STATISTICS, OPTIMIZATION AND INFORMATION COMPUTING
Stat., Optim. Inf. Comput., Vol. **x**, Month 202x, pp 0–22.
Published online in International Academic Press (www.IAPress.org)

Multi-Objective Genetic Optimization of Sharia-Compliant Portfolios: From Binary Screening to Compliance Scoring

AABBAD Zineb [1,*], CHIADMI Mohammed Salah [1], ABOULAICH Rajae [2]

[1] *LERMA, Ecole Mohammadia d'Ingenieurs, Mohammed V University , Rabat, Morocco*
[2] *Africa Business School,UM6P, Rabat, Morocco*

Abstract In this paper, we propose an integrated multi-objective optimization framework for Sharia-compliant portfolio selection that shifts asset screening from a traditional binary cut into a continuous quantitative compliance scoring paradigm. Departing from existing literature that treats continuous scores merely as descriptive ranking metrics, the primary contribution of this study is embedding the continuous compliance score directly into an evolutionary multi-objective optimization engine via the reference-direction algorithm NSGA-III.By formulating this score as an explicit third objective, the framework simultaneously maximizes expected return, minimizes risk, measured either by the variance or by the Conditional Value-at-Risk, and maximizes Sharia compliance, under realistic cardinality and position-size constraints that make the problem non-convex and motivate the evolutionary approach. The framework systematically mitigates boundary proximity risk to guard against threshold breaches and forced divestment. To demonstrate the efficiency of this continuous paradigm, we compare it, under identical constraints, with the traditional binary-screening approach, in which the same Sharia filters are applied and only return and risk are optimized, using fundamental and daily price data from the S&P 500 universe. The algorithmic configuration is selected by a grid search parameter study, and NSGA-III is benchmarked against NSGA-II. At equal risk, the continuous model reaches substantially higher compliance scores at an explicit cost in expected return. It also roughly halves the share of capital invested in firms close to the compliance thresholds, and these conclusions hold whether risk is measured by the variance or by the CVaR.

Keywords Multi-Objective Optimization, NSGA-III, Sharia-Compliant Portfolios, Continuous Compliance Scoring, Conditional Value-at-Risk, Cardinality Constraints, Das and Dennis Method, Portfolio Selection

AMS 2010 subject classifications 90C29, 91B28

DOI: 10.19139/soic-2310-5070-4222

1. Introduction

Sharia-compliant investing has witnessed steady growth, as more investors seek to align their ethical values with their financial objectives. Traditionally, the construction of Sharia-compliant portfolios has relied on binary screening frameworks, in which firms are classified as either compliant or non-compliant based on predefined thresholds applied to financial ratios, revenue sources, and sectoral activities. Such threshold-based methodologies are widely adopted by major institutions, including S&P, Dow Jones, and MSCI [1].

While computationally straightforward, this binary methodology suffers from a critical structural deficiency: it introduces an analytical limitation by treating all assets within the feasible region as ethically equivalent, irrespective of their proximity to the boundary limits [2]. Consequently, an asset operating marginally below a non-compliance threshold is assigned an identical status to an asset operating with substantial compliance margins.This binary approach fails to account for boundary proximity risk [3]: a firm sitting right next to the legal limit poses a

*Correspondence to: AABBAD Zineb (Email: zinebaabbad@research.emi.ac.ma). LERMA, Ecole Mohammadia d'Ingenieurs, Avenue Ibn Sina, B.P. 765, Rabat-Agdal

ISSN 2310-5070 (online) ISSN 2311-004X (print)
Copyright © 202x International Academic Press

significant risk of breaching the threshold during normal market movements. Such a breach forces the fund manager to hastily liquidate the asset at a loss to maintain compliance [4]. Consequently, this lack of granularity limits the investor's capacity to execute nuanced trade-offs between financial performance and ideological adherence.

To resolve this limitation, contemporary literature has initiated a transition toward quantitative scoring mechanisms designed to model compliance as a continuous spectrum. For instance, Hameed and Muneeza [2] formulated the Shariah Compliance Degree (SCD)—a percentage-based index derived from five structural screening ratios within the KMI Index—to provide a granular measure of corporate compliance. Similarly, Alnamlah et al. [5] developed the Weighted Average Shariah Compliance (WASC) score, which maps financial ratios and sectoral attributes into relative compliance percentiles. However, existing studies limit these scores to asset ranking, rather than an active input in portfolio construction. Building on these foundations, this paper presents a framework that incorporates continuous compliance scores directly into portfolio optimization. We first eliminate strictly non-permissible equities using deterministic sector and financial ratio filters to guarantee that the optimization algorithm only evaluates legally compliant assets. For the surviving asset universe, we formalize a continuous compliance metric defined as the normalized average of the financial ratios relative to their legal boundaries, after which the calculated compliance scores of the surviving universe are fed directly into the multi-objective optimization algorithm to guide asset allocation.

The contribution of this study therefore does not lie in the computation of a continuous compliance score per se, which has been proposed in the SCD, WASC and CSCI frameworks [2, 5, 6], but in its embedding as an optimizable objective of a multi-objective portfolio optimization problem solved by an evolutionary algorithm. In existing studies the score is computed to rank or describe assets; here it shapes the allocation itself and produces an explicit return–risk–compliance efficient surface. Specifically, the paper (i) provides a fully specified and reproducible threshold-relative score; (ii) embeds it as a third objective of an NSGA-III model with a risk objective and realistic portfolio constraints; (iii) compares the resulting allocations with the industry practice of binary screening followed by return–risk optimization, under identical constraints and at equal risk levels; and (iv) shows that these results hold for two alternative risk measures, the variance and the CVaR, and quantifies the exposure of the portfolios to firms close to the compliance thresholds.

Following this quantitative scoring phase, the asset allocation problem is formulated as an evolutionary multi-objective optimization problem. We deploy the Non-dominated Sorting Genetic Algorithm III (NSGA-III) to resolve the underlying non-linear trade-offs across a tri-criterion space: maximizing expected portfolio returns, minimizing risk, measured either by the variance or by the Conditional Value-at-Risk (CVaR), and maximizing the integrated Sharia compliance score, subject to bounds on the number of holdings and on individual position sizes. By embedding this continuous compliance score as an explicit third objective within NSGA-III, the algorithm actively favors assets that stay safely below the threshold limits, providing an algorithmic solution to the boundary proximity risk. The optimization process yields a multi-dimensional Pareto front that systematically defines the non-dominated trade-offs available to the decision-maker.

The remainder of this paper is organized as follows. The rest of this section reviews Sharia screening methodologies and multi-objective optimization in finance. Section 2 details the data, the continuous scoring function, the estimation of return and risk, and the two versions of the constrained optimization problem. The following section describes the NSGA-III implementation, including the treatment of the portfolio constraints. Section 4 reports the selection of the algorithmic parameters and the comparison with NSGA-II. Section 5 compares binary screening and continuous scoring under identical constraints, through an overview of the portfolios obtained with the two approaches and the two risk measures, their Pareto fronts, a comparison at equal risk and representative investor profiles. Section 6 concludes and discusses the limitations of the study and directions for future research.

1.1. Sharia screening methodologies

The governance rules dictating Sharia compliance for publicly traded equities are not explicitly defined in foundational Islamic legal texts [7], leading to regional variations and diverse methodological interpretations among legal scholars. Although the literature [1, 8, 9] presents varied operational criteria, the industry has largely converged on a standardized two-stage screening pipeline. The foundations of these screening norms were examined by Khatkhatay and Nisar [10], who compared the screening criteria of three Sharia-screening

organizations, assessed their impact on Bombay Stock Exchange data and proposed an alternative set of norms, while Derigs and Marzban [11] analyzed the effect of Sharia screening strategies on portfolio optimization and proposed a paradigm in which compliance is attributed to the portfolio rather than to the individual assets. The first phase consists of a qualitative sectoral filter designed to eliminate corporations engaged in non-permissible (*haram*) commercial activities. The second phase applies a quantitative financial filter to restrict exposure to interest-based leverage (*riba*) and speculative financial operations.

Qualitative screening systematically purges companies deriving capital from prohibited sectors, including alcohol production, gambling, conventional financial services, and weapons manufacturing. However, as noted by Md. Hashim et al. [12], corporate entities generally fall into three distinct operational categories: purely non-compliant, purely compliant, and mixed-activity (hybrid) enterprises. Hybrid corporations represent entities whose primary operations are permissible, yet they maintain secondary exposures to non-compliant revenue streams, such as a hospitality firm that operates a luxury hotel containing conventional beverage services.

Given that a strict requirement to include exclusively include businesses whose operations are entirely Sharia compliant would severely compromise asset diversification and restrict the investment universe, modern Islamic jurisprudence permits the inclusion of hybrid enterprises, provided their non-permissible exposures remain bounded within conservative tolerance limits [13, 14]. This aligns with major Islamic index standards. In practice, after the qualitative sectoral filter is applied, the remaining firms are subjected to a quantitative filter ensuring that non-permissible revenue accounts for less than 5% of gross income. Furthermore, three distinct financial ratios—measuring leverage, cash, and liquidity—are evaluated to ensure the firm's capital structure avoids excessive interest-bearing debt or illiquid positions, as Sharia law discourages the speculative trading of money-equivalent assets [15].

Table 1 provides a structural comparison of the financial screening ratios mandated by major global Islamic index providers, demonstrating the variations in denominator selection (Market Capitalization vs. Total Assets) and numerator composition.

