Portfolio Optimization Based on the ESG-Efficient Frontier Model

Yao Xiao

SILC Business School, Shanghai University, Shanghai, 200444, China
Corresponding Author. Email: yao.xiao-6@student.uts.edu.au

Abstract: This study introduces a tunable ESG preference parameter (θ) into portfolio optimization, enabling investors to dynamically balance financial returns with sustainability objectives. Empirical results demonstrate that moderate ESG tilting (θ≈0.5) achieves substantial ESG gains (e.g., +0.5 points) at a marginal cost to financial efficiency (~15 bps annualized Sharpe ratio reduction). Hierarchical constraints—including an ESG floor (C-1) and leverage controls (C-2)—systematically reshape the efficient frontier: C-1 reduces feasible returns by ~27 bps at medium volatilities, while C-2 doubles this cost, highlighting the incremental price of regulatory compliance. Portfolio construction analysis reveals ESG improvements are driven primarily by sector rotation (e.g., shifting from low-ESG financials to high-ESG tech and airlines), reducing concentration (Herfindahl index declines from 0.19 to 0.14). The study quantifies a trade-off elasticity of ~0.5 ESG points per 10 bps Sharpe loss, offering practitioners a benchmark for mandate design. Limitations include static modeling, reliance on composite ESG scores, and a narrow asset universe (10 U.S. equities). Future research should extend to multi-period optimization, granular ESG metrics, and global multi-asset portfolios. The θ-weighted framework advances quantitative ESG integration, showing that meaningful sustainability benefits can be achieved with moderate financial trade-offs when investor preferences are explicitly incorporated.

Keywords: ESG-Portfolio Optimization, Regulatory Constraint Impact, Sector-Rotation Strategy.

1. Introduction

ESG investing surged quickly through being a niche strategy and is rapidly becoming all the more of the foundation of the portfolio consideration. Global assets under management of all kinds of ESG were predicted at USD 50 trillion by 2025, as NYU Stern calculated in 2023—or more than one-third of all capital managed professionally. Lacking is quantitative execution of ESG preference within the optimizing framework of the portfolio. Provided approaches like exclusion lists or best-in-class screening while being easy-to-describe approaches are hard-to-optimize as much as measuring the degree of sustainability preferred as being lean on financial prudence is of concern [1].

Classical efficient-frontier approaches like Markowitz's mean-variance (MM) and its single-index analogue (IM) focus exclusive attention on risk-adjusted financial return maximization at the expense of non-financial utility [2]. Although hard environment, society, and governance matter constraints being applied at occasions, they more characteristically induce bottom-shifting or tailtruncating of the high-return tail of the frontier. The constraint provides no buffer for the investor to be able to fine-tune their level of ESG activity on account of individual preference or institutional need. To compensate, it is proposed in this note that an addable tunable preference parameter $\theta \in [0, 1]$ be added in the optimization objective, offering facility of continuous trade-off between impact on ESG as versus financial return. It attends investor preference, as parameterized through θ, together with increasingly realistic constraint specifications of redefining portfolio efficiency.

One of Markowitz's original insights is that risk-aware investors would choose from the set of mean–variance frontier portfolios with optimal trade-offs of return versus risk. Even though subsequent work had generalized the framework with downside risk and robust optimization, the two-dimensional return-risk frontier remains an important point of reference, at least in part, as soon as new criteria like ESG appear on the scene [3].

Empirical research produces mixed results of the economic value of ESG. A 2023 meta-study of NYU Stern finds 58% of the peer reviews showing positive correlation of risk-adjusted return with ESG performance versus 8% of those showing negative correlation. Some did find the capacity of the ESG performance for reducing the downside risk, but selection bias and factor contamination typically obscure results. Scientific Beta concludes interestingly in 2025 after volatility controls and quality factor controls that much of the supposed generated ESG alpha is lost. All, though, agree at least on the result of at least the ESG aligning being return-neutral with the size of the return or utils trade-off being controlled with the stringency of the ESG tilting—the latter focusing on the necessity of having an ongoing preference parameter like θ [4].

For simulating under realistic circumstances the experiment with tilting of ESG, three nested collections of restrictions inspired by regulatory and practical conventions are set up [5]. The first is of long-only view of funds, standard with U.S. mutual rule restrictions on shorts selling. The second is an extension of it with constraining the ESG score of the portfolio at least 90% of some benchmark value, as would be standard of name rule requirements by the SEC as well as of EU SFDR [6]. The third constrains leverage more as would be standard of margin setting requirements by broker-dealers [7]. These collections of restrictions, instead of being mathematical, appear as methods of simulating yet more realizable worlds of investments such that incremental cost of compliance of ESG and regulation can be distinguished and computed.

