios that are built to reduce "fat tail" climate risks. CVaR optimization looks at the expected loss above the Value-at-Risk (VaR) limit, which makes it stronger against extreme drops in the market.

The two-stage framework, first developed by researchers like Zopounidis and Doumpos (2002), has become more popular. In this approach, MCDM methods are used first to select a high-quality "green" or "sustainable" group of assets from the investment universe. Then, mathematical optimization is used to find the final weights for these assets. Recent studies, such as Taheripour et al. (2025), improved this method by adding fuzzy logic and forward-looking performance analysis to deal with the uncertainty and subjectivity in ESG data ratings.

Finally, a large part of the literature looks at the problem of "ESG rating divergence". Berg et al. (2022) found that ESG ratings from different providers often have low correlation. This happens mostly because of differences in scope, measurement, and weighting. Because of this, it is important to have studies that are neutral in their methods and compare how different selection algorithms, such as AHP and VIKOR, work under different data conditions. Also, Khan et al. (2016) introduced the idea of "materiality". They showed that companies with high ratings on important sustainability issues are more likely to impact financial results and do much better than companies with high ratings on less important issues.

In conclusion, the literature shows that including ESG is no longer only an ethical decision. It is now also an advanced strategy for managing risk and seeking higher returns. Combining MCDM for asset selection with CVaR for portfolio optimization is seen as a recent approach in financial research. This combination helps to balance the goals of profit and sustainability.

Methodology

Data and preprocessing

The empirical analysis uses daily equity price data together with firm-level ESG scores. Daily logarithmic returns are computed as

$$r_t = \ln \left(\frac{P_t}{P_{t-1}} \right) \quad (1)$$

and then annualized using the 252-day trading convention. Expected returns and volatilities are estimated as the annualized sample mean and standard deviation of returns. Only assets with both price histories and ESG scores are included in the final investment universe.

Multi-criteria asset selection

Asset selection is based on three criteria: ESG score, expected return, and volatility. Let x_{ij} denote the raw value of criterion $j \in \{ESG, \mu, \sigma\}$ for asset $i = 1, \dots, N$. ESG score and expected return are considered as benefit criteria, while volatility is considered as a cost criterion.

To ensure comparability across criteria with different units and scales, all variables are normalized to the unit interval using min-max scaling. For benefit criteria, the normalized value is defined as

$$\tilde{x}_{ij} = \frac{x_{ij} - \min_i x_{ij}}{\max_i x_{ij} - \min_i x_{ij}} \quad (2)$$

whereas for the cost criterion, normalization is performed as

$$\tilde{x}_{ij} = \frac{\max_i x_{ij} - x_{ij}}{\max_i x_{ij} - \min_i x_{ij}} \quad (3)$$

AHP-based scoring

Let $\tilde{x}_{ij}$ denote the normalized value of criterion j for asset i , obtained as described in the previous subsection. Let

$$w = (w_1, w_2, w_3) \quad \sum_{j=1}^3 w_j = 1 \quad (4)$$

be the vector of criterion weights reflecting the relative importance of ESG score, expected return, and volatility.

In its linear form, the Analytic Hierarchy Process (AHP) assigns to each asset a composite score defined as a weighted sum of normalized criteria:

$$AHP_i = \sum_{j=1}^3 w_j \tilde{x}_{ij} \quad (5)$$

Assets are ranked in descending order of AHP_i , with higher scores indicating a more favorable ESG-return-risk trade-off.

TOPSIS-based scoring

The Technique for Order Preference by Similarity to an Ideal Solution (TOPSIS) checks each asset by measuring how close it is to the best possible option and how far it is from the worst. Let $\tilde{x}_{ij}$ denotes the normalized decision matrix and define the weighted normalized matrix

$$v_{ij} = w_j \tilde{x}_{ij} \quad (6)$$

The ideal best and ideal worst solutions are given by

$$v_j^+ = \max_i v_{ij} \quad v_j^- = \min_i v_{ij} \quad (7)$$

respectively. For each asset i , the Euclidean distances to the ideal best and ideal worst points are computed as

$$d_i^+ = \sqrt{\sum_{j=1}^3 (v_{ij} - v_j^+)^2} \quad d_i^- = \sqrt{\sum_{j=1}^3 (v_{ij} - v_j^-)^2} \quad (8)$$

The TOPSIS closeness coefficient is then defined as

$$\text{TOPSIS}_i = \frac{d_i^-}{d_i^+ + d_i^-} \quad (9)$$

Assets with larger TOPSIS _{i} values are preferred, as they are closer to the ideal solution and farther from the worst-case alternative.