Index Provider	Debt Ratio	Cash Ratio	Liquidity Ratio
Dow Jones Islamic Market (DJIM) [16]	$\frac{\text{Total Debt}}{\text{24M Avg. Market Cap}} < 33\%$	$\frac{\text{Cash + Interest-Bearing Sec.}}{\text{24M Avg. Market Cap}} < 33\%$	$\frac{\text{Accounts Receivable}}{\text{24M Avg. Market Cap}} < 33\%$
S&P Shariah Indices [17]	$\frac{\text{Debt}}{\text{36M Avg. Market Cap}} < 33\%$	$\frac{\text{Cash + Interest-Bearing Sec.}}{\text{36M Avg. Market Cap}} < 33\%$	$\frac{\text{Accounts Receivable}}{\text{36M Avg. Market Cap}} < 49\%$
FTSE Yasaar Global Equity [18]	$\frac{\text{Total Debt}}{\text{Total Assets}} \leq 33.3\%$	$\frac{\text{Cash + Interest-Bearing Sec.}}{\text{Total Assets}} \leq 33.3\%$	$\frac{\text{Receivable + Cash}}{\text{Total Assets}} \leq 50\%$
MSCI Islamic Index Series Standard [19]	$\frac{\text{Total Debt}}{\text{Total Assets}} < 33.3\%$ (Entry 30%, Exit 35%)	$\frac{\text{Cash + Interest-Bearing Sec.}}{\text{Total Assets}} < 33.3\%$ (Entry 30%, Exit 35%)	$\frac{\text{Receivables + Cash}}{\text{Total Assets}} < 70\%$ (Entry 46%)
MSCI Islamic Index Series M-Series [19]	$\frac{\text{Total Debt}}{\text{36M Avg. Market Cap}} < 33.3\%$	$\frac{\text{Cash + Interest-Bearing Sec.}}{\text{36M Avg. Market Cap}} < 33.3\%$	$\frac{\text{Receivables + Cash}}{\text{36M Avg. Market Cap}} < 49\%$

Table 1. Comparison of Sharia Financial Ratio Screens Across Major Index Providers.

Beyond standard screening, recent quantitative literature has expanded into continuous scoring paradigms. For instance, Hameed and Muneeza [2] constructed a Sharia Compliance Degree (SCD) utilizing five asset ratios from the Karachi Stock Exchange KMI index to map detailed corporate profiles. Similarly, Alnamlah et al. [5] proposed the Weighted Average Shariah Compliance (WASC) metric to assign compliance percentiles, while Qadi et al. [6] formalized a continuous Shariah Compliance Index (CSCI) bounded strictly within the interval $[0, 1]$. In all three cases, however, the score is used as a descriptive or ranking device and is not optimized jointly with the financial objectives of the portfolio, which is the gap addressed in this paper.

Multicriteria methods have also been applied to the selection of Sharia-compliant portfolios, for instance on the Jakarta Islamic Index [20], but compliance is then treated as a condition of eligibility rather than as a criterion of the allocation.

1.2. Portfolio optimization

The mathematical foundation of modern portfolio selection was established by Markowitz [21], who formalized asset allocation as a bi-criterion optimization problem balancing expected return and variance. Subsequent portfolio literature has significantly extended this framework by incorporating higher-order statistical moments and robust risk metrics to handle non-Gaussian asset distributions [22]. These extensions include optimization models integrating skewness [23], kurtosis [24, 25], Value-at-Risk (VaR) [26, 27], Conditional Value-at-Risk (CVaR) [28], and non-linear reward-to-risk measures such as the Sharpe ratio [29].

In contemporary financial markets, investment mandates must frequently satisfy multidimensional criteria that extend beyond basic risk-return trade-offs. Institutional mandates must simultaneously reconcile liquidity constraints [30], regulatory boundaries, and sustainability or non-financial ethical metrics. This operational complexity requires robust decision-making frameworks capable of solving multi-objective optimization (MOO) challenges [31, 32].

As highlighted in the exhaustive review by Salo et al. [33], mapping 50 years of development in portfolio optimization, MOO has emerged as the dominant mathematical approach for evaluating conflicting investment objectives. Historically, multi-objective problems were solved by reducing the vector optimization space into a single-objective equivalent via scalarization, the ε-constraint method, or goal programming techniques.

While analytical scalarization is computationally inexpensive, it struggles to resolve non-convex or high-dimensional trade-off topologies. To overcome this limitation, evolutionary computation has been extensively adopted. Genetic Algorithms (GAs) [34] offer a powerful heuristic alternative for solving complex multi-objective optimization problems [32]. Guided by principles of natural selection, GAs iteratively evolve a population of candidate portfolio weight vectors through selection, crossover, and mutation operators to locate global optima across non-linear search landscapes [35].

Evolutionary methods are particularly relevant when realistic portfolio constraints are introduced. Limits on the number of assets held and minimum or maximum position sizes turn portfolio selection into a mixed-integer, non-convex problem for which exact methods become computationally demanding [36].Chang et al. [37] first addressed this cardinality-constrained problem with genetic algorithms that have since been widely used for this class of problems. In such settings, Lagrangian and weighted-sum scalarizations no longer guarantee the recovery of the whole efficient set, whereas a population-based algorithm approximates it in a single run.

To maintain adequate population diversity along the efficient frontier, Deb et al. [38] introduced the Non-dominated Sorting Genetic Algorithm II (NSGA-II). NSGA-II uses a fast non-dominated sorting mechanism and a crowding-distance sorting operator to prevent localized population collapse, making it a benchmark tool in quantitative finance. However, as the objective space expands beyond two dimensions, the effectiveness of the crowding distance metric drops significantly. To rectify this scaling limitation, Deb and Jain [39] developed NSGA-III, substituting local crowding distance with a structured grid of reference directions. This adaptation ensures uniform solution distribution across high-dimensional Pareto-optimal surfaces, establishing an essential computational methodology for complex multi-criteria financial allocation [35]. Recent theoretical work supports this advantage: Wietheger and Doerr [40] prove that, with sufficiently many reference points, NSGA-III computes the complete Pareto front of a three-objective benchmark, on which NSGA-II had been shown to lose a constant fraction of the front. NSGA-III has accordingly been applied to portfolio selection with higher-order moments [32] and to multi-period portfolio optimization under real-world constraints [41].

integration of a continuous non-financial score into portfolio selection has a close parallel in the responsible-investing literature. Pedersen et al. [42] characterize the ESG-efficient frontier, i.e. the highest Sharpe ratio attainable for each level of a continuous ESG score. Hirschberger et al. [43] and Utz et al. [44] compute the efficient surface of mean–variance problems extended by a sustainability criterion, and more recent studies approximate mean–variance–ESG fronts with multi-objective evolutionary algorithms [45, 46]. Sharia compliance scoring can be viewed as a faith-based instance of this broader responsible-investing problem, with the specificity that the score is anchored to regulatory thresholds rather than to the rating of an external provider.

2. Methodology

2.1. Methodology description

The methodology employed in this study introduces a novel computational framework that transcends traditional, binary Sharia screening by integrating a continuous compliance scoring mechanism with advanced evolutionary multi-objective optimization. Historically, portfolio construction in Islamic finance has relied entirely on a rigid "pass/fail" dichotomy: assets are either compliant or non-compliant, with no distinction made between a company hovering dangerously close to the maximum permissible debt limit and one operating entirely debt-free. Our approach bridges the critical gap between strict religious compliance and continuous risk management by employing a sequential "filter-then-score" methodology. First, a rigid quantitative filter defines the absolute feasible universe, ensuring that the optimization algorithm only evaluates assets that strictly satisfy Islamic jurisprudence. However, from a portfolio management perspective, holding borderline-compliant assets introduces significant forced-liquidation risk [4]; a minor market downturn could spike a firm's debt-to-asset ratio past the 33% threshold [3], forcing the fund manager to sell at a loss. Therefore, the second step transforms these compliance margins into a continuous compliance score. By maximizing this score during optimization, the framework establishes a critical "margin of safety," actively rewarding financial purity and shielding the portfolio from boundary breaches.

In practice, the framework follows five steps. First, a qualitative filter removes the firms whose activity is prohibited by Sharia, such as alcohol, gambling or conventional banking. Second, a quantitative filter keeps only the firms whose debt, cash, receivables and interest-income ratios are below the accepted thresholds. Third, for each remaining firm, we measure how far each ratio is from its threshold and average these distances into a single Sharia compliance score between 0 and 1: a firm with no debt receives a high score, while a firm close to the limit receives a low one. Fourth, the expected return of each asset and its daily returns over the sample period are computed from market prices; they are used to measure the risk of a portfolio, either by its variance or by its Conditional Value-at-Risk (CVaR), which is the average loss on the worst days.

Finally, the portfolios are optimized with the Non-dominated Sorting Genetic Algorithm III (NSGA-III), an extension of NSGA-II designed for problems with three or more objectives. For each risk measure, two types of portfolios are built under the same limits on the number and size of positions: portfolios that maximize the expected return and minimize the risk only, as in current binary screening, and portfolios that also maximize the compliance score.

Comparing these portfolios shows what the score adds, and whether this result holds with both risk measures. By treating the aggregate portfolio compliance score as an explicit third objective, maximized simultaneously with the expected return while the risk is minimized, and by imposing realistic bounds on the number of assets and size of positions, the algorithm generates a three-dimensional Pareto front of investable portfolios. Several portfolios are then compared, built with two different risk calculations, the variance and the CVaR, and, for each of them, with the compliance score used either as an objective or only as a screening filter. This comparison shows whether the benefit of the score holds whatever the way risk is measured, and provides investors with a granular capacity to trade off risk, expected return and the strictness of their Sharia adherence.

2.2. Data

The initial investment universe is based on the S&P 500 universe, which offers broad market coverage and is known to include a considerable share of Sharia-compliant firms. This choice is further motivated by the existence of Islamic index families constructed on the same foundations, such as the *S&P 500 Shariah Index*, which applies Sharia filters to the U.S. large-cap universe, and the *S&P Pan Arab Shariah Index*, which is derived from the S&P Pan Arab Index covering equities from the MENA region [47, 48]. The universe is then refined through a two-step filtering procedure in line with widely accepted Sharia standards.

The first step applies a **qualitative screen**, excluding firms engaged in sectors considered non-permissible, such as conventional banking, insurance, gambling, alcohol, tobacco, and defense. Furthermore, companies with controversial involvement in areas such as adult content, controversial weapons, or nuclear activities are also removed to preserve ethical alignment [19].

Second, a **quantitative screen** applies financial ratio thresholds to confirm structural compliance:

- **Debt ratio** $\leq 33\%$
- **Cash and interest-bearing assets ratio** $\leq 33\%$
- **Receivables ratio** $\leq 33\%$
- **Interest income ratio** $\leq 5\%$

These thresholds reflect widely adopted Sharia standards.

Fundamental data consist of annual balance sheets and income statements, together with their filing dates, for the 503 constituents of the S&P 500. Daily closing prices cover the period from 23 July 2024 to 18 July 2025 and are restricted to NYSE trading days, which yields 248 daily returns per asset.