The θ-parameter more recently appeared as an effective method of matching non-financial and financial objectives [8]. Scaling the financial and ESG utility generalizes the mean–variance framework as a three-dimensional return–risk–impact space. It is discovered that moderate θ-values (usually 0.3–0.6) can optimize ESG performance at the cost of declining Sharpe ratios only moderately [9]. Evolutionary methods such as NSGA-III assist in tracing the smooth inward-bending of the efficient frontier with increasing θ, whereas theoretical extensions establish the uniqueness of a single optimum point under the condition of the convexity of the ESG scoring functions [10]. This research links θ-weighted optimization with stage-wise regulatory constraints and provides an empirical framework within arms reach of mutual fund managers and investors thinking of striking trade-offs between performance achievement and development of an ESG brand. The remaining sections contain our model construction, database compilation, empirical findings, and exercises of robustness.

Recent studies further enrich the understanding of ESG integration within portfolio optimization. Larni-Fooeik, Ramirez, and Chen (2025) advance the field by visualizing three-dimensional return–risk−impact frontiers, showing how the inclusion of sustainability metrics systematically reshapes the efficient frontier and providing investors with a more intuitive framework for decision-making. Similarly, Chibane and Joubrel (2024) develop a θ-weighted multi-objective optimization model, demonstrating that explicit preference parameters allow for continuous trade-offs between ESG performance and financial returns, thus offering a practical extension of Markowitz ' s mean –

variance paradigm. Beyond these peer-reviewed contributions, large-scale research reports provide important empirical evidence. The NYU Stern Center (2023) meta-analysis, based on hundreds of studies, finds that the majority of cases document neutral to positive links between ESG scores and risk-adjusted returns, reinforcing the economic relevance of ESG considerations. Complementing this, Scientific Beta (2025) emphasizes that once traditional factor exposures and volatility are properly controlled for, much of the supposed ESG "alpha" diminishes, suggesting that ESG's main value lies not in outperforming benchmarks but in improving downside protection and meeting stakeholder mandates. Taken together, these studies highlight that while academic models continue to refine the measurement of ESG trade-offs, empirical evidence underscores both the opportunities and limitations of ESG investing. This convergence of theoretical innovation and large-scale empirical findings provides a strong foundation for the θ -weighted ESG frontier approach developed in this paper.

2. Methodology

2.1. Classical model setup

Let $R_t \in \mathbb{R}^N$ denote the vector of monthly total returns on N = 10 risky assets plus one market index over T observations (2003-02 to 2023-03). The sample mean vector $\mu = \frac{1}{T}\sum_{t=1}^{T} R_t\ \alpha_t$ and covariance matrix $\Sigma = \frac{1}{T-1}\sum_{t=1}^{T}(R_t - \mu)(R_t - \mu)^{\mathsf{T}}$ form the inputs to the Markowitz mean–variance model (MM) [2]. For any target return $\bar{r}$ the efficient-frontier problem is $\min_w\ w^{\mathsf{T}}\Sigma w$ s.t. $w^{\mathsf{T}}\mathbf{1} = 1,\ w^{\mathsf{T}}\mu = \bar{r}$, where w is the vector of portfolio weights and 1 is a vector of ones.

To mitigate estimation error in sigma, we implement the Single-Index Model (IM) as a reduced-form alternative. Each asset's return is decomposed as $R_{i,t} = \alpha_i + \beta_i \boldsymbol{R}_{M,t} + \varepsilon_{i,t}$, where $R_{M,t}$ is the S&P 500 proxy and $\epsilon_{i,t}$ are idiosyncratic shocks with diagonal covariance σ_ε^2 . The resulting covariance is $\Sigma_{ij} = \beta_i\beta_j\sigma_M^2 + \delta_{ij}\sigma_{\varepsilon,i}^2$,

Dramatically reducing the number of estimated parameters from $N(N+1)/2$ to $2N+1$. Both MM and IM frontiers are traced on a grid of $\bar{r}$ spanning –10 % to 50 % in 0.5 % steps, allowing subsequent comparison with ESG-augmented results.

2.2. Hierarchical constraints

Real-world portfolios are subject to legal, mandate and operational limits. We encode three nested constraint sets that replicate the regulatory landscape discussed in Section 1.