VIKOR-based scoring

The VIKOR method identifies compromise solutions by balancing overall group utility and individual regret. Let x_{ij} denote the raw (non-normalized) value of criterion j for asset i . For each criterion, the best and worst reference values are defined as

$$f_j^* = \begin{cases} \max_i x_{ij}, & \text{if } j \text{ is a benefic criterion,} \\ \min_i x_{ij}, & \text{if } j \text{ is a cost criterion,} \end{cases} \quad (10)$$

$$f_j^- = \begin{cases} \min_i x_{ij}, & \text{if } j \text{ is a benefic criterion,} \\ \max_i x_{ij}, & \text{if } j \text{ is a cost criterion.} \end{cases}$$

The normalized regret for each criterion is computed as

$$\frac{f_j^* - x_{ij}}{f_j^* - f_j^-}, \quad (11)$$

and aggregated into two measures:

$$S_i = \sum_{j=1}^3 w_j \frac{f_j^* - x_{ij}}{f_j^* - f_j^-} \quad R_i = \max_j \left\{ w_j \frac{f_j^* - x_{ij}}{f_j^* - f_j^-} \right\} \quad (12)$$

Here, S_i captures the overall group utility loss, while R_i represents the maximum individual regret. The final VIKOR index is defined as

$$Q_i = v \frac{S_i - \min_i S_i}{\max_i S_i - \min_i S_i} + (1 - v) \frac{R_i - \min_i R_i}{\max_i R_i - \min_i R_i} \quad (13)$$

where $v \in [0, 1]$ is a compromise parameter controlling the trade-off between collective performance and individual regret. Assets are ranked in ascending order of Q_i , with lower values indicating more desirable compromise solutions.

AHP, TOPSIS, and VIKOR offer a consistent and comparable way to select assets using multiple criteria. All three methods use the same normalized ESG, return, and risk data and give each asset a single ranking score. Although they share a common decision base, the methods are different in how they combine the information: AHP uses a simple weighted sum, TOPSIS ranks assets by how close they are to the ideal solution, and VIKOR focuses on finding a compromise between overall performance and the worst-case result. This combination helps the selection process include different decision viewpoints while keeping a clear and consistent method.

To avoid having too many assets from one sector, assets are chosen separately within each sector. For each Multi-Criteria Decision-Making (MCDM) method, only the top- k ranked assets from every sector are kept. This step creates three different, sector-balanced sets of assets, each one based on AHP, TOPSIS, or VIKOR selection.

Portfolio construction and efficient frontiers

Given the sector-balanced asset sets obtained from the MCDM selection stage, portfolio construction is performed within a mean–variance optimization framework. Let $\mu \in \mathbb{R}^n$ denote the vector of expected asset returns and $\Sigma \in \mathbb{R}^{n \times n}$ the corresponding covariance matrix, both estimated from daily log returns and annualized. Portfolio weights are denoted by $\mathbf{w} = (w_1, \dots, w_n)$, $\sum_{i=1}^n w_i = 1$, where $w_i \geq 0$ enforces a long-only allocation.

Classical mean-variance efficient frontier

The classical efficient frontier is constructed by solving a sequence of quadratic optimization problems of the form

$$\begin{aligned} \min_{\mathbf{w}} \quad & \mathbf{w}^T \Sigma \mathbf{w} \\ \text{s. t.} \quad & \mathbf{w}^T \mu = R, \\ & \sum_{i=1}^n w_i = 1, \\ & w_i \geq 0 \quad \forall i \end{aligned} \tag{14}$$

where R denotes a target expected return. By varying R over a feasible grid, the set of minimum-variance portfolios traces out the classical Markowitz efficient frontier, which represents the optimal trade-off between risk and return within the selected asset universe. This frontier serves as a natural benchmark, capturing the best achievable diversification under purely financial considerations.