Of the 503 constituents, 383 pass the qualitative screen and 328 of them report the interest income required by the quantitative screen for their most recent fiscal year. The quantitative screen excludes 153 firms on the debt ratio, 21 on the cash ratio, 8 on the receivables ratio and 25 on the interest-income ratio (a firm may fail several screens), leaving 148 financially compliant firms, of which 146 have a complete price history and constitute the investment universe. All 146 retained assets have a strictly positive compliance score, ranging from 0.376 to 0.889 (median 0.640, interquartile range $[0.571, 0.706]$); 13 assets score below 0.5, 93 between 0.5 and 0.7 and 40 above 0.7, so that the score discriminates between assets that all pass the binary screen.

Point-in-time fundamental data with comparable coverage for Muslim-majority markets (e.g. Bursa Malaysia, Saudi Tadawul or the Indonesia Stock Exchange) were not available to the authors; the empirical analysis is therefore restricted to the S&P 500 universe.

2.3. Sharia compliance score

To quantify Sharia compliance, we normalize each financial ratio relative to its screening threshold, thereby placing heterogeneous measures on a common, dimensionless scale that can be aggregated into a composite score [49]. Normalization also mitigates differences in variation ranges across ratios [50], which is crucial here because the admissible thresholds differ significantly (for example, a 33% limit for debt ratio versus a 5% limit for interest income).

Various normalization techniques have been developed. One of the most used approaches is min–max linear normalization, which rescales values into the [0,1] interval relative to their minimum and maximum bounds [51]. Other common methods are: Z-score normalization that standardizes values by subtracting the mean and dividing by the standard deviation, thereby highlighting deviations from average behavior [52], vector normalization which rescales values by dividing them by the Euclidean norm of the vector [53], and finally logarithmic non-linear normalization which uses logarithmic functions to scale data, which is helpful when values within a criterion are vastly different from one another [50].

Min–max linear normalization is particularly effective when the lower and upper bounds of the data remain relatively stable over time [54]. This approach is well-suited to the current study, as our financial ratios are naturally bounded by the defined Sharia thresholds.

Prior to quantitative evaluation, a qualitative filter removes non-compliant assets. A firm is disqualified if it belongs to the financial services sector, if its industry belongs to a list of prohibited activities (diversified and regional banks, credit services, insurance, capital markets, tobacco, alcoholic beverages, gambling and casinos, aerospace and defense, and entertainment), or if it is flagged for involvement in adult entertainment, alcohol, controversial weapons, small arms, gambling, military contracting, nuclear activities or tobacco. For qualitatively eligible firms, the ratios are computed from the most recent annual financial statements. Let asset i ($i = 1, \ldots, n$) be evaluated across $M = 4$ financial criteria, with thresholds $\tau_1 = \tau_2 = \tau_3 = 0.33$ and $\tau_4 = 0.05$. The raw financial ratios $R_{i,j}$ are constructed from balance sheet and income statement items as detailed in Table 2.

A score is obtained from each observed ratio via a linear threshold-based normalization, bounded to the unit interval:

$$s_{i,j} \;=\; \min\Big(1,\; \max\Big(0,\; 1 - \frac{R_{i,j}}{\tau_j}\Big)\Big), \qquad j = 1, \ldots, 4.$$

Table 2. Screening ratios computed from accounting items and their thresholds.

Index (j)	**Screening criterion**	**Ratio** ($R_{i,j}$)	**Threshold** (τ_j)
1	Debt	$\frac{\text{Total short- and long-term debt}_i}{\text{Total assets}_i}$	0.33
2	Cash and interest-bearing assets	$\frac{\text{Cash and short-term investments}_i}{\text{Total assets}_i}$	0.33
3	Accounts receivable	$\frac{\text{Net receivables}_i}{\text{Total assets}_i}$	0.33
4	Impermissible income	$\frac{\text{Interest income}_i}{\text{Total revenue}_i}$	0.05

The overall compliance score for a stock i is the simple average(equal-weighted arithmetic mean) of these four components:

$$S_i \;=\; \frac{1}{4}\sum_{j=1}^{4} s_{i,j}.$$

By construction, each component $s_{i,j}$ equals 1 when the corresponding ratio is 0 and equals 0 exactly at its threshold. This makes the composite score S_i sensitive to how comfortably a firm sits within the financial screens. Since the score is computed only for firms that pass the binary screen ($R_{i,j} \le \tau_j$ for all j), every component lies in $[0, 1]$. No negative score therefore enters the optimization, and the score of a portfolio, which is a weighted average of the asset scores, also lies in $[0, 1]$.

Robustness to outliers. Unlike the usual min–max normalization, which rescales each ratio by its smallest and largest observed values, the mapping above uses fixed regulatory bounds (0 and τ_j), so that an extreme value in one firm cannot change the scores of the other firms. In the retained universe, the largest observed ratios are 0.330 (debt), 0.312 (cash), 0.326 (receivables) and 0.046 (interest income), all within their thresholds, so that no further treatment of extreme values was needed.

2.4. Justification of the Linear Robustness Indicator for Sharia Compliance

To justify the use of the linear min–max normalization for our Sharia compliance objective, we evaluated three different scaling methods: truncated, logarithmic and linear min–max. Our goal was to find a scoring metric that accurately measures compliance while functioning perfectly within the NSGA-III optimization algorithm.

First, we discarded the Truncated method. While it correctly mimics binary rules (a stock is either compliant or not), it assigns a flat, perfect score of 1.0 to any stock safely below the limit. If many stocks share the same perfect score, NSGA-III can no longer distinguish them through their objective values: if all compliant stocks look identical, the algorithm loses its ability to select the best options.

Next, we tested if a Logarithmic curve would perform better than a simple linear min–max line. To do this, we computed several **robustness indicators** to see how differently the two methods scored the exact same stocks. The results proved the two methods are virtually identical:

- **Spearman Rank Correlation (**0.9984**):** This metric checks if both methods rank the companies in the exact same order (from best to worst). A score of almost 1.0 means the ranking is identical. Since evolutionary algorithms care most about relative ranking, the logarithmic math adds no new value.
- **Pearson Correlation (**0.9991**) and R-squared (**0.9635**):** These measure how tightly the scores move together. A result this close to 1.0 confirms the two formulas produce the exact same trends.
- **Absolute Errors (Max:** 0.0269**, Mean:** 0.0186**):** This shows the actual numerical difference between a stock's linear score and its logarithmic score. An average difference of 0.018 is negligible.

Because the Truncated method blinds the algorithm and the Logarithmic method is mathematically redundant, we selected the linear min–max normalization.

Choice of weights and aggregation function
The final score is the equally weighted arithmetic mean of the four normalized components. Two remarks qualify this choice. First, the binary screen remains non-compensatory: an asset breaching any single threshold is excluded before scoring, and no margin on one ratio can offset a breach of another [55]. The averaging only operates inside the admissible region, where it measures the average safety margin of an asset that is already compliant. Second, the major index methodologies treat all quantitative screens as equally binding and do not prescribe any relative importance among them; in the absence of such guidance, equal weights constitute the neutral choice. The weights were not calibrated against the opinion of Sharia scholars, which was not available to us; since the weights are explicit parameters of the score, a Sharia supervisory board may substitute its own weighting without any change to the optimization model.

2.5. Estimation of return and risk

For each compliant stock, we compute return and risk based on the Markowitz [21] framework, where μ_i represents the expected annual return of stock i, obtained by annualizing the mean daily return over 252 trading days, and Σ denotes the annualized covariance matrix, which serves as the first risk estimator.

These quantities are obtained by converting daily price series into annualized statistical measures. Let $r_{i,t}$ denote the daily return of asset i on day $t = 1, \ldots, T$, with $T = 248$, and $\mathbf{r}_t = (r_{1,t}, \ldots, r_{n,t})^\top$ the vector of daily asset returns. Then

$$\mu_i = 252 \times \frac{1}{T} \sum_{t=1}^{T} r_{i,t}, \qquad \Sigma = 252 \times \mathrm{Cov}(\mathbf{r}_t),$$

where Cov denotes the sample covariance matrix of the daily returns.

As a second risk measure, we use the historical Conditional Value-at-Risk (CVaR) [28], defined as the average loss over the worst 5% of trading days. For a portfolio $\mathbf{w}$, the daily loss is the opposite of the daily portfolio return, $L_t(\mathbf{w}) = -\mathbf{w}^\top \mathbf{r}_t$, so that a negative return corresponds to a positive loss. The CVaR at the 95% confidence level, denoted CVaR 95%, is then

$$\mathrm{CVaR}_{95\%}(\mathbf{w}) = \frac{1}{\lceil 0.05\, T \rceil} \sum_{k=1}^{\lceil 0.05\, T \rceil} L_{(k)}(\mathbf{w}),$$

where $L_{(1)} \geq L_{(2)} \geq \cdots \geq L_{(T)}$ are the daily portfolio losses sorted in decreasing order, so that the sum runs over the $\lceil 0.05\, T \rceil = 13$ worst days.

The variance depends only on the second moment of the return distribution and penalizes gains and losses symmetrically, whereas the CVaR is a coherent risk measure [56, 57] that focuses on the left tail. When returns are normally distributed, minimizing either measure leads to the same efficient portfolios [28]; the two approaches differ when returns are asymmetric and heavy-tailed, in which case the mean–variance framework can substantially underestimate the tail risk captured by the mean–CVaR framework, as shown by Agarwal and Naik [58] for hedge funds with option-like payoffs.

2.6. Multi-Objective Portfolio Optimization problem definition

Let $\mathbf{w} \in \mathbb{R}^n$ denote the portfolio weight vector. The investor wishes to maximize the expected return $\mathbf{w}^\top \mu$, to minimize the risk of the portfolio and to maximize its compliance score $\mathbf{w}^\top S$. Two alternative measures of risk are considered: the variance, as in the classical mean–variance framework, and the CVaR, which focuses on extreme losses. Since NSGA-III, like most multi-objective solvers, is formulated for minimization, the quantities to be maximized are replaced by their negatives, which leaves the set of Pareto-optimal portfolios unchanged. The

objective functions are therefore defined as:

$$\begin{aligned} f_1(\mathbf{w}) &= -\mathbf{w}^\top \mu && \text{(negative expected return)} \\ f_2(\mathbf{w}) &= \mathbf{w}^\top \Sigma \mathbf{w} && \text{(portfolio variance)} \\ f_3(\mathbf{w}) &= -\mathbf{w}^\top S && \text{(negative compliance score)} \\ f_4(\mathbf{w}) &= \mathrm{CVaR}_{95\%}(\mathbf{w}) && \text{(downside risk)} \end{aligned} \tag{1}$$

where $\mu \in \mathbb{R}^n$ is the vector of expected asset returns μ_i and $S \in \mathbb{R}^n$ represents the vector of compliance scores of the assets. The two risk objectives f_2 and f_4 are never optimized simultaneously: each optimization problem contains a single risk objective, so that the effect of the compliance score can be assessed for each specification of the financial trade-off.