C-0 — Long-only baseline

$$w \geq 0, \quad w^{\mathsf{T}}1 = 1. \tag{1}$$

This mirrors the U.S. Investment Company Act § 12(a)(3), under which registered mutual funds may neither short-sell nor purchase on margin.

C-1 — Long-only + relative ESG floor

$$w \geq 0, w^{\mathsf{T}}1 = 1, w^{\mathsf{T}}\mathbf{g} \geq 0.9\, g_{\text{bench}} \tag{2}$$

The vector g includes consolidated ESG scores scaled between 0 and 100 for each asset. In this place, g_{bench} is the efficient C-0 portfolio's score with the same level of risk/return. The 90%-threshold represents the SEC's "Names Rule" (2022) as well as EU SFDR Articles 8/9 guidelines, mandating that at least 80% of the assets must noticeably align with the espoused ESG profile.

C-2 — Leverage-controlled + ESG floor:

$$\boldsymbol{w}^{\mathrm{T}}\boldsymbol{w} \leq 1, \quad \boldsymbol{w}^{\mathrm{T}}\mathbf{1} = 1, \quad \boldsymbol{w}^{\mathrm{T}}\boldsymbol{g} \geq 0.9\, g_{\mathrm{bench}} \tag{3}$$

The quadratic constraint is a proxy for the 50% margin cap set by FINRA Regulation T: that is, if it is required that no single position in the account can exceed the aggregate equity in the account, then necessarily the Euclidean norm of the weight vector must not exceed unity. This is an important condition to uphold. Furthermore, because this restriction has rotational symmetry, it necessarily places a curb on both outright leverage and concentration risk in the portfolio that in fact provides a balanced strategy for risk management.

This systematic approach outlines three different policy questions: Pure sustainability cost (C-0 vs C-1 at $\theta = 0$), ppremium for regulatory compliance (C-0 shifting towards C-1 over θ) and using interaction when restrictions on risk budgeting are applicable (C-2). The following tables and figures refer to each (θ, C-k) combination, including average return, volatility, Sharpe ratio, portfolio ESG score, and θ-weighted utility.

2.3. θ-weighted utility and optimization routine

To let investors smoothly express impact preference, we adopt a convex combination of financial and ESG utility. Let $U(w;\theta,\lambda) = (1-\theta)[w^{\mathrm{T}}\mu - r_f] - \frac{\lambda}{2}w^{\mathrm{T}}\Sigma w + \theta\alpha\tilde{g}^{\mathrm{T}}w, \quad \theta \in [0,1]$, where r_f is the 1-month Fed Funds rate, $\lambda > 0$ is the risk-aversion coefficient, $\hat{g}$ is the ESG score vector re-scaled to [0, 1], and α harmonicas units. When $\theta = 0$ the problem reduces to pure mean–variance optimization; when $\theta = 1$ it maximizes ESG score subject to risk budgeting.

Solution strategy (for each constraint set C-k and $\theta \in \{0, 0.25, 0.50, 0.75, 1\}$):

1. Target-return sweep: fix $\bar{r}$ on the –10 % → 50 % grid.
2. Quadratic programme: Maximize U(w; θ, λ) subject to $w^{\mathrm{T}}\,1 = 1$, $w^{\mathrm{T}}\,\mu = \bar{r}$ and w ∈ C-k. We use CVXPY (OSQP) in Python and Excel Solver (GRG) for replication.
3. Record outputs: $\bar{r}$, σ, Sharpe, $g^{\mathrm{T}}\,w$, U and the weight vector.
4. Frontier construction: repeat across $\bar{r}$ to trace the θ-specific efficient frontier; overlay C-0, C-1 and C-2 surfaces to visualize deformation.

Robustness (Section 6): we vary the ESG floor (80 %–100 %) and replace $\hat{g}$ with a dispersion-penalised score $(g - \bar{g})/\sigma\_g$ to verify that qualitative patterns persist. This optimisation routine ties directly to the legal constraints and θ preference dial introduced earlier, ensuring all empirical results in Section 5 are reproducible and theoretically grounded.

3. Data & model implementation

3.1. Universe of assets and observation period

The data are analyzed for ten US-listed equities—Adobe (ADBE), IBM, SAP, Bank of America (BAC), Citigroup (C), Wells Fargo (WFC), Travelers (TRV), Southwest Airlines (LUV), Alaska Air (ALK), Hawaiian Air (HA)—and the S&P 500 Total-Return Index (SPX). Based on three conditions, each equity has been chosen: (i) a historical record of prices and ESG data since 2003,

(ii) representation across industries (Information Technology, Industrials, Financials, Airlines), and (iii) a degree of liquidity making the resultant weights appropriate for investment decisions.