To give a broader view of the possible portfolio space, many random portfolios are created using Dirichlet-distributed weights. This Monte Carlo approach creates a dense set of portfolios that fill the inside of the feasible region. These random portfolios are used as a reference to compare with the efficient portfolios, both visually and numerically.

ESG-constrained (Green) efficient frontier

To assess the impact of sustainability constraints on portfolio efficiency, an ESG-constrained (Green) frontier is constructed by augmenting the optimization problem with a portfolio-level ESG requirement. Let $\mathbf{e} \in \mathbb{R}^n$ denote the vector of asset ESG scores and let $\bar{e}$ be their cross-sectional average. The Green frontier solves

$$\begin{aligned} \min_{\mathbf{w}} \quad & \mathbf{w}^T \Sigma \mathbf{w} \\ \text{s. t.} \quad & \mathbf{w}^T \mu = R, \\ & \mathbf{w}^T \mathbf{e} \geq \bar{e}, \\ & \sum_{i=1}^n w_i = 1, \\ & 0 \leq w_i \leq \bar{w}, \\ & \mathbf{A}_s \mathbf{w} \geq \mathbf{c} \end{aligned} \tag{15}$$

where $\bar{w}$ denotes a maximum weight per asset and $\mathbf{A}_s$ encodes minimum sector exposure constraints. This formulation enforces sustainability at the portfolio level while maintaining diversification and feasibility.

Low-ESG (Brown) efficient frontier

For symmetry and comparison, a Low-ESG (Brown) frontier is constructed by replacing the ESG constraint with $\mathbf{w}^T \mathbf{e} \leq \bar{e}$, while keeping all other constraints identical. This frontier represents portfolios tilted toward lower-ESG assets and provides a counterfactual benchmark against which the cost or benefit of ESG integration can be evaluated.

In addition to optimized portfolios, two simple benchmark strategies are constructed to serve as reference points for performance comparison. These benchmarks are computed using the same selected asset universe and return estimates as the optimized portfolios.

The first benchmark is an equal-weight (EW) portfolio, in which each asset receives the same allocation,

$$w_i^{EW} = \frac{1}{n}, \quad i = 1, \dots, n \quad (16)$$

This portfolio represents a naïve diversification strategy and is commonly used as a baseline in empirical asset allocation studies.

The second benchmark is an inverse-volatility-weighted (IVW) portfolio, where portfolio weights are proportional to the inverse of individual asset volatilities,

$$w_i^{IVW} = \frac{\sigma_i^{-1}}{\sum_{j=1}^n \sigma_j^{-1}} \quad (17)$$

with σ_i denoting the annualized volatility of asset i . This strategy places lower weight on more volatile assets and provides a simple risk-based allocation rule.

Thus, for each asset selection method, the classical, green, and brown efficient frontiers are built using the same set of assets and the same return and risk estimates. This means that, in each scenario, any differences in the frontiers' position or shape are only caused by ESG constraints. Comparing frontiers across AHP, TOPSIS, and VIKOR also shows how different ways of selecting assets affect overall portfolio performance. In summary, this framework makes it possible to study the effects of ESG constraints and asset selection methods both separately and together.

Performance and risk metrics

Portfolio performance is evaluated using a concise yet comprehensive set of metrics capturing return, volatility, downside risk, and tail behavior. All measures are computed from daily log-returns and annualized using a 252-trading-day convention. Let r_t denote the daily portfolio log-return and r_f the annualized risk-free rate.

Annualized portfolio return is computed as

$$R = \exp(252 \cdot \mathbb{E}[r_t]) - 1 \quad (18)$$

while risk is measured by annualized volatility,

$$\sigma = \sqrt{252} \cdot \text{Std}(r_t) \quad (19)$$

These two quantities summarize the classical risk–return profile.