All portfolios satisfy the budget and no-short-selling constraints, together with two operational constraints: the portfolio may hold at most K assets, and every asset held must have a weight between a minimum $w_{\min}$, which avoids economically insignificant positions, and a maximum $w_{\max}$, which prevents excessive concentration:

$$\sum_{i=1}^{n} w_i = 1, \quad w_i \geq 0 \qquad \text{(budget, no short selling)} \tag{2}$$

$$\mathrm{card}(\mathbf{w}) \leq K \qquad \text{(maximum number of assets)} \tag{3}$$

$$w_{\min} \leq w_i \leq w_{\max} \ \text{ for every asset held } (w_i > 0) \qquad \text{(position bounds)} \tag{4}$$

where $\mathrm{card}(\mathbf{w})$ denotes the number of assets with a positive weight, $K = 20$, $w_{\min} = 2\%$ and $w_{\max} = 10\%$. An asset is therefore either excluded from the portfolio ($w_i = 0$) or held with a weight between 2% and 10%. The maximum weight corresponds to the 10% single-issuer limit of the European UCITS diversification rules and implies that every portfolio holds at least 10 assets, so that the admissible portfolios contain between 10 and 20 assets.

The versions of the optimization problem are summarized in Table 3. For each version, the benchmark representing current practice, referred to as *binary screening*, optimizes only the expected return and the risk objective over the same screened universe and under the same constraints, so that the compliance information is used only to filter the universe. Comparing each continuous-scoring model with its binary-screening counterpart isolates the contribution of the score.

Table 3. Versions of the optimization problem. Each version contains a single risk objective; its binary-screening benchmark drops the compliance objective f_3. Both versions are subject to constraints (2)–(4).

Version	Risk measure	Continuous scoring	Binary screening
P1	Variance	(f_1, f_2, f_3)	(f_1, f_2)
P2	$\mathrm{CVaR}_{95\%}$	(f_1, f_4, f_3)	(f_1, f_4)

The maximum-number-of-assets and minimum-position constraints (3)–(4) make the feasible set of both versions non-convex, since the subset of assets held must be selected among the 146 candidates. This is the setting in which a population-based evolutionary algorithm is most useful.

3. Multi-Objective Portfolio Optimization with NSGA-III

In multi-objective optimization, the aim is not to identify a single solution, but rather a set of *non-dominated solutions*, also known as the Pareto-optimal set. A solution $\mathbf{w}^{(1)}$ is said to dominate another solution $\mathbf{w}^{(2)}$ if $f_j(\mathbf{w}^{(1)}) \leq f_j(\mathbf{w}^{(2)})$ for all objectives j, with strict inequality for at least one objective. A solution is non-dominated if no other solution in the feasible region dominates it. The set of all such non-dominated solutions forms the Pareto front, which represents the efficient trade-offs among objectives [39].

The Non-dominated Sorting Genetic Algorithm III (NSGA-III) is particularly suitable for this problem because it was specifically designed to handle three or more conflicting objectives, where classical bi-objective evolutionary algorithms such as NSGA-II may be inadequate. By incorporating a reference-point-based approach, NSGA-III ensures both convergence to the Pareto front and diversity of solutions across high-dimensional objective spaces [32, 39].

The computational optimization phase of this study is executed using the pymoo multi-objective optimization framework in Python [59]. To determine the optimal NSGA-III parameters, an exhaustive grid search was conducted across discrete candidate values for each hyperparameter, as described in Section 4.2; the resulting configuration is summarized in Table 4.

Table 4. NSGA-III configuration used for all reported results.

Element	**Setting**
Decision variables	146 continuous genes in $[0, 1]$, mapped to admissible weights (see below)
Population size N	450
Reference directions	Das–Dennis, $p = 30$ partitions ($H = 496$ directions; 31 for two objectives)
Initialization	equally weighted portfolios of 10 assets (three anchors + random subsets)
Selection	binary tournament
Crossover	simulated binary crossover, probability 1.0, distribution index 30
Mutation	polynomial mutation, distribution index 20, probability $1/146$ per gene
Stopping criterion	1,000 generations
Random seeds	seed 1 for the reported fronts; seeds 1–5 for all comparative experiments

The core mechanics of the NSGA-III algorithm are detailed below. Note that while the step-by-step breakdown is explicitly formalized for the primary three-objective space ($J = 3$), the routine adapts identically to the two-objective binary-screening benchmark ($J = 2$) by dropping the compliance dimension.

1. **Population Initialization:** The search begins with an initial population $P_0 = \{\mathbf{x}^1, \ldots, \mathbf{x}^N\}$ of size $N = 450$, where each individual is a vector $\mathbf{x} \in [0, 1]^n$ and $n = 146$ is the number of Sharia-compliant assets. Since every admissible portfolio holds at least ten assets, the population is initialized with equally weighted portfolios of ten assets: three anchor portfolios, made of the ten assets with the highest expected return, the lowest individual risk and the highest compliance score, respectively, and 447 portfolios of ten randomly drawn assets.
2. **Constraint handling (repair):** Because the genetic operators generate vectors whose components do not sum to 100%, each vector is normalized inside the evaluation:

$$_i = \frac{x_i}{\sum_{j=1}^{n} x_j + \epsilon}$$

Where $\epsilon = 10^{-12}$ is introduced as an infinitesimal regularizer to avoid numerical division-by-zero errors.
To ensure the generated portfolios respect the maximum allowed number of assets (K), a sequential adjustment process is applied to the weights.

- First, the algorithm retains only the K largest weights and sets all others to zero.
- It then removes any weights that fall below the minimum threshold ($w_{\min}$). Because these removals reduce the total portfolio weight below 1, the remaining weights must be redistributed so they sum to 100% while strictly staying within the valid boundaries ($w_{\min} \leq w_i \leq w_{\max}$). Rather than simply multiplying the weights—which could push larger allocations over the maximum limit—the algorithm adds a constant value, or "shift," to each retained asset. The exact value of this shift is calculated using the bisection method, an iterative search that finds the precise mathematical adjustment needed for the weights to sum to 1 without breaking the $w_{\max}$ constraint.

- Finally, a safety fallback is implemented to guarantee that the allocation remains mathematically feasible. Because an upper bound constraint ($w_{\max}$) of 10% strictly requires a minimum of 10 assets to reach a total portfolio weight of 100%, the algorithm checks the remaining asset count. If the filtering process leaves fewer than 10 assets, it overrides the earlier steps and simply retains the 10 largest components of the original vector $\mathbf{x}$, ensuring the upper weight limit can still be satisfied.

3. **Parent and Offspring Populations combination:** At each generation t, the parent population P_t and its generated offspring Q_t are merged into a combined pool $R_t = P_t \cup Q_t$ of size $2N = 900$.
4. **Fast Non-Dominated Sorting:** The pool is sorted into hierarchical layers $(F_1, F_2, \dots)$ based on the Pareto dominance rule where the first front contains non-dominated individuals, the second front contains those dominated only by the first front, and so on. A portfolio $\mathbf{w}^{(1)}$ strictly dominates $\mathbf{w}^{(2)}$ ($\mathbf{w}^{(1)} \succ \mathbf{w}^{(2)}$) if and only if:
$$\forall j \in \{1,2,3\},\ f_j(\mathbf{w}^{(1)}) \leq f_j(\mathbf{w}^{(2)}) \quad \& \quad \exists j \in \{1,2,3\},\ f_j(\mathbf{w}^{(1)}) < f_j(\mathbf{w}^{(2)})$$
5. **Selection of the Next Generation:** A new population P_{t+1} is constructed by sequentially transferring individual vectors from consecutive fronts $(F_1, F_2, \dots)$ until the addition of a marginal front F_l causes the population size to exceed $N = 450$ ($|P_{t+1}| > N$)
6. **Niching Strategy for the Last Front:** To select individuals from F_l, NSGA-III employs a specialized reference-point mechanism:

   - Normalize Objectives: The algorithm first normalize the objective values of all selected individuals to account for differing scales so that no objective can dominate the distance calculations. therefore each raw objective vector $f(\mathbf{w}) = \big(f_1(\mathbf{w}), f_2(\mathbf{w}), f_3(\mathbf{w})\big)$ is mapped to its normalized equivalent
   $$f^{norm} = \big(f_1^{norm}(\mathbf{w}), f_2^{norm}(\mathbf{w}), f_3^{norm}(\mathbf{w})\big)$$
   Following [39], the ideal point, i.e. the best value of each objective in the population, is first subtracted from each objective before scaling. The normalized objectives are therefore non-negative whatever the sign of the original objectives, so that the negated return and compliance score (f_1 and f_3, which are negative) are handled exactly like the risk objective and do not affect the association with the reference directions.
   - Das–Dennis Reference Grid Coordinate Generation: To maintain a uniform spread across the 3D surface, a fixed set of reference directions Y is generated. For $J = 3$ objectives and $p = 30$ coordinate partitions, the combinatorial number of reference lines (H) is defined precisely via the Das–Dennis formulation:
   $$H = \binom{J+p-1}{p} = \binom{3+30-1}{30} = \binom{32}{30} = 496 \text{ reference lines}$$
   The population ($N = 450$) is slightly smaller than the number of reference directions, so that a few directions remain without an associated solution; the parameter study of Section 4.2 shows that this does not degrade the approximation.
   - Perpendicular Association Distance Metric : Each generated portfolio individual is mapped to the reference grid by calculating its orthogonal (perpendicular) distance $d(\mathbf{w}, y)$ to each tracking line $y \in Y$:
   $$d(\mathbf{w}, y) = \left\| f^{norm}(\mathbf{w}) - \frac{y^T f^{norm}(\mathbf{w})}{\|y\|^2} y \right\|_2$$
   A portfolio is then assigned to the specific reference line y^* that minimizes this distance field: $y^* = \arg\min_{y \in Y} d(\mathbf{w}, y)$.
   - Each individual is thus assigned to the reference point with the shortest perpendicular distance.
   - Niche Count Computation: The algorithm computes how many existing members of the population are associated with each reference point

- Population Selection via Niche: The algorithm Iteratively choose members from F_l that are associated with the least-represented reference points, ensuring the final population remains diverse and evenly distributed.