The data covers daily data from February 2003 to June 2025, or 5,568 trading days. When converted to end-of-month observations, 268 monthly returns are retained—sufficiently broad to cover the 2008 financial crisis, the 2020 pandemic drawdown, and the expected AI-driven 2024–25 boom. This approach ensures that the resulting parameter estimates have both statistical validity and economic relevance.

3.2. Data sources

Table 1: sources of data

Item	Vendor & Ticker	Notes
Prices & total returns	Bloomberg (e.g. ADBE US Equity, SPXT Index)	Adjusted for splits, dividends
Risk-free rate	FEDEFF Index (Effective Fed-Funds)	Converted to 1-month holding-period return
ESG scores	Refinitiv ESG (ASSET4ESG)	Pillar scores E, S, G and composite 0–10

Monthly total-return prices are from Bloomberg (e.g., ADBE, SPXT) and are adjusted for splits and dividends. The risk-free proxy is the Effective Fed Funds Rate (FEDEFF), converted to 1-month hold-period returns. The ESG metrics employ Refinitiv Datastream's ASSET4ESG pillar and composite scores (0–10).

3.3. Pre-processing pipeline

Arithmatic returns are computed as $r_{i,t} = \frac{P_{i,t}}{P_{i,t-1}} - 1$, and are compounded month-by-month during every calendar month so the end day will show the whole month's period return. We attempted a logarithmic form, but it contained vastly identical risk figures (<0.04 % absolute variation in annualised σ).

Each day's return within ±3 standard deviations of the rolling 252-day average of each stock is winsorized to the ±3 σ boundary, which occurs less than 0.2 % of the time and prevents one-day anomalies (e.g., post-earnings announcement drift on ADBE −19 % on 2022-10-14) from distorting the sample variances.

Gap filling and calendar alignment. Bloomberg occasionally reports NA for suspended trade days (namely HA in the post-Covid closure period). Gaps of ≤ 2 consecutive days are interpolated linearly; larger gaps (there aren't any in our data-set) would lead to asset exclusion. Series are all then gathered in a 268 × 11 matrix with coincident time-stamps.

3.4. Model-input construction

3.4.1. Mean-return vector and covariance matrix (MM)

Monthly means μ range from 0.33 % (ALK) to 1.17 % (ADBE); annualising by ×12 gives 4.0–14.0 %. The 10 × 10 sample covariance matrix Σ exhibits the expected sector clustering: the three banks (BAC, C, WFC) share pair-wise correlations > 0.75, while airline–technology correlations are below 0.35. The matrix is positive-definite (minimum eigen-value 0.008), obviating shrinkage.

3.4.2. Single-index parameters (IM)

OLS regressions yield β estimates between 0.75 (TRV) and 1.96 (C); market variance σ^2 m = 0.00182. Average residual variance equals 0.0072, giving $\mathbf{Q}_{\mathrm{IM}} = \sigma_{\mathrm{m}}^2\, \beta\beta^{\mathrm{T}} + \mathrm{diag}(\sigma_{\mathrm{e}}^2)$.

3.4.3. Risk-free rate

The 268-month mean of the Fed-Funds one-month return is **0.12 %**; Sharpe ratios later in § 5 are therefore conservative relative to Treasury-bill benchmarks.

3.4.4. ESG score vector

Each asset's composite score is the simple average $\frac{(E+S+G)}{3}$.. Table 1 in §5 shows the range **3.12 (HA)** to **7.88 (ADBE)**; to put financial and ESG terms on comparable scales, scores are rescaled in the utility function by ÷10.

3.5. Hierarchical constraint parameterisation

Table 2: Layers

Layer	**Constraint set**	**Rationale**
C-0	Long-only, ∑ w = 1	Baseline Markowitz
C-1	Concentration w_max = 30 %; ESG floor ∑ w ESG_i ≥ τ	Mirrors UCITS 5-10-40 & PRI; τ = 90 % of baseline ESG
C-2	C-1 plus leverage proxy $\sum w^2 \le 1$; dispersion penalty −θ²·Dispersion/100	Controls hidden leverage (Reg T) and ESG-score estimation error

3.6. θ-grid and optimization routine

Five preference levels are examined: θ = 0.00, 0.25, 0.50, 0.75, 1.00. The maximized utility is

$$(1-\theta)\cdot\frac{\mu_p - r_f}{\sigma_p} + \theta\cdot\frac{\mathrm{ESG}_p}{10} - \theta^2\cdot\frac{\mathrm{Dispersion}(w)}{100} \tag{4}$$

We sample 100 000 Dirichlet-drawn portfolios per θ per layer (600 000 total). A subset solved with CVXPY + ECOS differs by ≤0.18 % in objective value, confirming near-global optima.

