Downside Risk Measures and ESG Factors in Optimal Portfolio Construction: Evidence from European Equity Markets

Abstract

In this paper, we implement an integrated framework for constructing ESG-constrained, downside-risk-optimized equity portfolios in the European stock market. Extending traditional mean-variance approaches, we employ downside-oriented risk measures—conditional value at risk (CVaR) and semi-variance—to better capture investors' asymmetric aversion to losses. ESG scores are introduced as binding constraints based on percentile thresholds, ensuring that portfolios comply with predefined sustainability standards. Semi-variance and CVaR objectives are formulated as convex programs to enable tractable optimization. Using data from Euro Stoxx 50 and Euronext 100 constituents, our empirical analysis reveals that: (i) integrating downside risk measures enhances tail-risk protection and may improve performance for loss-averse investors; but (ii) enforcing ESG constraints, particularly at stricter thresholds, leads to reduced diversification and a decline in risk-adjusted returns (e.g., Sharpe and Sortino ratios). These findings highlight the inherent trade-off between sustainability and financial efficiency, underscoring the importance of moderate ESG integration when balancing performance and ethical objectives.

Key words: ESG factors, Downside Risk Measures, portfolio optimization.
JEL Classification: C61, G11, G18.

1. Introduction

Modern portfolio theory (MPT), as pioneered by Markowitz (1952), frames portfolio construction as a mean–variance (MV) trade-off, defining risk via return variance and deriving an efficient frontier that balances expected return and volatility. However, variance treats upside and downside deviations symmetrically, despite investors' aversion to loss, and fails to capture extreme tail events that often drive realized drawdowns (Sortino & Price, 1994). These limitations have spurred the adoption of downside-focused risk measures—Conditional Value at Risk (CVaR), Value at Risk (VaR), and semi-variance—which isolate losses below a threshold and align risk assessment with investors priorities (Acerbi & Tasche, 2002; Rockafellar & Uryasev, 2000).

Concurrently, Environmental, Social, and Governance (ESG) considerations have moved to the forefront of investment decision-making. Investors and regulators now emphasize sustainability metrics alongside financial factors. Empirical studies report that companies with high ESG scores often deliver financial returns equal to or better than lower scoring peers, supporting the "fiduciary duty" of sustainability integration (Coqueret, 2022). As Chen and Mussalli (2020), Chen et al. (2021), and Coqueret (2022) found that the rise of SRI and ESG investing has galvanized both practitioners and academics, prompting a wide range of portfolio-level studies on performance, abnormal returns, and ESG-risk mitigation.

A growing strand of this literature focuses on ESG-aware portfolio optimization. Hirschberger et al. (2013), Utz et al. (2014), and Gasser et al. (2017) extend Markowitz's bi-criterion model to a tri-criterion framework that optimizes return, risk, and ESG scores. In this line, Alessandrini and Jondeau (2020) introduce turnover, tracking-error, and factor-exposure constraints to their ESG optimization; Chen et al. (2023) use a cross-efficiency data envelopment analysis (DEA) to combine ESG scores with financial measures; and Pedersen et al. (2021) derive an "ESG-efficient frontier" to quantify the trade-off between sustainability levels and maximum Sharpe ratios. More recently, Steuer and Utz (2023) map a three-dimensional efficient surface with non-contour cross-sections for ESG, return, and risk, while Abate et al. (2021) investigate downside-risk extensions of ESG-efficient models. Xidonas and Essner (2022) further explore multi-objective formulations that maximize ESG metrics subject to risk constraints.

Despite recent advances, existing approaches tend to focus either on mean–variance ESG models or on downside-risk optimization without sustainability constraints, often overlooking the interplay between both. However, the integration of ESG criteria—particularly through binding constraints—can significantly alter portfolio composition, often reducing diversification and deteriorating performance. This creates a practical tension between ethical investing and financial efficiency. To address this gap, we propose a unified and tractable convex optimization framework that jointly minimizes downside risk (via CVaR and semi-variance objectives) and enforces ESG compliance through weighted percentile-based constraints. Our approach allows for a systematic comparison of the trade-offs involved, highlighting how stricter ESG standards may lead to higher portfolio concentration and lower Sharpe and Sortino ratios, particularly in the context of European equity markets.