7. **Variation operators**: NSGA-III applies selection, crossover (mixing portfolio weights) and mutation (small perturbations of the weights) to create offspring that are mapped to admissible portfolios by the constraint handling repair of step 2.
8. **Iteration and termination**:Repeat steps until the stopping criterion (maximum generations $= 1000$) is met, and obtain a well-distributed approximation to the Pareto set .

4. Parameter selection and Algorithmic validation

4.1. Choice of the test problem

The parameter study is carried out on the continuous-scoring model of P2, and the selected configuration is used for all models. P2 is the most demanding of the models considered: it has the same decision variables, constraints and number of objectives as P1, but its risk objective, the CVaR, depends only on the worst days of the sample and is therefore non-smooth, whereas the variance of P1 is a smooth quadratic function of the weights. A configuration able to approximate the front of P2 is therefore expected to be at least as effective on P1, and repeating the full factorial design for each model would multiply the computational cost without changing the selected configuration.

4.2. Parameter selection

To find the best parameter settings for the NSGA-III algorithm, we systematically tested a predefined list of values, a process known as grid search. Because this specific problem is highly complex, we do not know what the true Pareto front is. Instead, we created an empirical reference front by combining all the solutions from every test we ran and keeping only the very best, non-dominated ones.We then scored our results against this reference front. After standardizing the values on a common scale (using the best and worst points of the front), we used two main metrics to judge performance:

- Hypervolume ratio: This calculates how much of the solution space our results cover compared to the reference front [59, 60]. A higher score is better because it proves the solutions are both close to optimal (convergence) and well-spread out (coverage). We used a standard reference point of $(1.1, 1.1, 1.1)$ for this calculation.
- IGD$^+$ Indicator: The Inverted Generational Distance Plus measures the average distance from the reference front to our solutions [61]. A lower score is better, as it indicates our approximations accurately reached the reference solutions.

Parameter study. An exhaustive grid search was conducted to evaluate all discrete parameter combinations on the constrained problem P2. The search space consisted of the number of Das–Dennis partitions $p \in \{15, 20, 23, 30\}$, population size $N \in \{300, 450\}$, and number of generations $G \in \{500, 1000\}$, and crossover distribution index $\eta_c \in \{15, 30\}$, with five random seeds applied per configuration (32 total configurations, yielding 160 runs).

The number of reference directions is by far the most influential parameter: it explains 74.6% of the variation of the hypervolume ratio, which increases from 0.9312 for $p = 15$ (136 directions) to 0.9592 for $p = 30$ (496 directions), because a finer grid retains more non-dominated portfolios . A larger population brings a smaller but significant gain , whereas doubling the number of generations has a marginal effect ,which indicates that the algorithm has essentially converged after 500 generations.

The configuration with the highest mean hypervolume ratio, $p = 30$, $N = 450$, $1{,}000$ generations and a crossover distribution index of 30, is retained for all reported results.Since the crossover distribution index has a negligible effect (0.2% of the variance), its initial value $_c = 30 is kept$.

Table 5. Sensitivity of NSGA-III to its parameters: mean hypervolume ratio for each level tested and share of the variance explained by each parameter (factorial ANOVA, five seeds per configuration).

Parameter	Level	Mean HV ratio	Share of variance (%)
Number of reference-direction partitions	15	0.9312	74.6
	20	0.9489	
	23	0.9516	
	30	0.9592	
Population size	300	0.9446	6.8
	450	0.9508	
Number of generations	500	0.9461	1.8
	1000	0.9493	
Crossover distribution index	15	0.9483	0.2
	30	0.9472	
Residual	–	–	16.6

4.3. Algorithmic validation

To assess whether the results depend on the choice of algorithm, NSGA-II [38] was run on the same problem P2 with the same population, operators, initialization and number of generations, and with the same five random seeds; only the survival mechanism differs (crowding distance instead of reference directions). Table 6 reports the mean and standard deviation over the five seeds of both algorithms, evaluated against their common reference front.

Table 6. Comparison of NSGA-III and NSGA-II under identical parameters, over five seeds. SD: standard deviation; ↑: higher is better, ↓: lower is better.

	Hypervolume ratio ↑		IGD^+ ↓			
Algorithm	Mean	SD	Mean	SD	Number of solutions	Runtime (s)
NSGA-II	0.9598	0.0024	0.0149	0.0007	450.0	97.1
NSGA-III	**0.9742**	0.0034	**0.0112**	0.0008	266.2	144.6

NSGA-III outperforms NSGA-II on both indicators: its hypervolume ratio is higher (0.9742 against 0.9598) and its IGD^+ lower (0.0112 against 0.0149), with comparable seed-to-seed variability, at the cost of a longer run time. NSGA-II returns more solutions, since it retains its whole non-dominated population, but these solutions cover the reference front less completely, which reflects the known weakness of the crowding distance beyond two objectives. Together with the stability of the indicators across seeds, this comparison indicates that the fronts analysed below are not an artefact of a particular algorithm. Run times are indicative, as several experiments were executed concurrently.

5. Results and Discussion

This section compares the portfolios obtained with binary screening and with continuous scoring, first through an overview of the four optimized models, then through the shape of their Pareto fronts, a comparison at equal risk and a set of representative investor profiles.

5.1. Overview of the binary-screening and continuous-scoring models in P1 and P2

Before comparing the models, we verify that the compliance score is not redundant with the financial objectives: if compliant firms were systematically less risky or more profitable, maximizing the score would amount to minimizing risk or maximizing return, and the three-objective problem would effectively reduce to a two-objective one. Table 7 reports the correlations between the expected return, the two risk measures and the compliance score of the 146 assets of the screened universe. The compliance score is unrelated to the expected return ($r = 0.02$) and only weakly associated with lower risk ($r = -0.24$ with the variance and $r = -0.28$ with the CVaR), whereas the two risk measures are strongly correlated with each other ($r = 0.93$). The score therefore carries information that is distinct from risk and return, which justifies treating it as a separate objective.

Table 7. Pearson correlations between the characteristics of the 146 assets of the screened universe.

	Expected return	Variance	CVaR 95%	Compliance score
Expected return	1			
Variance	0.04	1		
CVaR 95%	−0.14	0.93	1	
Compliance score	0.02	−0.24	−0.28	1

Table 8 summarizes the four models of Table 3: the binary-screening and continuous-scoring models of version P1, in which risk is measured by the variance, and of version P2, in which it is measured by the CVaR. For each model, the table reports medians computed over all the portfolios of the front, and each portfolio is evaluated with both risk measures, so that the four models can be compared on the same basis. For readability, the risk of P1 is reported as the annualized volatility, i.e. the square root of the variance.

The table also reports an indicator of the exposure of each portfolio to future losses of compliance. A firm i is called *near-boundary* if at least one of its financial ratios exceeds 80% of the corresponding threshold, i.e. if $R_{ij} > 0.8\,\tau_j$ for at least one ratio j, where R_{ij} denotes the value of ratio j for firm i and τ_j its threshold. Denoting by $\mathcal{N}$ the set of near-boundary firms, the near-boundary weight of a portfolio is the share of its capital invested in these firms:

$$NB(\mathbf{w}) = \sum_{i \in \mathcal{N}} w_i. \tag{5}$$

With thresholds of 33% for debt, cash and receivables and 5% for interest income, a firm is near-boundary when one of the first three ratios exceeds 26.4% of total assets or when interest income exceeds 4% of revenue. Such firms are compliant today, but a moderate deterioration of their balance sheet could make them non-compliant and force the manager to sell them. They represent 62 of the 146 assets of the screened universe, which shows that a large part of the universe remains fragile after binary screening.

Table 8. Summary of the four models (medians over the portfolios of each front). Every portfolio is evaluated with both risk measures; near-boundary weight: share of capital invested in near-boundary firms, Eq. (5).

Version	Model	Portfolios	Score range	Median score	Median return	Median volatility	Median CVaR 95%	Median near-boundary weight	Assets per portfolio
P1	Binary screening	31	[0.577, 0.669]	0.642	0.491	0.171	0.024	0.443	10–14
	Continuous scoring	291	[0.587, 0.799]	0.732	0.350	0.162	0.022	0.262	10–15
P2	Binary screening	31	[0.577, 0.671]	0.648	0.492	0.185	0.023	0.480	10–13
	Continuous scoring	255	[0.594, 0.799]	0.731	0.369	0.170	0.022	0.200	10–16

In both versions, the median compliance score is higher by 0.08 to 0.09 (0.732 against 0.642 in P1, 0.731 against 0.648 in P2). This gain translates directly into a lower exposure to near-boundary firms: the median near-boundary weight falls from 44.3% to 26.2% in P1 and from 48.0% to 20.0% in P2. Because binary screening ignores the distance to the thresholds once an asset is admitted, the binary-screening portfolios place almost half of their capital in firms that are close to losing their compliance. The gain is not obtained by taking more risk: the median

volatility of the continuous-scoring portfolios is 16.2% against 17.1% in P1 and 17.0% against 18.5% in P2, and their median daily CVaR 95% is 2.2% against 2.4% and 2.3%, respectively. It comes, however, at the cost of a lower median expected return, by 14.1 percentage points in P1 (35.0% against 49.1%) and 12.3 points in P2 (36.9% against 49.2%). All portfolios hold between 10 and 16 assets.

The number of portfolios per front is bounded by the number of reference directions and the medians depend on the distribution of the portfolios along each front; they are reported for descriptive purposes, the comparison at equal risk being given in Table 10.

5.2. Pareto fronts of the two versions

Figures 1 and 2 show, for P1 and P2 respectively, the front of the binary-screening model and that of the continuous-scoring model, both obtained under identical constraints. The two versions exhibit the same structure: whether risk is measured by the variance or by the CVaR, the continuous-scoring fronts reach almost identical median compliance scores (0.732 and 0.731), so that the contribution of the score does not depend on the risk measure, which mainly affects the composition of the portfolios.