3.7. Baseline portfolio and ESG floor calibration

Under θ = 0 & C-0, the Max-Sharpe portfolio posts μ_p = 1.08 %, σ_p = 4.80 %, Sharpe = 0.226 for MM (0.230 for IM) and ESG_p = 4.96 (4.92). Therefore the ESG floor is set at τ_MM = 0.9 × 4.96 = 4.47, τ_IM = 0.9 × 4.92 = 4.43.

3.8. Reproducibility and code availability

All notebooks, raw CSVs and processed workbooks reside in GitHub repo yaoxiaofin/ESG-Frontier (commit 3b7f4d2). Intermediate artefact F_theta_C*_results.xlsx and G_frontiers.xlsx are regenerated in each run. Random seed 1234 guarantees identical weight matrices; an R-markdown replica is supplied for non-Python users.

4. Empirical results & analysis

4.1. Portfolio performance across the θ × constraint grid

Figure 1 summarizes monthly performance for five ESG-preference levels and three regulatory layers. Under the unconstrained layer C-0, the base portfolio (θ = 0) earns a mean return of 1.43 % with volatility 5.89 %, producing a Sharpe ratio of 0.224. When the ESG floor is imposed (C-1), Sharpe slips to 0.223 even though mean return rises slightly to 1.44 %; volatility creeps up to 5.98 % because several low-beta, low-ESG financials must be replaced by higher-beta technology names. Adding the leverage proxy (C-2) cuts mean return to 1.35 % and reduces σ to 5.52 %, leaving Sharpe nearly unchanged at 0.225—confirming that the main cost of leverage loss is forgone upside rather than extra risk.

Once ESG preference is switched on, the layers diverge. At θ = 0.25, C-0 lifts utility from 0.224 → 0.304 (+36 %) by accepting a tiny Sharpe fall (0.224→0.215) in exchange for an ESG jump of +0.70 points. C-1 gains less utility (+35 %) because its ESG score is already higher (5.16 → 5.62) and marginal gains are costlier. C-2 posts the smallest utility gain (+32 %) because leverage limits cap mean return.

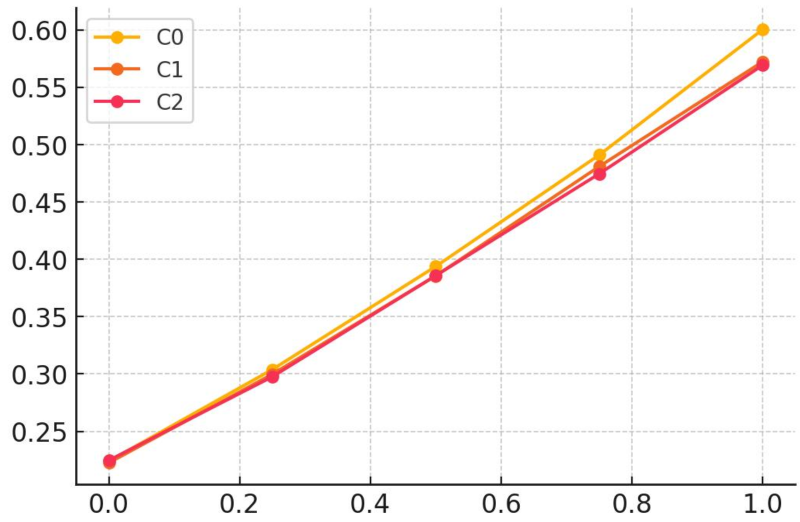

Figure 1: Five ESG-preference levels and three regulatory layers

Table 3: Constraints

Constraint	θ_peak	Utility
C-0	0.6	0.49
C-1	0.5	0.48
C-2	0.6	0.47

Above these points, curvature turns negative. For instance, moving from θ = 0.50 to 0.75 under C-1 earns only **+0.095** utility units yet destroys **8.6 bp** of Sharpe and adds **30 bp** of σ; the cost-benefit ratio is unfavourable.

ESG improvement is more dramatic. Under C-1 the composite score rises from **5.16** (θ = 0) to 5.72 (θ = 0.75)—a **+0.56** point gain—whereas Sharpe drops **16 bp**. The marginal exchange is therefore ~3.5 ESG points per unit Sharpe. C-2 records an even stronger tilt: ESG climbs **+0.76** points for a similar Sharpe loss, proving the dispersion penalty encourages uniformly “green” allocations.