Risk-adjusted performance is assessed using the Sharpe and Sortino ratios. The Sharpe ratio is defined as

$$\text{Sharpe} = \frac{R - r_f}{\sigma} \quad (20)$$

while the Sortino ratio replaces total volatility with downside volatility, computed from negative returns only. The latter focuses explicitly on downside risk and is therefore informative when return distributions are asymmetric.

To capture path-dependent risk, maximum drawdown is computed from cumulative portfolio value based on simple returns. The Calmar ratio relates annualized return to maximum drawdown,

$$\text{Calmar} = \frac{R}{|\text{MDD}|} \quad (21)$$

providing a measure of performance relative to severe losses.

Extreme downside risk is quantified using Expected Shortfall at the 95% confidence level,

$$\text{ES}_{0.95} = \mathbb{E} [\tilde{r}_t \mid \tilde{r}_t \leq \text{VaR}_{0.05}] \quad (22)$$

where $\tilde{r}_t$ denotes simple returns. In addition, the Rachev ratio compares upper and lower tail expectations, highlighting asymmetry between extreme gains and losses.

All metrics are computed for tangency portfolios derived from the classical, Green, and Brown efficient frontiers under each asset selection method. This ensures a consistent and informative comparison of performance across ESG constraints and selection approaches.

Results and discussions

This study uses a complete dataset that combines sustainability metrics and financial market data. The main source is an ESG dataset from Kaggle, which gives firm-level ESG scores and sector labels. Daily price data for these firms was collected from Yahoo Finance for the period from 19 January 2016 to 16 January 2026, covering about 10 years.

Table 1 shows the main features of the dataset. The initial sample has 418 stocks from 11 sectors, covering a wide range of the market and different sectors. The average ESG score is 21.55, and the median is 21.1. This means sustainability profiles are quite balanced across the firms. There are both highly sustainable and low-ESG assets in the sample, showing a lot of variety in ESG characteristics.

Table 1. Database summary

Metric	Value
Number of stocks	418
Number of sectors	11
Mean ESG score	21.55
Median ESG score	21.1
Greenest asset (max ESG)	41.66
Least green asset (min ESG)	7.08

Source: Authors' own research. 

Looking at the sector-level breakdown of the dataset, stocks are not evenly spread across sectors. Financials and Industrials have the most firms in the sample, while sectors like Communication Services and Energy have much fewer stocks. Besides these differences in sector size, the average ESG scores are also different across industries. This means sector representation does not always match with sustainability performance. Figure 2 shows both the number of stocks in each sector (bar height) and the average ESG score (color intensity). This diversity shows the risk of sector concentration and sector-level ESG bias if asset selection does not use extra controls.




Figure 2. Sector distribution and average ESG scores

Source: Author's own research. 

Going further with the definition of investable asset universes, the full set of 418 stocks is reduced using structured selection criteria, as portfolio optimization over such a large universe would lead to fragmented allocations and limited interpretability. Three multi-criteria decision-making methods—AHP, TOPSIS, and VIKOR—are employed to rank assets based on ESG score, expected return, and volatility, allowing the impact of different ranking mechanisms to be isolated. Criterion importance is fixed at 0.40 for ESG, 0.35 for expected return, and 0.25 for volatility, reflecting a preference for sustainability that remains balanced with financial considerations. For VIKOR, the compromise parameter is set to $\nu = 0.5$, assigning equal weight to overall performance and worst-case regret.

To keep diversification and prevent any sector from dominating, a sector-balanced rule is used. For each selection method, only the top three assets in each sector are kept. This creates small and diversified sets of assets that still represent the wider market. It also makes it easier to compare how different selection methods affect portfolio construction and performance.

Figure 3 shows the overlap between assets chosen by each method. The central area has 25 assets selected by all three methods, which means there is a strong group of attractive stocks. Smaller groups are selected by pairs of methods: six assets are common to AHP and TOPSIS, and two are common to AHP and VIKOR. No assets are chosen only by AHP, while six assets are unique to VIKOR and two to TOPSIS. These differences come from the way each method selects assets. VIKOR and TOPSIS find a small set of assets that are not picked by the other methods. Even though there is a large common group, the presence of assets unique to each method suggests that portfolio composition and risk–return profile can change depending on the selection approach.