In both versions, the binary-screening model produces a curve of 31 portfolios, on which a higher expected return can only be obtained at the cost of a higher risk, while the compliance score is not controlled and remains in the lower part of its range. The continuous-scoring model, by contrast, produces a surface in the return–risk–compliance space, which extends towards substantially higher compliance scores, up to 0.799. The upper edge of this surface lies close to the binary-screening curve: the third objective does not create new return–risk opportunities, but makes explicit how much expected return must be forgone, at a given level of risk, to move away from the compliance thresholds.

These differences are statistically significant (Table 9): two-sided Wilcoxon rank-sum and Kolmogorov–Smirnov tests reject the equality of the distributions of the compliance score and of the expected return between the two fronts (binary vs continuous) in both versions (P1, P2), whereas they do not reject the equality of the risk distributions. The tests compare the sets of portfolios returned by the algorithm, whose size depends on the number of reference directions (Section 3); they are therefore interpreted as evidence of a systematic shift between the two fronts, which the comparison at equal risk of Section 5.3 confirms. The third objective therefore moves the portfolios towards higher compliance rather than towards higher risk.

Table 9. Tests of equality of the distributions of the binary-screening and continuous-scoring portfolios. Risk is the variance in P1 and the CVaR 95% in P2.

Version	Objective	Wilcoxon p-value	Kolmogorov–Smirnov p-value
P1	Compliance score	2.3×10^{-16}	1.3×10^{-25}
	Expected return	1.4×10^{-8}	7.4×10^{-7}
	Risk	0.20	0.70
P2	Compliance score	3.0×10^{-17}	3.9×10^{-26}
	Expected return	7.6×10^{-9}	3.4×10^{-7}
	Risk	0.57	0.90

5.3. Binary screening versus continuous scoring at equal risk

To isolate the contribution of the continuous score, Table 10 compares the two models at the same level of risk. Three risk levels are considered, low, medium and high, defined as the 25th, 50th and 75th percentiles of the risk of the binary-screening front; risk is measured by the annualized volatility in P1 and by the daily CVaR 95% in P2. For each level, the table first reports the expected return and the compliance score of the binary-screening portfolio whose risk is closest to this level, i.e. the portfolio that current practice would select. It then considers the continuous-scoring portfolios whose risk lies within $\pm 10\%$ of the same level, and reports their number (n), their median and maximum compliance scores, and their median expected return. The median score describes a

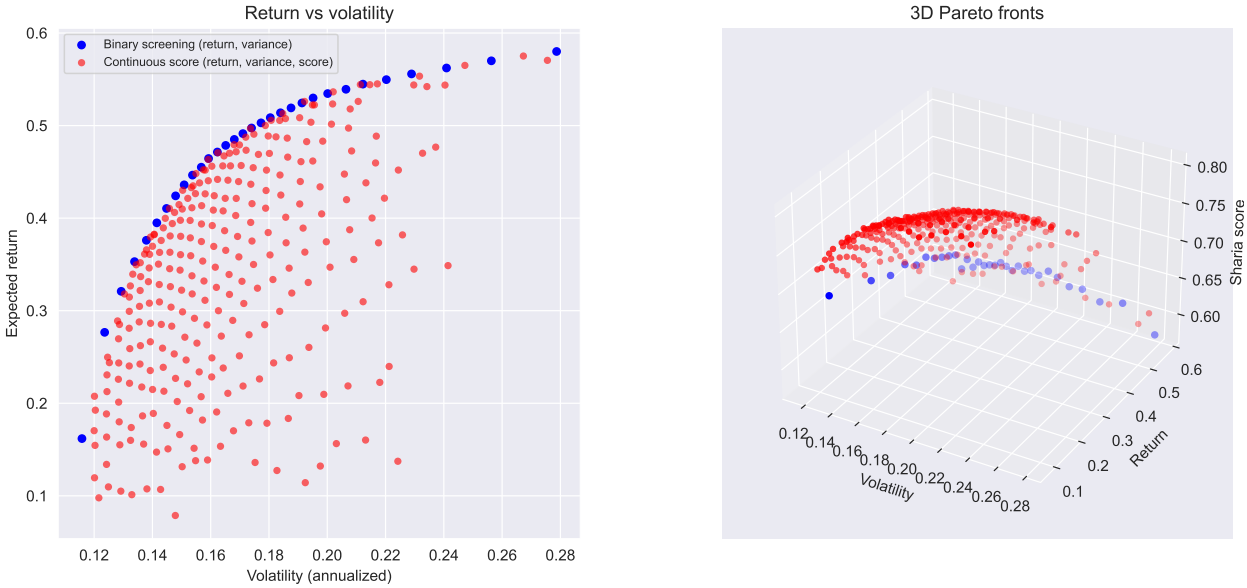


Figure 1. Version P1 (variance): Pareto fronts of the binary-screening model (return and variance, blue) and of the continuous-scoring model (return, variance and compliance score, red) under identical constraints. Left: projection on the return–volatility plane, where volatility is the square root of the variance; right: three-dimensional view including the compliance score.

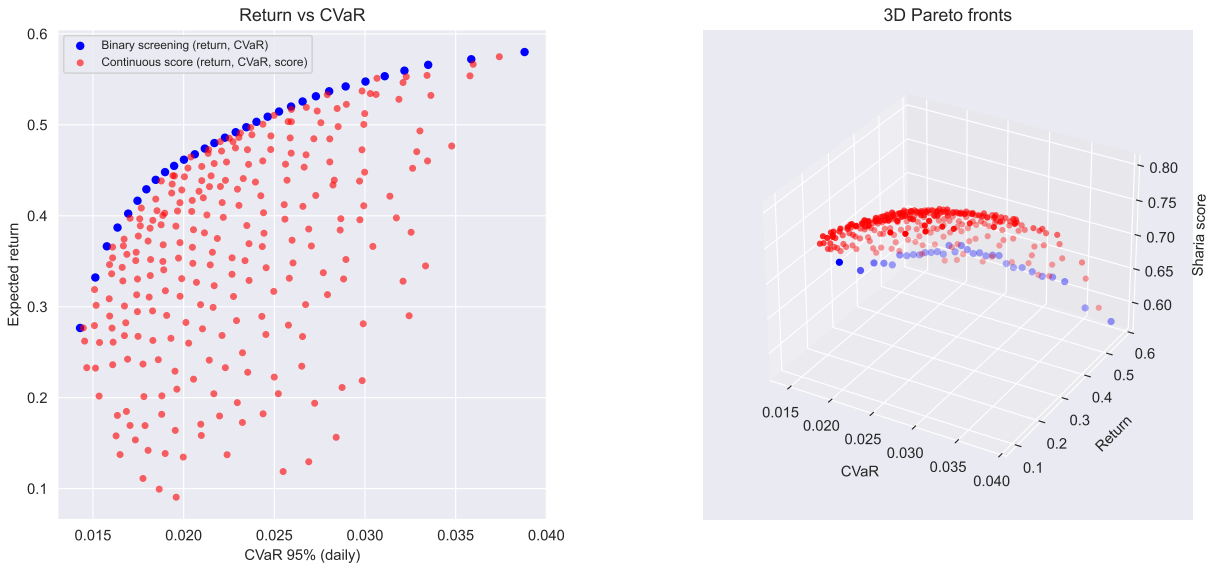


Figure 2. Version P2 (CVaR): Pareto fronts of the binary-screening model (return and CVaR, blue) and of the continuous-scoring model (return, CVaR and compliance score, red) under identical constraints. Left: projection on the return–CVaR plane; right: three-dimensional view including the compliance score.

typical portfolio of the continuous model at this level of risk, the maximum score indicates the highest compliance attainable without increasing risk, and the median return measures the corresponding cost in expected return.

Table 10. Equal-risk comparison of binary screening and continuous scoring under identical constraints.

			Binary screening		Continuous scoring (same risk band)			
Version	Risk level	Risk	Return	Score	n	Median score	Max score	Median return
P1	Low	0.149	0.424	0.660	119	0.733	0.791	0.331
	Medium	0.171	0.491	0.642	111	0.739	0.797	0.386
	High	0.198	0.535	0.630	76	0.740	0.799	0.427
P2	Low	0.019	0.440	0.648	73	0.732	0.782	0.350
	Medium	0.023	0.492	0.661	74	0.739	0.792	0.394
	High	0.028	0.531	0.637	60	0.740	0.799	0.417

Three conclusions emerge from Table 10, and they hold in both versions. First, at every level of risk, the continuous-scoring model offers portfolios that are markedly more compliant than the binary-screening portfolio, with a gain in median score of about 0.07 to 0.11. Second, a high level of compliance does not require accepting more risk: scores close to the maximum remain attainable at every risk level. The gain is even largest for the riskiest portfolios, because binary screening tends to select less compliant firms as risk increases, whereas the continuous model keeps the compliance level stable. Third, this gain has a cost in expected return, of about 10 percentage points at equal risk. Binary screening therefore represents only one extreme of the trade-off, the highest return with an uncontrolled level of compliance, whereas the continuous model allows the investor to choose the balance between return and compliance at each level of risk.

5.4. Binary screening versus continuous scoring for similar investor profiles

Since an investor holds a single portfolio, representative portfolios are selected on each front according to four investor profiles: return-priority, risk-priority, compliance-priority and balanced. Each profile is expressed as a vector of preference weights $(\lambda_1, \lambda_2, \lambda_3)$ on the expected return, the risk and the compliance score (Table 11). To choose the portfolio of a profile, each objective is first rescaled on the front between 0 (worst value) and 1 (best value), so that the three objectives are comparable, and the portfolio with the highest weighted sum of the rescaled objectives is selected:

$$\mathbf{w}^* = \arg\max_{\mathbf{w}\in\mathcal{F}} \sum_j \lambda_j \, \tilde{f}_j(\mathbf{w}), \tag{6}$$

where $\mathcal{F}$ is the set of portfolios of the front, λ_j the weight of objective j and $\tilde{f}_j(\mathbf{w})$ its rescaled value. The binary-screening models have no compliance objective: their profiles keep the same relative importance of return and risk, and they have no compliance-priority profile.

Table 11. Preference weights of the investor profiles (return, risk, compliance score).