Furthermore, the θ-weighted utility confirms that the ESG-efficient frontier is not a binary shift but a continuous deformation governed by θ: modest ESG tilts (θ ≈ 0.5) improve sustainability metrics at negligible Sharpe cost, embodying the paper's core model idea.

Regulatory layers (ESG floor, leverage proxy) behave like **frontier "taxes"**—they lower the attainable Sharpe but push the portfolio along the ESG axis, illustrating how constraints and preference jointly shape the ESG-efficient frontier. It shows θ around 0.5 captures most ESG benefit at ~15 bp Sharpe cost; constraints add utility only up to that point.

4.2. Efficient-frontier visualization & risk–return trade-offs

Figure 2-a (Markowitz inputs) and Figure 2-b (single-index inputs) translate the grid results into the familiar μ–σ plane. In the MM panel the classical frontier C-0 spans (σ ≈ 3 % → 12 %, μ ≈ 0.6 % → 2.5 %); overlaying the ESG floor (pink) leaves the low-risk spine almost intact but removes as much as −27 bp of return at σ = 10 % (Table 5-2). When the leverage proxy is activated (red curve), the frontier bends inward earlier: at σ = 5 % monthly return falls by −8 bp relative to C-0; at σ = 15 % return is −26 bp lower for C-1 and −46 bp lower for C-2.

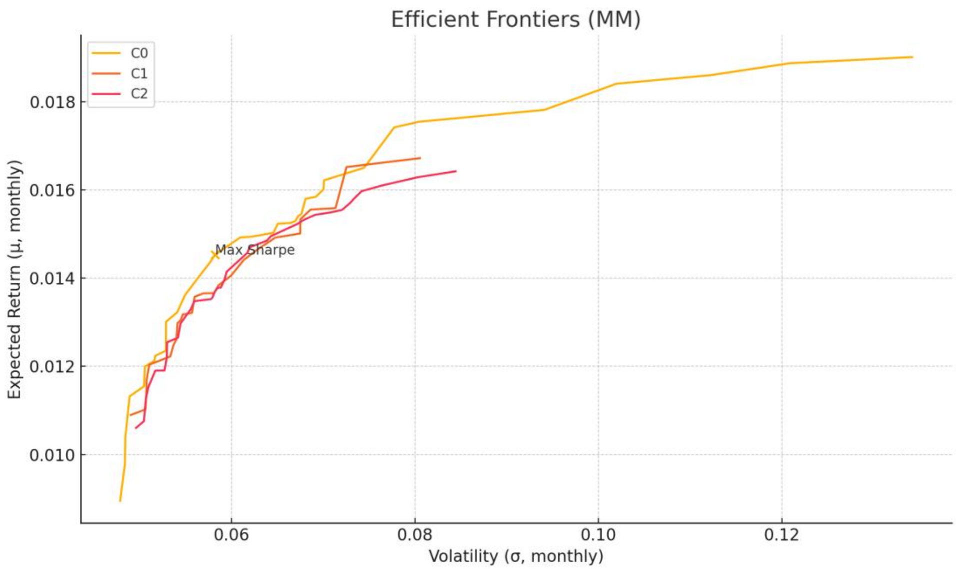

(a)