The paper is organized as follows. Section 2 presents optimization models and ESG integration. Section 3 describes the data used and the empirical results on the Euro Stoxx 50 and Euronext 100 datasets, including comparisons of Sharpe and Sortino ratios across models. Finally, Section 4 concludes and discusses implications for sustainable risk management in European equity portfolios.

2. Portfolio Models and ESG integration

2.1 Markowitz's MV model and Sharpe formulation

Markowitz (1952) proposed the traditional mean-variance (MV) framework, in which the portfolio variance is minimized for a given expected return. This model assumes that investors are risk-averse and base their decisions on the trade-off between expected return and risk, measured by variance. The Markowitz' optimization problem is solved by minimizing the variance, $\sigma_P^2 = w'\Sigma w$, which is the risk measure of the portfolio, for a given expected return, $\mu_p = w'\mu$, as follows:

$$\min_{\{w\}} \{w'\Sigma w\} \qquad \begin{aligned} s.t.\ \mu' w &= \mu_p, \\ w'\mathbf{1} &= 1, \\ w &\geq 0 \end{aligned} \tag{1}$$

Where, μ_p is the expected return of the portfolio, $w \in \mathbb{R}^{n\times 1}$ is the vector of weights of the assets, $\mathbf{1} \in \mathbb{R}^{n\times 1}$ is a vector of ones. The constraint $w \geq 0$ implies no short-selling.. Alternatively, the Sharpe model scales the portfolio to achieve an expected return adjusted by risk, simplifying the tangent optimal portfolio construction as:

$$\min_{\{w\}} \{w'\Sigma w\} \qquad \begin{aligned} s.t.\ w'\mu &= 1, \\ w &\geq 0 \end{aligned} \tag{2}$$

Both formulations are solved using quadratic programming (QP), leveraging the convexity of the objective function and the linearity of the constraints.

2.2 Semivariance Framework and Sortino model

The Sortino model refines the Markowitz framework by replacing variance with semi-variance, focusing only on downside deviations from a target return or a threshold. This is particularly relevant for risk-averse investors concerned more with losses than with total variability. Optimization proceeds by minimizing the semi-variance subject to return and budget constraints:

$$\min_{\{w\}} \{w'\tilde{\Sigma} w\} \qquad \begin{aligned} s.t.\ \mu' w &= \mu_p, \\ w'\mathbf{1} &= 1, \\ w &\geq 0 \end{aligned} \tag{3}$$

where $\tilde{\Sigma}$ is the semi-variance matrix constructed using returns below the threshold. The main difference with the Markowitz model lies in the risk measure: while variance penalizes all volatility symmetrically, semi-variance focuses solely on downside risk, often producing more conservative portfolios in the presence of asymmetric return distributions.

2.3 VaR and CVaR models

Conditional Value at Risk (CVaR), also known as Expected Shortfall (ES), provides a coherent measure of downside risk, addressing limitations of Value at Risk (VaR) such as non-subadditivity. Following Rockafellar and Uryasev (2000), CVaR-based portfolio optimization can be expressed as a convex linear programming problem. For a loss function

$L(w,r) = -w'r$ and a confidence level α\alphaα, the CVaR minimization problem is formulated as:

$$\min_{\{w,\xi_i\}} \left\{ \eta + \frac{1}{(1-\alpha)N} \sum_{i=1}^{N} \xi_i \right\} \qquad \begin{aligned} s.t.\ \xi_i &> \mu' r^{(i)} - \eta, \\ \mu' w &= \mu_p, \\ w' \mathbf{1} &= 1, \\ w, \xi_i &\geq 0 \end{aligned} \tag{4}$$

where $\eta \in R$ is the VaR at level α, ξ_i are auxiliary variables representing tail losses, and $r^{(i)}$ are historical return scenarios. This model minimizes the expected loss in the worst $(1-\alpha)$ of cases. This formulation is convex and solvable with standard linear programming solvers.