Profile	Continuous scoring	Binary screening
Return-priority	$(0.6,\ 0.2,\ 0.2)$	$(0.75,\ 0.25)$
Risk-priority	$(0.2,\ 0.6,\ 0.2)$	$(0.25,\ 0.75)$
Compliance-priority	$(0.2,\ 0.2,\ 0.6)$	–
Balanced	$(1/3,\ 1/3,\ 1/3)$	$(0.5,\ 0.5)$

Table 12 compares the selected portfolios with two benchmarks built on the same screened universe: the equally weighted portfolio [62] and the capitalization-weighted portfolio, which mimics the construction of Sharia indices such as the S&P 500 Shariah. For each portfolio, the table reports the annualized expected return and volatility, the CVaR 95% (average loss over the 5% worst days, computed on daily returns), the compliance score, the near-boundary weight defined in (5) and the number of assets held. The risk-adjusted performance of each portfolio is

measured by the Sharpe ratio [63], i.e. its excess return per unit of total risk:

$$SR(\mathbf{w}) = \frac{\mu_p(\mathbf{w}) - r_f}{\sigma_p(\mathbf{w})}, \tag{7}$$

where μ_p and σ_p are the annualized mean return and volatility of the portfolio, computed from its daily returns, and $r_f = 2\%$ is the annual risk-free rate. Since the Sharpe ratio is not optimized by the models, it provides an independent assessment of the financial quality of the selected portfolios.

Table 12. In-sample indicators of the portfolios selected for each investor profile (Eq. (6)) and of two benchmarks (23 July 2024 – 18 July 2025). CVaR 95%: average loss over the 5% worst days, computed on daily returns.

Version	Model	Profile	Return	Volatility	CVaR 95%	Sharpe	Score	Near-boundary weight	Assets
P1	Continuous	Return-priority	0.508	0.191	0.026	2.56	0.690	0.400	10
		Risk-priority	0.331	0.139	0.018	2.25	0.721	0.349	11
		Compliance-priority	0.170	0.168	0.022	0.90	0.791	0.000	11
		Balanced	0.408	0.164	0.022	2.36	0.731	0.298	11
	Binary	Return-priority	0.535	0.200	0.029	2.57	0.630	0.476	11
		Risk-priority	0.376	0.138	0.017	2.58	0.658	0.469	14
		Balanced	0.485	0.168	0.023	2.77	0.642	0.436	12
P2	Continuous	Return-priority	0.503	0.190	0.026	2.54	0.696	0.400	10
		Risk-priority	0.262	0.139	0.015	1.74	0.697	0.178	11
		Compliance-priority	0.173	0.176	0.023	0.87	0.792	0.000	11
		Balanced	0.390	0.152	0.019	2.43	0.720	0.400	11
	Binary	Return-priority	0.566	0.248	0.033	2.20	0.622	0.486	11
		Risk-priority	0.366	0.139	0.016	2.50	0.648	0.501	13
		Balanced	0.462	0.164	0.020	2.69	0.650	0.497	11
Benchmarks		Equally weighted $(1/N)$	0.050	0.206	0.030	0.15	0.634	0.425	146
		Capitalization-weighted	0.139	0.226	0.033	0.53	0.609	0.389	37

For the balanced profile, the continuous-scoring portfolio raises the compliance score from 0.642 to 0.731 in P1 and from 0.650 to 0.720 in P2, and reduces the near-boundary weight from 43.6% to 29.8% and from 49.7% to 40.0%, respectively. Moving from the balanced binary-screening portfolio to the balanced continuous-scoring portfolio changes the weight of 15 assets in P1 and 12 in P2, which shows that the score changes the composition of the portfolios substantially rather than marginally.

The compliance-priority portfolios go further: with a score of 0.79 in both versions, they hold no near-boundary firm at all, but their expected return falls to about 17%. This gain in compliance has a financial cost: the binary-screening profiles achieve Sharpe ratios between 2.20 and 2.77, against 0.87 to 2.56 for the continuous-scoring profiles, which measures the price paid for a larger margin of compliance.

All optimized portfolios, however, outperform the passive benchmarks, whose Sharpe ratios do not exceed 0.53 over the sample.

Sector composition. In both versions, the binary-screening portfolios are mainly invested in Technology, Industrials and Healthcare, whereas the continuous-scoring portfolios shift capital towards Basic Materials, Consumer Defensive and Energy.

6. Conclusion

This paper embeds a continuous, threshold-relative Sharia compliance score as an objective of a multi-objective portfolio model solved with NSGA-III, and compares it with the practice of binary screening followed by return–risk optimization. Both approaches are applied to the same screened universe, under the same realistic constraints on the number and size of positions, and with two alternative risk measures, the variance (P1) and the CVaR (P2).

Five findings emerge. First, the evolutionary approach is reliable: the quality of the approximation depends mainly on the number of reference directions, and NSGA-III outperforms NSGA-II .

Second, the compliance score is not redundant with the financial objectives, since it is almost uncorrelated with the expected return and only weakly correlated with risk at the asset level.

Third, optimizing the score produces markedly more compliant portfolios without additional risk: at equal risk, the continuous model offers scores higher by 0.07 to 0.11, and the share of capital invested in firms close to the compliance thresholds is roughly halved. This gain has an explicit cost, of about 10 to 14 percentage points of expected return, which binary screening leaves invisible.

Fourth, these conclusions are identical whether risk is measured by the variance or by the CVaR, so that the contribution of the score does not depend on the risk measure.

Fifth, for every investor profile, the continuous model changes the composition of the portfolio substantially, and all optimized portfolios outperform the equally weighted and capitalization-weighted benchmarks in terms of Sharpe ratio, although the binary-screening portfolios achieve higher risk-adjusted performance in sample. The continuous score therefore does not replace binary screening, which remains necessary to define the admissible universe, but complements it by turning the distance to the thresholds into an explicit dimension of the investment decision.

The study has several limitations. (i) The empirical analysis is confined to a single market, the S&P 500, and to one year of daily prices, because comparable data on Muslim-majority markets were not available to the authors. (ii) The score relies on publicly available accounting ratios and on an equally weighted aggregation; it does not capture qualitative aspects of compliance, such as governance or business ethics, and its weights were not calibrated against the judgement of Sharia scholars. (iii) Transaction costs, turnover and rebalancing frequency are not included in the optimization model. (iv) Expected returns and risk are estimated from historical data, and optimized portfolios are known to be sensitive to estimation errors in these inputs, in particular in expected returns.

These limitations suggest several directions for future research. The framework should first be applied to other universes, in particular to Muslim-majority markets. It should also be tested with other Sharia standards, such as the AAOIFI standard based on market capitalization, the MSCI Islamic, DJIM or S&P methodologies (Table 1), or the benchmarks of the Securities Commission Malaysia; since the score depends on the standard only through the financial ratios, their denominators and their thresholds, this only requires re-specifying these three elements, and would show how the choice of standard affects the composition of the optimized portfolios. Further work could evaluate the optimized portfolios out of sample over a longer history, include transaction costs and turnover constraints directly in the model, and add higher-order moments, such as skewness, as additional objectives, a setting in which reference-direction algorithms such as NSGA-III are particularly appropriate.

REFERENCES

[1] Ulrich Derigs and Shehab Marzban. Review and analysis of current shariah-compliant equity screening practices. *International Journal of Islamic and Middle Eastern Finance and Management*, 1(4):285–303, 2008.

[2] Asim Hameed and Aishath Muneeza. Enhancing shariah stock screening criteria: A proposed framework for evaluating shariah compliance degree. *International Journal of Management and Applied Research*, 11: 126–139, 12 2024. doi: 10.18646/2056.112.24-006.

[3] Abdulrahman Qadi, Akash Sharma, and Francesca Medda. The price of permission: Classification uncertainty in constrained capital markets. *arXiv preprint arXiv:2608.12634*, 2026.

[4] Karim Ginena. Sharī ‘ah risk and corporate governance of islamic banks. *Corporate Governance*, 14(1): 86–103, 2014.

[5] Abdullah Alnamlah, M Kabir Hassan, Asem Alhomaidi, and Edib Smolo. A new model for screening shariah-compliant firms. *Borsa Istanbul Review*, 22:S10–S23, 2022.

[6] Abdulrahman Qadi, Akash Sharma, and Francesca Medda. From binary screens to continuous compliance: A shariah screening measure for portfolio design. *arXiv preprint arXiv:2512.22858*, 2025.

[7] Idris Babatunde Adeyemi and Ömer Faruk Tekdoğan. Understanding the screening criteria for shariah-compliant stocks. *Adam Academy Journal of Social Sciences*, 14(2):371–395, 2024.

[8] Sheila Nu Nu Htay, Zafar Abedeen, and Syed Ahmed Salman. Towards standardization of shari'ah screening norms and practices. *International Journal of Humanities and Social Science Invention*, 2(11):23–30, 2013.

[9] F. Gamaleldin. *Shariah-compliant stocks screening and purification*. PhD thesis, University of Liverpool, Liverpool, 2015.

[10] Mohammed Husain Khatkhatay and Shariq Nisar. Shariah compliant equity investments: An assessment of current screening norms. *Islamic Economic Studies*, 15(1), 2007.

[11] Ulrich Derigs and Shehab Marzban. New strategies and a new paradigm for shariah-compliant portfolio optimization. *Journal of Banking & Finance*, 33(6):1166–1176, 2009.

[12] Ashraf Md. Hashim, Farrukh Habib, Ziyaat Isaacs, and Mohamed Anouar Gadhoum. Isra-bloomberg sharī'ah stock screening and income cleansing methodologies: a conceptual paper. *ISRA International Journal of Islamic Finance*, 9(1):27–42, 2017.

[13] Engku Ngah S Engkuchik. A review of securities commission of malaysia's revised shariah stock-screening criteria. *Journal of Islamic Banking and Finance*, 4(1):71–80, 2016.

[14] Catherine SF Ho. Comparison of quantitative shari'ah-compliant screening methods. *ISRA International Journal of Islamic Finance*, 3(2):91–110, 2011.

[15] Dila Puspita, Adam Kolkiewicz, and Ken Seng Tan. Markowitz-based shariah compliant portfolio model with stochastic purification and probabilistic compliance screening constraints. *Journal of Islamic Accounting and Business Research*, 14(8):1300–1323, 2023.

[16] S&P Dow Jones Indices. Dow jones islamic market indices methodology. Technical report, S&P Dow Jones Indices, 2023. https://www.spglobal.com/spdji/en/documents/methodologies/methodology-dj-islamic-market-indices.pdf.