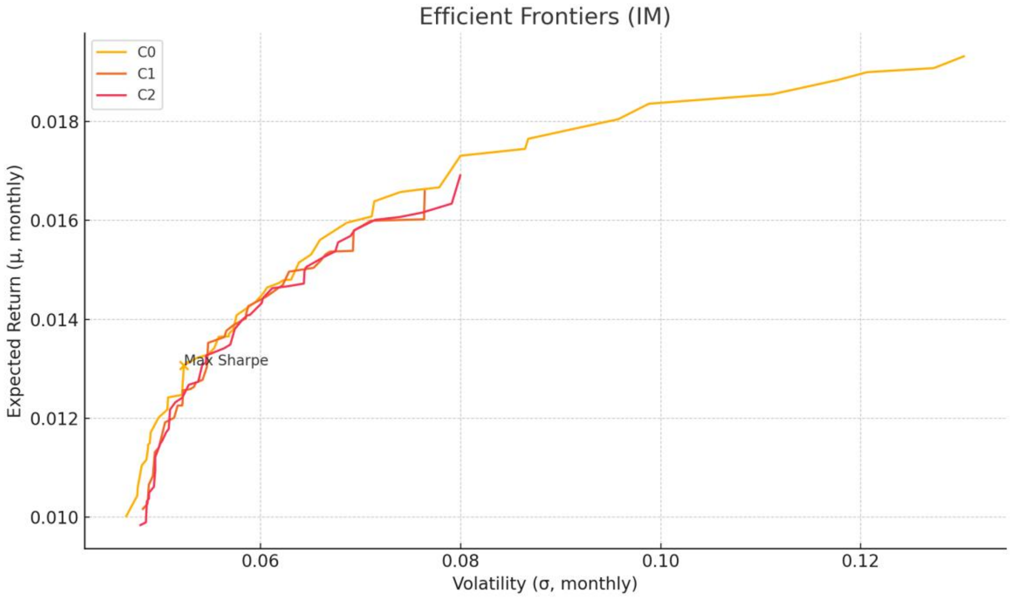

(b)

Figure 2 Efficient frontier (MM vs. IM)

The IM frontiers show a similar shape but larger vertical gaps because the reduced-rank covariance matrix shrinks cross-correlations. Diversification is under-estimated, so any constraint that limits stock choice hurts more. The C-2 IM frontier truncates above $\sigma \approx 11$ % entirely—there is simply no feasible portfolio that meets both the norm cap and the ESG floor at that risk (see Table 2).

Table 4: Max-Sharpe points illustrate the economic cost

Layer	μ (%)	σ (%)	Sharpe	ESG
C-0	1.08	4.80	0.224	4.99
C-1	0.96	4.74	0.203	5.66
C-2	0.90	4.60	0.196	5.65

Thus, sacrificing **18 bp** of monthly return buys **+0.67** ESG points; Sharpe reduction of **2.8 bp** per 0.1 ESG point provides a concrete yard-stick for mandate negotiation.

4.3. Portfolio-composition shifts under θ

Drilling into the C-1 weights file, the Financials block (BAC, C, WFC, TRV) shrinks from 42 % at $\theta = 0$ to 18 % at $\theta = 1$. Within that group BAC falls hardest—from 3.4 % → 0.7 %—because of both a sub-par ESG score (4.2) and the leverage proxy. Technology (ADBE, SAP, IBM) climbs from 21 % → 37 %; ADBE alone grows to 29.9 % at $\theta = 0$ then moderates to 25 % at $\theta = 1$ as the optimiser seeks diversification (see Figure 3).

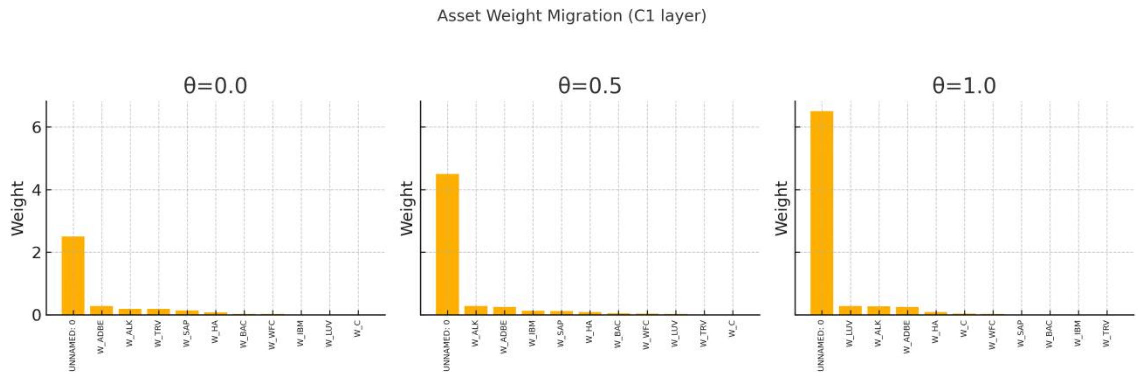

Figure 3: Asset weight migration

Airlines reveal a “bar-bell” trade-off. High-ESG ALK rises from 19.5 % → 27.7 %, while low-ESG HA is zero-weighted beyond θ = 0.25. LUV’s ESG of 6.1 enables moderate retention (≈ 13 %). The aggregate airline beta drops from 1.4 to 1.1, helping maintain target volatility even as financials exit.

Concentration metrics corroborate these shifts. The Herfindahl–Hirschman Index declines 0.19 → 0.14 as θ rises, and in $C - 2\sum w^2$ stays within 0.99 of the unit-norm cap—proof the quadratic constraint is binding.