2.4 ESG Integration into the portfolio optimization models

ESG considerations are integrated into the portfolio optimization models as linear constraints that enforce a minimum environmental, social, and governance (ESG) score for the portfolio. Let $e \in \mathbb{R}^{n\times 1}$ be the vector of ESG scores of the assets, and $\mathcal{E}_\mathrm{p}$ is the ESG target from the optimal portfolio. Following to Utz et al. (2014) and Gasser et al. (2017), the MV-ESG optimal portfolio with a given ESG target $\mathcal{E}_\mathrm{p} = w'e$, for the convex function is given by:

$$\min_{\{w\}} \{w' \Sigma\, w\} \qquad \begin{aligned} s.t.\ w'\mu &= \mu_p, \\ w'e &= \mathcal{E}_p, \\ w'\mathbf{1} &= 1, \\ w &\geq 0 \end{aligned} \tag{5}$$

In this sense, the previous Sharpe, Sortino and M-CVaR models can be reformulated as follows:

$$\min_{\{w\}} \{w' \Sigma\, w\} \qquad \begin{aligned} s.t.\ w'\mu &= 1 \\ w'e &= \mathcal{E}_p, \\ w &\geq 0 \end{aligned} \tag{6}$$

$$\min_{\{w\}} \{w' \tilde{\Sigma}\, w\} \qquad \begin{aligned} s.t.\ w'\mu &= 1 \\ w'e &= \mathcal{E}_p, \\ w &\geq 0 \end{aligned} \tag{7}$$

$$\min_{\{w,\xi_i\}} \left\{ \eta + \frac{1}{(1-\alpha)N} \sum_{i=1}^{N} \xi_i \right\} \qquad \begin{aligned} s.t.\ \xi_i &> \mu' r^{(i)} - \eta, \\ \mu' w &= \mu_p, \end{aligned} \tag{8}$$

$$w'e = \mathcal{E}_p,$$
$$w'\mathbf{1} = 1,$$
$$w, \xi_i \geq 0$$

This integration ensures that the portfolio not only meets financial risk-return objectives but also satisfies sustainability criteria. As demonstrated by Utz et al. (2014), Gasser et al. (2017), and Steuer & Utz (2023), ESG-constrained models can be used to construct an efficient frontier in three dimensions—return, risk, and ESG score—enabling investors to navigate the trade-offs among these competing goals.

3. Numerical example and results

3.1 Data

The proposed model is applied to the European equity market, with a specific focus on the Euronext 100 and Euro Stoxx 50 indices, which serve as representative benchmarks of large-cap and highly liquid stocks across the Eurozone. The Euronext 100 (NDX 100) captures the performance of the most actively traded companies listed on the Euronext exchanges, while the Euro Stoxx 50 (SX 50E) reflects the market capitalization and sectoral breadth of leading blue-chip firms in the region. The empirical analysis relies on historical monthly prices of the constituent stocks from January 2015 to August 2024. All financial and ESG-related data were obtained from Bloomberg, a widely recognized and reputable source of global market information. In the case of ESG indicators, we utilized the composite ESG score available through Bloomberg's proprietary methodology, which integrates environmental, social, and governance dimensions into a single performance metric. This score facilitates consistent cross-sectional comparisons and serves as a crucial input for ESG-integrated portfolio optimization models.

Given the well-documented challenges surrounding ESG data heterogeneity and rating divergence across providers, the selection of a robust and transparent data source such as Bloomberg is essential to ensure the reliability and replicability of the results. Prior to analysis, the dataset was systematically cleaned, and missing values were addressed using standard statistical imputation techniques, preserving the integrity of the time series structure. Table 1 summarizes the main descriptive statistics—including expected return, standard deviation (volatility)—of the monthly returns for the stocks with the highest ESG scores in each index.

Table 1. Descriptive data for the top 10 stocks and ESG scores.