Figure 3. Overlap of Selected Assets Across MCDM Methods

Source: Author's own research. 

The earlier figure showed that the selected asset universes are different in their composition. Looking at their descriptive statistics, these universes also have clear differences in sustainability and financial features. Table 2 shows that, on average, TOPSIS has the most sustainable asset universe, followed by AHP. On the other hand, VIKOR has slightly lower ESG levels. However, VIKOR also has the lowest ESG variation, as shown by its standard deviation, meaning that assets selected by VIKOR have a more similar sustainability profile.

From a financial point of view, the VIKOR universe has the highest average expected return, while all three methods have the same level of volatility. This means that VIKOR gives more weight to return than to ESG, compared to AHP and TOPSIS. This is also seen in the ESG/volatility ratio, where AHP has the highest value. This means AHP achieves higher sustainability for each unit of risk.

Table 2. Asset universe statistics

	AHP	TOPSIS	VIKOR
Mean ESG score	29.67	30.00	28.86
Median ESG score	29.93	31.79	29.26
ESG std. dev.	6.52	6.49	6.01
Mean expected return	0.16	0.16	0.17
Mean volatility	0.30	0.30	0.30
ESG/Volatility ratio	98.55	97.87	94.69
Unique asset share	0.00	0.06	0.18

Source: Author's own research. 

Overall, the results show that sustainable portfolio efficiency depends not only on ESG constraints during optimization, but also on the way sustainability, return, and risk information are used to build the investable asset universe.

Looking at the efficient frontiers, Figure 4 shows that adding ESG and sector constraints moves the green and brown frontiers only a little inward compared to the classical Markowitz

benchmark. The biggest differences appear at higher risk levels. The fact that the low- and high-ESG frontiers are close to the classical frontier means that sustainability constraints do not have a big cost in terms of risk–return efficiency. Also, ESG-aware optimization gives much better performance than simple benchmarks, like equal-weight and inverse-volatility portfolios, without losing efficiency.




Figure 4. Classical vs Green vs Brown frontiers (AHP, TOPSIS, VIKOR)

Source: Author's own research.

Zooming in on the tangency portfolios (Figure 5), the results indicate that the choice of asset selection method has a limited, but not negligible, impact on portfolio performance. Under AHP and VIKOR, the tangency portfolios of the classical, green, and brown frontiers are tightly clustered in risk–return space, suggesting very similar risk-adjusted outcomes across ESG specifications. In these two cases, differences between portfolios are minor and mainly reflect variations in asset composition rather than substantial changes in efficiency.

In contrast, the TOPSIS-based universe shows a different pattern. The classical tangency portfolio is clearly further to the right and higher than the green and brown tangency portfolios, meaning it has both higher volatility and higher expected return. Because of this, there is a wider gap between tangency portfolios under TOPSIS compared to AHP and VIKOR. However, even here, the overall differences are still moderate, and the ESG-constrained tangency portfolios remain close to the classical solution.




Figure 5. Tangency portfolio (AHP, TOPSIS, VIKOR)

Source: Author's own research.

Table 3 presents the performance metrics for the tangency portfolios built with each asset selection method and for each type of frontier. As expected, classical portfolios have the highest expected returns and Sharpe ratios in all three methods, because they do not have sustainability constraints. However, the difference in performance between classical and green portfolios is small. For AHP, TOPSIS, and VIKOR, green portfolios keep Sharpe ratios close to those of classical portfolios. This result shows that adding ESG constraints only slightly reduces risk-adjusted performance.

Differences between asset selection methods appear primarily as small changes in risk and return, not as big changes in efficiency. Portfolios made with VIKOR usually have a bit lower volatility and competitive Sharpe and Sortino ratios, which means they have a slightly better risk-adjusted profile. On the other hand, TOPSIS portfolios tend to have higher volatility and can face bigger drawdowns. This explains why their risk-adjusted metrics are a bit lower, even though they sometimes have higher expected returns. AHP portfolios are usually in between, showing a balance between risk and return.