[17] S&P Dow Jones Indices. S&p shariah indices compliance criteria update & september 2023 rebalance. Technical report, S&P Dow Jones Indices, 2023. https://www.spglobal.com/spdji/en/documents/indexnews/announcements/20230804-1465479/1465479_spshariahindices-8-4-2023.pdf.

[18] FTSE Russell. Ftse yasaar global equity shariah index series — ground rules. Technical report, FTSE Russell, 2025. https://www.lseg.com/content/dam/ftse-russell/en_us/documents/ground-rules/ftse-yasaar-global-equity-shariah-index-series-ground-rules.pdf.

[19] MSCI Inc. Msci islamic index series methodology. Technical report, MSCI, 2025. https://www.msci.com/index/methodology/latest/Islamic.

[20] Eka Kristya Rahayu, Farikhin, and Moch. Fandi Ansori. Portfolio selection on the Jakarta Islamic Index using the UTASTAR method and compromise programming. *Statistics, Optimization & Information Computing*, 16(5):4629–4666, 2026. doi: 10.19139/soic-2310-5070-3724.

[21] Harry Markowitz. Portfolio selection. *Journal of Finance*, 7(11):77–91, 1952.

[22] Mohammed Ziane, Chillali Sara, Belhabib Fatima, Chillali Abdelhakim, and Karim El Moutaouakil. Portfolio selection problem: main knowledge and models (a systematic review). *Statistics, Optimization & Information Computing*, 12(3):799–816, 2024. doi: 10.19139/soic-2310-5070-1961.

[23] Alex Kane. Skewness preference and portfolio choice. *Journal of Financial and Quantitative Analysis*, 17 (1):15–25, 1982.

[24] Kin Keung Lai, Lean Yu, and Shouyang Wang. Mean-variance-skewness-kurtosis-based portfolio optimization. In *First International Multi-Symposiums on Computer and Computational Sciences (IMSCCS'06)*, volume 2, pages 292–297. IEEE, 2006.

[25] Ilyes Abid, Christian Urom, Jonathan Peillex, Majdi Karmani, and Gideon Ndubuisi. Pgp for portfolio optimization: application to esg index family. *Annals of Operations Research*, 347(1):405–417, 2025.

[26] Alexei A Gaivoronski and Georg Pflug. Value-at-risk in portfolio optimization: properties and computational approach. *Journal of risk*, 7(2):1–31, 2005.

[27] Rachel Campbell, Ronald Huisman, and Kees Koedijk. Optimal portfolio selection in a value-at-risk framework. *Journal of Banking & Finance*, 25(9):1789–1804, 2001.

[28] R Tyrrell Rockafellar, Stanislav Uryasev, et al. Optimization of conditional value-at-risk. *Journal of risk*, 2: 21–42, 2000.

[29] Ross A Maller, Robert B Durand, and Hediah Jafarpour. Optimal portfolio choice using the maximum sharpe ratio. *The journal of Risk*, 12(4):49, 2010.

[30] Fabio Caccioli, Imre Kondor, Matteo Marsili, and Susanne Still. Liquidity risk and instabilities in portfolio optimization. *International Journal of Theoretical and Applied Finance*, 19(05):1650035, 2016.

[31] Alireza Azarberahman, Malihe Tohidinia, and Hossein Aliakbarzadeh. Multi-objective portfolio optimization strategy using the spea-ii algorithm. *Journal of Information and Communication Technology*, 24(1):52–78, 2025.

[32] John Weirstrass Muteba Mwamba, Leon Mishindo Mbucici, and Jules Clement Mba. Multi-objective portfolio optimization: An application of the non-dominated sorting genetic algorithm iii. *International journal of financial studies*, 13(1):15, 2025.

[33] Ahti Salo, Michalis Doumpos, Juuso Liesiö, and Constantin Zopounidis. Fifty years of portfolio optimization. *European Journal of Operational Research*, 318(1):1–18, 2024.

[34] Prisadarng Skolpadungket, Keshav Dahal, and Napat Harnpornchai. Portfolio optimization using multi-obj ective genetic algorithms. In *2007 IEEE congress on evolutionary computation*, pages 516–523. IEEE, 2007.

[35] Mahmoud Awad, Mohamed Abouhawwash, and HN Agiza. On nsga-ii and nsga-iii in portfolio management. *Intelligent Automation & Soft Computing*, 32(3), 2022.

[36] Khalid Belabbes, Mostafa El Hachloufi, Zine El Abidine Guennoun, and Abderrahim El Attar. Uncertain portfolio optimization problems: Systematic review. *Statistics, Optimization & Information Computing*, 13: 961–976, 2025. doi: 10.19139/soic-2310-5070-2179.

[37] T.-J. Chang, N. Meade, J. E. Beasley, and Y. M. Sharaiha. Heuristics for cardinality constrained portfolio optimisation. *Computers & Operations Research*, 27(13):1271–1302, 2000.

[38] Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and TAMT Meyarivan. A fast and elitist multiobjective genetic algorithm: Nsga-ii. *IEEE transactions on evolutionary computation*, 6(2):182–197, 2002.

[39] Kalyanmoy Deb and Himanshu Jain. An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, part i: solving problems with box constraints. *IEEE transactions on evolutionary computation*, 2014.

[40] Simon Wietheger and Benjamin Doerr. A mathematical runtime analysis of the non-dominated sorting genetic algorithm III (NSGA-III). In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-23*, pages 5657–5665, 2023. doi: 10.24963/ijcai.2023/628.

[41] Yihe Qian and Jinpeng Wang. Multi-period portfolio optimization: A parallel NSGA-III algorithm with real-world constraints. *Finance Research Letters*, 60:104868, 2024. doi: 10.1016/j.frl.2023.104868.

[42] Lasse Heje Pedersen, Shaun Fitzgibbons, and Lukasz Pomorski. Responsible investing: The ESG-efficient frontier. *Journal of Financial Economics*, 142(2):572–597, 2021.

[43] Markus Hirschberger, Ralph E. Steuer, Sebastian Utz, Maximilian Wimmer, and Yue Qi. Computing the nondominated surface in tri-criterion portfolio selection. *Operations Research*, 61(1):169–183, 2013.

[44] Sebastian Utz, Maximilian Wimmer, and Ralph E. Steuer. Tri-criterion modeling for constructing more-sustainable mutual funds. *European Journal of Operational Research*, 246(1):331–338, 2015. doi: 10.1016/j.ejor.2015.04.035.

[45] Ana Garcia-Bernabeu, Adolfo Hilario-Caballero, Fabio Tardella, and David Pla-Santamaria. ESG integration in portfolio selection: A robust preference-based multicriteria approach. *Operations Research Perspectives*, 12:100305, 2024. doi: 10.1016/j.orp.2024.100305.

[46] Yeudiel Lara Moreno and Carlos Ignacio Hernández Castellanos. A hierarchical approach to a tri-objective portfolio optimization problem considering an ESG index. *Mathematics*, 12(19):3145, 2024. doi: 10.3390/math12193145.

[47] S&P Dow Jones Indices. S&p 500 shariah index, 2025. https://www.spglobal.com/spdji/en/indices/equity/sp-500-shariah-index/.

[48] S&P Dow Jones Indices. S&p pan arab shariah index, 2025. https://www.spglobal.com/spdji/en/indices/equity/sp-pan-arab-shariah-index/.

[49] Silvia Muzzioli, Lorenzo Vitale, et al. Esg criteria for the evaluation of european countries: a multicriteria analysis (mcda). *DEMB WORKING PAPER SERIES*, 2025.

[50] Ahmet Aytekin. Comparative analysis of the normalization techniques in the context of mcdm problems. *Decision Making: Applications in Management and Engineering*, 4(2):1–25, 2021.

[51] Xi Hang Cao, Ivan Stojkovic, and Zoran Obradovic. A robust data scaling algorithm to improve classification accuracies in biomedical data. *BMC bioinformatics*, 17(1):359, 2016.

[52] Jiawei Han, Micheline Kamber, Jian Pei, Jian Pei, et al. *Data mining: concepts and techniques*, volume 2. Elsevier, 2006.

[53] Nazanin Vafaei, Rita A Ribeiro, and Luis M Camarinha-Matos. Normalization techniques for multi-criteria decision making: analytical hierarchy process case study. In *doctoral conference on computing, electrical and industrial systems*, pages 261–269. Springer, 2016.

[54] Google Developers. Numerical data: Normalization — machine learning, 2025. URL https://developers.google.com/machine-learning/crash-course/numerical-data/normalization. Last updated 2025-12-03.

[55] AAOIFI. Shari'ah standard no. 21: Financial paper (shares and bonds). In: Shari'ah Standards, Accounting and Auditing Organization for Islamic Financial Institutions, Manama, Bahrain, 2017.

[56] Philippe Artzner, Freddy Delbaen, Jean-Marc Eber, and David Heath. Coherent measures of risk. *Mathematical Finance*, 9(3):203–228, 1999.

[57] Carlo Acerbi and Dirk Tasche. On the coherence of expected shortfall. *Journal of Banking & Finance*, 26(7): 1487–1503, 2002.

[58] Vikas Agarwal and Narayan Y. Naik. Risks and portfolio decisions involving hedge funds. *The Review of Financial Studies*, 17(1):63–98, 2004. doi: 10.1093/rfs/hhg044.

[59] Julian Blank and Kalyanmoy Deb. Pymoo: Multi-objective optimization in python. *Ieee access*, 8:89497–89509, 2020.

[60] Charles Audet, Jean Bigeon, Dominique Cartier, Sébastien Le Digabel, and Ludovic Salomon. Performance indicators in multiobjective optimization. *European journal of operational research*, 292(2):397–422, 2021.

[61] Hisao Ishibuchi, Hiroyuki Masuda, Yuki Tanigaki, and Yusuke Nojima. Modified distance calculation in generational distance and inverted generational distance. In *International conference on evolutionary multi-criterion optimization*, pages 110–125. Springer, 2015.

[62] Victor DeMiguel, Lorenzo Garlappi, and Raman Uppal. Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy? *The review of Financial studies*, 22(5):1915–1953, 2009.

[63] William F. Sharpe. The Sharpe ratio. *The Journal of Portfolio Management*, 21(1):49–58, 1994.