NDX 100				**SX 50E**			
Stocks	**Return (%)**	**Volatility (%)**	**ESG score**	**Stocks**	**Return (%)**	**Volatility (%)**	**ESG score**
RYA ID	0.47	9.32	7.35	SU FP	1.38	6.67	7.75
SHELL NA	0.58	6.88	7.15	TTE FP	0.81	6.63	6.96

AKRBP NO	2.01	11.72	7.11	ASML NA	1.98	7.72	6.83
TTE FP	0.81	6.63	6.96	ENI IM	0.51	7.37	6.8
ASML NA	1.98	7.72	6.83	ENEL IM	0.96	6.19	6.68
LDO IM	1.02	11.16	6.83	DB1 GY	1.26	5.42	6.53
ENI IM	0.51	7.37	6.8	NOKIA FH	-0.23	9.21	6.46
BESI NA	2.58	12.83	6.78	MUV2 GY	1.30	6.01	6.4
EDP PL	0.60	6.08	6.76	ITX SQ	0.85	7.41	6.38
ENEL IM	0.96	6.19	6.68	SIE GY	0.87	7.15	6.37

Note: Ryanair Holdings (RYA ID), Shell PLC (SHELL NA), Aker BP ASA (AKRBP NO), TotalEnergies SE (TTE FP), ASML Holding NV (ASML NA), Leonardo SpA (LDO IM), ENI SpA (ENI IM), BE Semiconductor Industries NV (BESI NA), EDP - Energias de Portugal SA (EDP PL), and Enel SpA (ENEL IM), chneider Electric SE (SU FP), Deutsche Börse AG (DB1 GY), Nokia Oyj (NOKIA FH), Münchener Rückversicherungs-Gesellschaft AG (MUV2 GY), Industria de Diseño Textil S.A. (Inditex) (ITX SQ), and Siemens AG (SIE GY).

Source: own elaboration.

3.2 Results and discussion

This section presents the empirical results of the portfolio optimization process across five key methodologies: Mean-Variance (MV), Maximum Sharpe Ratio (Max. SR), Maximum Sortino Ratio (Max. SoR), Conditional Value at Risk (CVaR), and Mean-CVaR (M-CVaR) portfolios, initially without incorporating ESG constraints. We begin by reporting the baseline performance metrics—expected return, volatility, CVaR, and risk-adjusted measures such as the Sharpe and Sortino ratios—of portfolios optimized solely on financial criteria. This provides a benchmark for subsequent comparison. Next, we analyse the performance of portfolios optimized using the same methodologies but under the inclusion of ESG constraints, implemented through minimum ESG score thresholds. This allows us to evaluate the impact of ESG integration on portfolio efficiency and risk-adjusted returns. We compare how downside-risk-optimized portfolios (CVaR and M-CVaR) perform relative to traditional MV portfolios, both with and without ESG considerations.

We then explore the variations in portfolio composition resulting from the different optimization models and degrees of ESG integration. This analysis highlights how the choice of risk metric and the introduction of ESG constraints influence asset allocation decisions. To ensure the robustness of our findings, we conduct the optimization for two distinct European equity benchmarks—NDX 100 and SX 50E—and examine results under varying ESG selection criteria by applying percentile-based thresholds. Specifically, we select the top-performing firms based on ESG scores using the 60th, 70th, and 80th percentiles, thereby identifying sustainability leaders within the investment universe.

Table 2. Results for the MV, Max. SR, Max. SoR, CVaR, and M-CVaR portfolios without ESG constraints.

NDX 100

	Mean	Volatility	CVaR	Sharpe Ratio	Sortino Ratio
MV	0.0094	0.0281	-0.0597	0.333	0.616
Max. SR	0.0201	0.0384	-0.069	0.5223	1.0983
Max. SoR	0.0211	0.0433	-0.0641	0.4875	1.2119
CVaR	0.0119	0.0353	-0.0445	0.3358	0.7733
M-CVaR	0.0195	0.0425	-0.0534	0.4573	1.2439
			SX 50E		
	Mean	**Volatility**	**CVaR**	**Sharpe Ratio**	**Sortino Ratio**
MV	0.011	0.0337	-0.0606	0.3254	0.6811
Max. SR	0.0153	0.0378	-0.0631	0.4047	0.8651
Max. SoR	0.0153	0.0384	-0.0633	0.398	0.8703
CVaR	0.0128	0.0383	-0.0495	0.3334	0.7712
M-CVaR	0.0137	0.0371	-0.0517	0.3683	0.8374

Source: own elaboration.