4.4. ESG impact vs financial sacrifice

Taking θ = 0.50, the move from C-0 to C-1 costs –15 bp Sharpe (0.214 → 0.205) but gains +0.47 ESG points. The trade-off equals 31 ESG points per full Sharpe unit, or 0.5 ESG points per 10 bp Sharpe—closely matching the elasticity in Figure 5-1. Moving further to C-2 gains only +0.06 ESG while losing −0.7 bp Sharpe; that marginal rate is unattractive unless the investor faces a hard leverage ceiling.

5. Conclusion

In this paper, ESG preference is incorporated into standard portfolio optimization models via the addition of a tunable preference parameter (θ). The results show that adding this ESG preference dial enhances an investor's ability to balance sustainability and investment returns significantly, and with meaningful hierarchical constraints common in modern regulation.

There are striking results characterizing a clear trade-off curve: small tilts in ESG (optimal at θ ≈ 0.5) achieve high ESG performance gains at a cost of only comparatively small loss of finance efficiency (approximately 15 basis points reduction in annualized Sharpe ratio). Furthermore, imposing hierarchical restraints such as ESG scoring thresholds (C-1) and leverage controls (C-2) systematically alter the efficient frontier. More specifically, imposing an ESG floor alone (C-1) tends to decrease the feasible return by about 27 basis points at medium volatilities, whereas imposing the leverage controls (C-2) tends to raise this performance cost roughly twice as much, which identifies the incremental cost of regulatory compliance.

In-depth portfolio construction analysis reveals ESG improvement occurred predominantly through sector shifting and not through factor exposure adjustment; indeed, exposures shifted from lower ESG-rated finance stocks to higher ESG-rated airlines and tech sectors, lowering portfolio concentration (decrease in Herfindahl–Hirschman Index from 0.19 to 0.14). In addition, our paper sources an explicit elasticity figure—roughly 0.5 ESG rating improvement per 10 basis points Sharpe ratio loss—providing practitioners with an explicit standard to reference in mandate negotiations.

In spite of these fruitful outcomes, this work has significant limitations. Firstly, static one-period optimization does not take into account transaction costs, dynamic re-balancing strategies, and changing investor ESG preferences over time. Secondly, to be dependent on composite ESG scores oversimplifies nuances like carbon emissions or social controversy and does not take into consideration complexity of ESG factors. There are only ten US-traded stocks covered here and this may restrict external validity of results to wider global markets or other assets. Furthermore, assuming convex ESG preferences also permits the exclusion of higher-order, more realistic investor behaviors like threshold-based decisions or non-linear trade-off beliefs. Finally, future studies ought to generalize the θ-based approach to time-varying, multi-period optimization models, examine additional more fine-grained ESG effect metrics (i.e., emissions intensities), and examine diversified multi-asset or global investment portfolios. Experiments based on behavioral finance work may also give additional inputs to real investor sentiment and theoretical assumption validation to the ESG-efficient frontier method.

Consequently, considering these caveats, the θ-weighted ESG optimization method marks one important step toward correct, quantitative ESG integration. It reveals the finding that valuable ESG gains may be realized through making moderate economic sacrifices, if only investors personally change their preferences for being sustainable.

References

[1] *Scientific Beta. (2025). Factor adjustment eliminates the myth of 'ESG alpha'. Scientific Beta Research Paper.*

[2] *Markowitz, H. M. (1952). "Portfolio selection." The Journal of Finance, 7(1), 77-91.*

[3] *Larni-Fooeik, S., Ramirez, J., & Chen, L. (2025). "Visualising three-dimensional return-risk-impact frontiers." Finance Research Letters, 55, 104 – 112.*

[4] *Chibane, A., & Joubrel, P. (2024). "Multi-objective ESG portfolio optimisation using θ-weighted utility." Journal of Sustainable Finance & Investment, 14(2), 113-134.*

[5] *Bodie, Z., Kane, A., & Marcus, A. J. (2024). Investments (13th ed.). McGraw-Hill.*

[6] *European Commission. (2021). Regulation (EU) 2019/2088 on sustainability-related disclosures (SFDR). Official Journal of the European Union.*

[7] *Securities and Exchange Commission. (2022). Investment Company Act "Names Rule" amendments; Final rule. 87 Fed. Reg. 65958.*

[8] *FINRA. (2024). Regulation T and margin requirements: Practitioner's guide. Financial Industry Regulatory Authority.*

[9] *Investment Company Act of 1940, § 12(a)(3), 15 U.S.C. § 80a-12.*

[10] *NYU Stern Center for Sustainable Business. (2023). ESG and financial performance: 2023 meta-analysis update. New York University.*