From a downside-risk perspective, green portfolios often perform comparably to, or better than, brown portfolios. This is reflected in higher Sortino ratios and less severe maximum drawdowns in several cases. Importantly, tail-risk measures such as the 95% Expected Shortfall and Rachev ratio remain very similar across classical, green, and brown portfolios, indicating that ESG constraints do not materially increase exposure to extreme losses.

Taken together, the results show that different asset selection methods slightly shift the risk and return of optimal portfolios and portfolio optimization preserves a high degree of efficiency across methods, allowing ESG-aligned portfolios to remain competitive.

Table 3. Metrics of portfolios

Method		Ret	Vol	Sharpe	Sortino	MMD	Calmar	ES ₉₅	RR ₉₅
AHP	Classical	0.32	0.19	1.49	1.94	-0.26	1.24	-0.027	1.05
	Brown	0.29	0.19	1.39	1.73	-0.30	0.97	-0.027	1.00
	Green	0.29	0.18	1.43	1.79	-0.27	1.09	-0.026	1.01
TOPSIS	Classical	0.34	0.21	1.47	1.90	-0.31	1.12	-0.030	1.05
	Brown	0.29	0.19	1.34	1.67	-0.33	0.88	-0.028	0.99
	Green	0.31	0.20	1.38	1.70	-0.33	0.94	-0.029	0.99
VIKOR	Classical	0.31	0.19	1.50	1.94	-0.25	1.26	-0.027	1.04
	Brown	0.28	0.18	1.41	1.77	-0.28	1.03	-0.026	1.00
	Green	0.30	0.18	1.44	1.82	-0.27	1.11	-0.026	1.00

Source: Author's own research. 

This study also has some limitations. First, ESG scores are used as fixed values from only one data source, which may not show the changing and varied nature of sustainability assessment. Second, the classical mean–variance framework makes distributional assumptions that might underestimate tail risk, even though downside-risk metrics are included. Third, the results depend on choices like criterion weights, sector balancing, and using only U.S. stocks, which may limit how widely the findings apply. Future research could improve on these points by using time-varying ESG data, different risk-based optimization models, and a larger, more diverse set of assets.

Conclusion

This study examines the interaction between ESG constraints, asset selection methodologies, and mean–variance portfolio optimization. Using AHP, TOPSIS, and VIKOR to construct sector-balanced investable universes, and comparing classical, green, and brown efficient frontiers, the analysis reveals several consistent insights.

First, adding ESG constraints does not significantly lower risk–return efficiency. For all selection methods, the green efficient frontiers stay close to the classical benchmark, with only small inward shifts when risk is high. The Sharpe and Sortino ratios for ESG-constrained tangency portfolios are similar to those for unconstrained portfolios. This means that sustainability can be included without notable loss in performance.

Second, the asset selection method has a smaller but still important impact on portfolio results. Most assets are shared across AHP, TOPSIS, and VIKOR, but each method adds some unique assets, leading to different risk and return profiles. Portfolios built with VIKOR tend to have a bit lower volatility and good risk-adjusted performance, while TOPSIS portfolios often show higher volatility and larger drawdowns. AHP gives results that fall in between, balancing both sustainability and financial goals. These findings show that ESG portfolio performance depends not only on the optimization step, but also on how information about sustainability, return, and risk is combined during the asset selection stage.

Overall, the results show that ESG investing can be implemented using standard quantitative methods without harming portfolio performance. ESG constraints mainly change how assets are allocated within efficient portfolios, rather than reducing overall efficiency, while different selection methods lead to small but consistent performance differences.