Table 3. Results for the MV, Max. SR, Max. SoR, CVaR, and M-CVaR portfolios with ESG constraints (60th percentile)

			NDX 100		
	Mean	**Volatility**	**CVaR**	**Sharpe Ratio**	**Sortino Ratio**
MV	0.0097	0.029	-0.061	0.3324	0.6145
Max. SR	0.0198	0.0387	-0.07	0.5121	1.0928
Max. SoR	0.0198	0.041	-0.064	0.4833	1.1822
CVaR	0.0135	0.036	-0.047	0.3752	0.9319
M-CVaR	0.0193	0.0449	-0.054	0.4307	1.204
			SX 50E		
	Mean	**Volatility**	**CVaR**	**Sharpe Ratio**	**Sortino Ratio**
MV	0.0112	0.034	-0.063	0.3298	0.6808
Max. SR	0.015	0.0374	-0.061	0.4001	0.8594
Max. SoR	0.015	0.0383	-0.063	0.3906	0.8584
CVaR	0.0128	0.0369	-0.053	0.347	0.7714
M-CVaR	0.0142	0.04	-0.055	0.3561	0.8111

Source: own elaboration.

The empirical results show that while the integration of ESG constraints into the optimization process leads to a moderate erosion of performance metrics - particularly in terms of expected return, Sharpe ratio and Sortino ratio - the resulting ESG-constrained portfolios continue to outperform their respective benchmarks. For both the NDX 100 and SX 50E universes, portfolios constructed without ESG constraints delivered higher risk-adjusted returns in most models, particularly under the Sortino and M-CVaR frameworks. However, when ESG criteria are imposed (60th percentile threshold), the underperformance remains relatively limited. For example, the SR portfolio for the NDX 100 universe declines slightly in mean return (from 0.0201 to 0.0198) and Sharpe ratio (from 0.5223 to 0.5121), while maintaining its relative efficiency. Similar patterns are observed in all the other optimization models and in the SX 50E results. Importantly, as shown in the cumulative returns (Figure 1), the ESG-integrated portfolios continue to outperform the benchmarks over the entire sample period, demonstrating that sustainability constraints, while imposing certain trade-offs, do not compromise the portfolios' ability to deliver robust long-term performance.

Figure 1. Cumulative returns for optimal portfolios and benchmarks

a. NDX 100

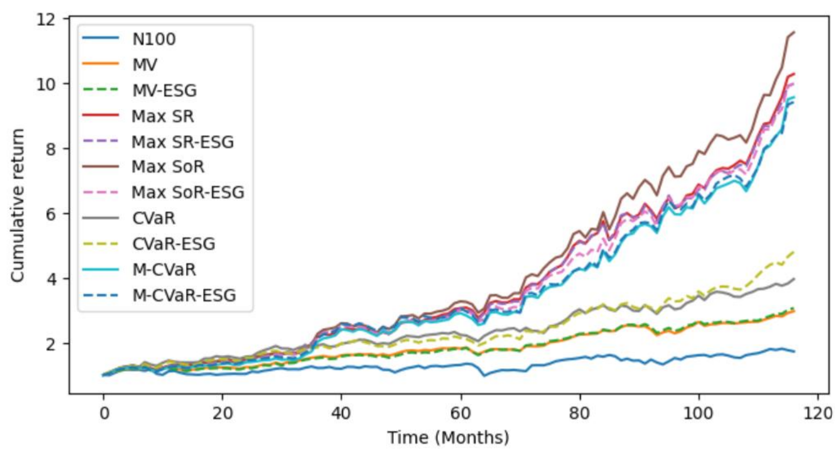

a. SX 50E

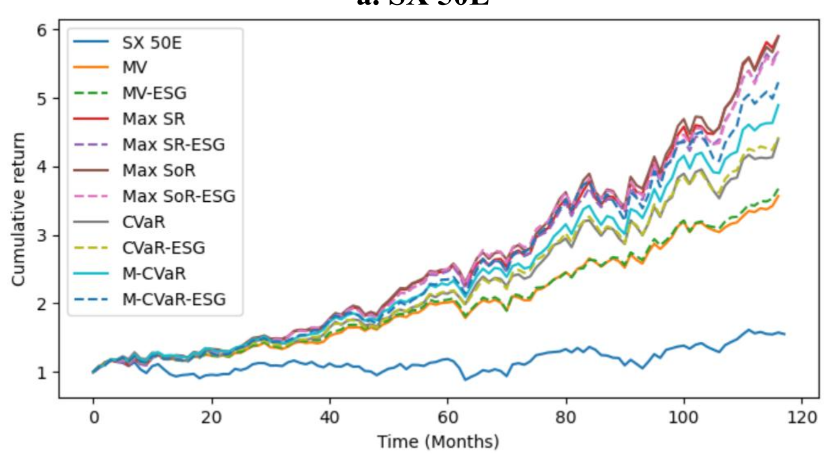

Source: own elaboration.

3.3 Sensitivity Analysis: ESG Constraints with 70th and 80th Percentiles

To evaluate the robustness of the ESG integration in the portfolio construction process, we conducted a sensitivity analysis by tightening the ESG filter to retain only the top-performing companies in each reference index (NDX 100 and SX50E), based on the 70th and 80th percentiles of ESG scores. Table 4 presents the results when the top 30% (≥70th percentile) of companies by ESG score are held. In both indices, the Sharpe and Sortino ratios for the Max. SR and Max. SoR portfolios remain superior, highlighting the efficiency of these optimization approaches under ESG constraints. Notably:

In the NDX100, the Sortino Ratio reaches 1.132 for the Max. SoR strategy, indicating strong downside risk-adjusted returns. In the SX50E, although absolute returns are slightly lower, the Max SR and Max SoR strategies still show improved ratios compared to MV and CVaR.

Table 4. Optimal portfolios with ESG constraints by taking 70th percentile

			NDX 100		
	Mean	**Volatility**	**CVaR**	**Sharpe Ratio**	**Sortino Ratio**
MV	0.0097	0.0298	-0.062	0.3265	0.6015
Max. SR	0.0196	0.0391	-0.072	0.5002	1.063
Max. SoR	0.0196	0.0409	-0.066	0.4792	1.132
CVaR	0.0142	0.0369	-0.048	0.3846	0.9551
M-CVaR	0.0181	0.0424	-0.052	0.4261	1.1462
			SX 50E		
	Mean	**Volatility**	**CVaR**	**Sharpe Ratio**	**Sortino Ratio**
MV	0.0115	0.035	-0.063	0.3273	0.6703
Max. SR	0.0148	0.0381	-0.063	0.388	0.8332
Max. SoR	0.0148	0.0389	-0.064	0.3801	0.831
CVaR	0.0125	0.0379	-0.057	0.3289	0.7184
M-CVaR	0.0151	0.0432	-0.06	0.3509	0.7904

Source: own elaboration.

Table 5 further tightens the constraint by selecting only the top 20% (≥80th percentile) of ESG-compliant firms. As expected, this leads to a marginal decline in mean returns and risk-adjusted performance, due to reduced diversification. However, in this case, the NDX100 portfolio still maintains a Sharpe Ratio above 0.45 for Max SR, and the Sortino Ratio stays close to 1. The SX50E results show resilience in the M-CVaR strategy, which achieves the highest mean return (0.0162) and a reasonable Sortino Ratio (0.7327), despite the stricter ESG filter.

Overall, the results suggest that ESG integration through percentile-based filtering preserves portfolio efficiency, especially when using strategies that explicitly target risk-adjusted returns (Max. SR, Max. SoR, M-CVaR). While there is some trade-off between ESG strictness and diversification, the performance remains robust even at the 80th percentile.

Table 5. Optimal portfolios with ESG constraints by taking 80th percentile

NDX 100 -					
	Mean	**Volatility**	**CVaR**	**Sharpe Ratio**	**Sortino Ratio**
MV	0.0094	0.0321	-0.067	0.2934	0.5368
Max. SR	0.0188	0.0414	-0.079	0.4549	0.9199
Max. SoR	0.0188	0.0425	-0.072	0.4427	0.9701
CVaR	0.014	0.0392	-0.054	0.3575	0.8267
M-CVaR	0.0186	0.0453	-0.063	0.4096	0.964
SX 50E					
	Mean	**Volatility**	**CVaR**	**Sharpe Ratio**	**Sortino Ratio**
MV	0.0115	0.0382	-0.07	0.3015	0.6091
Max. SR	0.0146	0.0422	-0.069	0.3464	0.7379
Max. SoR	0.0146	0.0424	-0.07	0.3445	0.7421
CVaR	0.0137	0.0444	-0.064	0.3081	0.6769
M-CVaR	0.0162	0.0485	-0.07	0.3349	0.7327

Source: own elaboration.