Acknowledgements

This paper is supported through the project “IDA Institute of Digital Assets”, CF166/15.11.2022, contract number CN760046/23.05.2023; the project “AI for Energy Finance (AI4EFin)”, CF162/15.11.2022, contract number CN760048/23.05.2023, financed under the Romania’s National Recovery and Resilience Plan, Apel nr. PNRR-III-C9-2022-I8; and the Marie Skłodowska-Curie Actions under the European Union's Horizon Europe research and innovation program for the Industrial Doctoral Network on Digital Finance, acronym: DIGITAL, Project No. 101119635. Also, the paper was co-financed by The Bucharest University of Economic Studies during the PhD program.

During the preparation of this work, the author(s) utilized Grammarly to enhance readability and language clarity. Following the use of this tool, the author(s) meticulously reviewed and revised the content as necessary and assume(s) full responsibility for the final content of the publication.

References

- Berg, F., Koelbel, J. F., & Rigobon, R. (2022). *Aggregate Confusion: The Divergence of ESG Ratings*. *Review of Finance*, 26(6), 1315–1344.
- Bolton, P. and Kacperczyk, M. (2021). Do investors care about carbon risk?. *Journal of financial economics*, 142(2), pp.517–549.
- Bonami, P. and M. A. Lejeune (2009). *An exact solution approach for portfolio optimization problems under stochastic and integer constraints*. *Operations research* 57 (3), 650–670.
- Capital Group. (2025). *ESG global study: Fifth edition (2025)*.

- Ceccarelli, M., Ramelli, S., & Wagner, A. F. (2023). *Low Carbon Mutual Funds. Review of Finance*, 28(1), 45-74.
- Chang, T.-J., N. Meade, J. E. Beasley, and Y. M. Sharaiha (2000). *Heuristics for cardinality constrained portfolio optimisation*. *Computers & Operations Research* 27 (13), 1271–1302.
- Friede, G., T. Busch, and A. Bassen (2015). *Esg and financial performance: aggregated evidence from more than 2000 empirical studies*. *Journal of sustainable finance & investment* 5 (4), 210–233.
- Hwang, C. L., & Yoon, K. (1981). Methods for multiple attribute decision making. In *Multiple attribute decision making: methods and applications a state-of-the-art survey* (pp. 58-191). Berlin, Heidelberg: Springer Berlin Heidelberg.
- Khan, M., Serafeim, G., & Yoon, A. (2016). *Corporate Sustainability: First Evidence on Materiality*. *The Accounting Review*, 91(6), 1697-1724.
- Markowitz, H. M. (1952). *Portfolio selection*, the journal of finance. 7 (1). N 1, 71–91.
- Morningstar Sustainability. (2025). *Global sustainable fund flows: Q3 2025 in review*. Morningstar Sustainability.
- Oprić, S. (1998). Multicriteria optimization of civil engineering systems (in Serbian, Višekriterijumska optimizacija sistema u građevinarstvu). Belgrade: Faculty of Civil Engineering.
- Pástor, L., Stambaugh, R. F., & Taylor, L. A. (2021). *Sustainable investing in equilibrium*. *Journal of Financial Economics*, 142(2), 550-571.
- Pedersen, L. H., S. Fitzgibbons, and L. Pomorski (2021). *Responsible investing: The esg-efficient frontier*. *Journal of financial economics* 142 (2), 572–597.
- Rockafellar, R. T., & Uryasev, S. (2000). *Optimization of conditional value-at-risk*. *Journal of Risk*, 2, 21-42.
- Saaty, T. L., & Wind, Y. (1980). Marketing applications of the analytic hierarchy process. *Management science*, 26(7), 641-658.
- Taheripour, E., Sadjadi, S. J., & Amiri, B. (2025). *A multi-criteria approach to ESG-based portfolio optimization*. *Scientific Reports*, 15(1).
- United Nations (2015a). *Paris Agreement*. Available at: https://treaties.un.org/doc/Treaties/2016/02/20160215%2006-03%20PM/Ch_XXVII-7-d.pdf.
- United Nations (2015b). *Sustainable Development Goals*. Available at: <https://sdgs.un.org/goals>.
- Zopounidis, C., & Doumpos, M. (2002). Multicriteria classification and sorting methods: A literature review. *European Journal of Operational Research*, 138(2), 229-246.