Finally, Figure 2 illustrates the shift of the efficient frontiers (EFs) as higher ESG criteria are applied—specifically by raising the selection threshold from the 60th to the 70th and 80th percentiles. The analysis is performed separately for the NDX100 and SX50E indices, and for three optimization frameworks: Mean-Variance, Mean-Semivariance, and M-CVaR. As expected, increasing the ESG threshold reduces the number of eligible companies, concentrating the portfolio in ESG leaders. This narrowing of the investment universe leads to:

i. Reduced diversification: As the pool of assets decreases, the opportunity to balance risk across sectors and risk profiles diminishes.
ii. Higher portfolio risk: The EFs shift downward and to the right in the portfolio models, reflecting higher volatility and tail risk for a given level of expected return.

These results highlight a critical trade-off between ESG integration and portfolio efficiency. Portfolios constrained at the 60th percentile offer a better balance of ESG quality and diversification. However, at higher percentiles, although the ESG quality of the portfolio improves, the concentration in a few leaders reduces the ability to mitigate risk, shifting the frontier unfavourably.

Figure 2. Efficient-Frontiers for the portfolio models

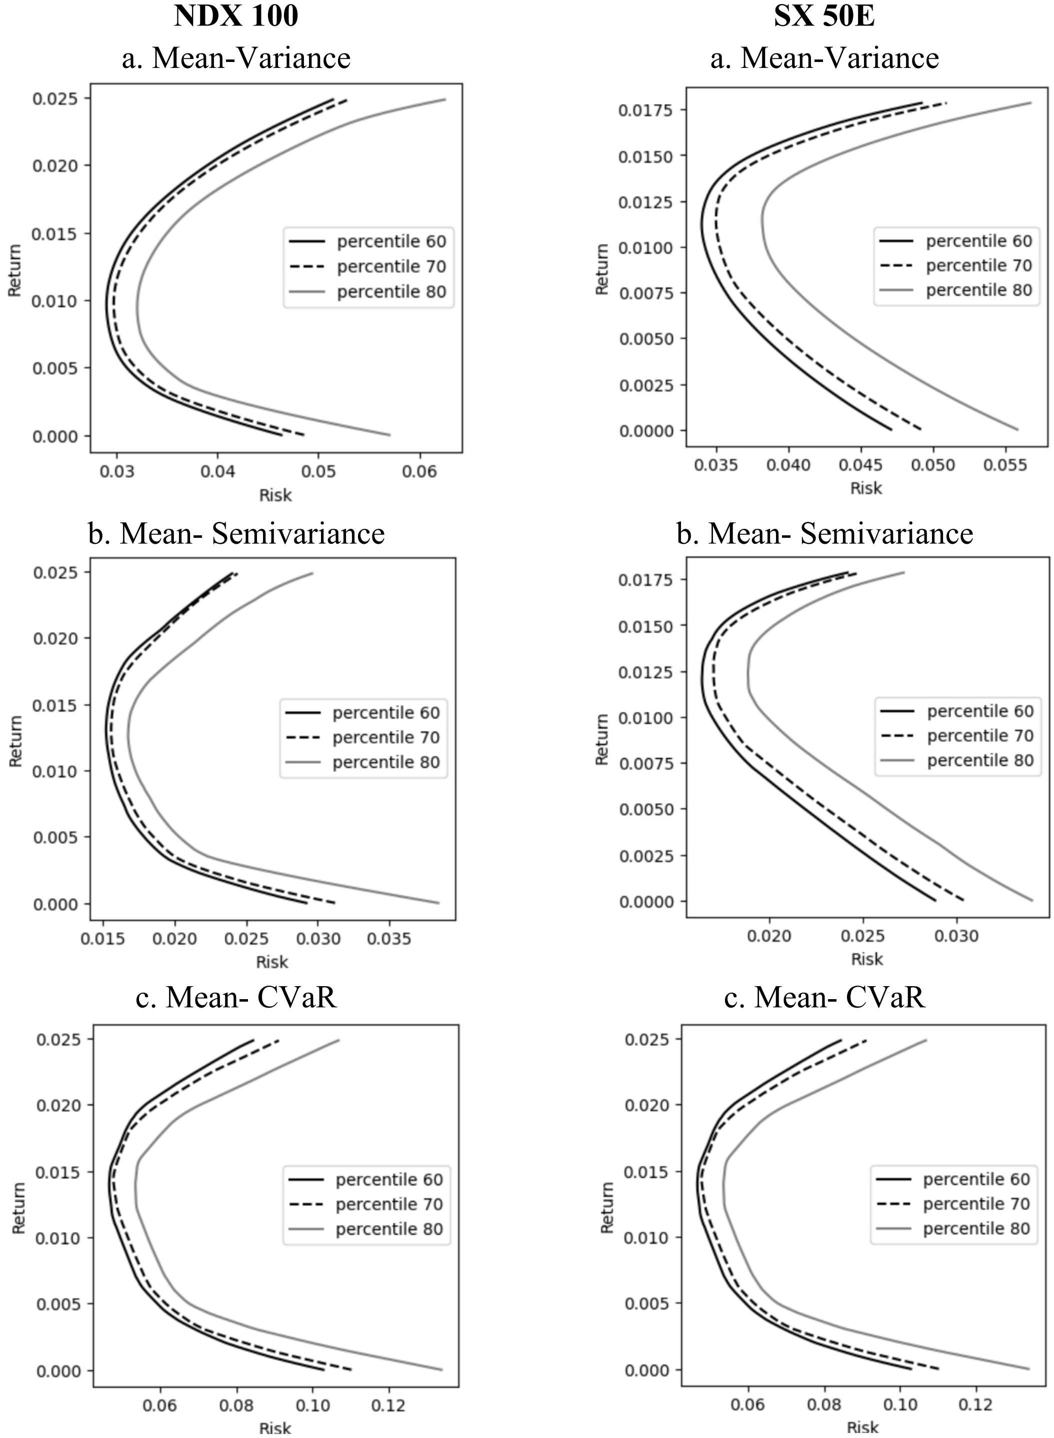

Source: own elaboration.

4. Conclusions

This research offers a comprehensive analysis of optimal portfolio construction by incorporating downside risk measures and ESG criteria within European equity markets. Using a comparative framework across Mean-Variance, Mean-Semivariance (Sortino), CVaR, and M-CVaR optimization models, and applying ESG constraints based on percentile thresholds, we evaluate their impact on portfolio performance and efficient frontiers using the SX 50E and NDX 100 indices as benchmarks.

The findings reveal that while downside risk measures such as semivariance and CVaR provide better risk control for loss-averse investors, the integration of ESG constraints—particularly at higher thresholds (e.g., 80th percentile)—results in reduced diversification, higher volatility, and ultimately lower risk-adjusted returns. Specifically, ESG portfolios tend to exhibit inferior Sharpe and Sortino ratios compared to their unconstrained counterparts, highlighting a trade-off between sustainability goals and financial performance. The sensitivity analysis confirms that portfolios under higher ESG scores become concentrated in a small number of leading companies, which increases idiosyncratic risk and undermines the benefits of diversification. In this context, moderate ESG constraints (e.g., 60th or 70th percentile) may offer a more balanced approach, preserving the sustainability profile without severely compromising portfolio efficiency.

By systematically comparing downside risk models within an ESG-integrated framework, this study contributes to the growing academic discourse on sustainable investing and offers relevant insights for institutional investors navigating the tensions between ethical considerations and return optimization. Future work could advance this analysis in several directions. First, incorporating higher-order moments, such as skewness and kurtosis, into the optimization process (e.g., MVSK models) would provide a more robust treatment of asymmetric and fat-tailed return distributions, especially relevant for concentrated ESG portfolios. Second, leveraging advanced techniques, such as robust or Bayesian optimization (e.g., Meucci's entropy pooling), and machine learning models, including autoencoders, variational inference, or Kolmogorov-Arnold Networks (KAN), could help uncover non-linear patterns in ESG-risk relationships and improve out-of-sample performance. Finally, testing the robustness of the results using alternative ESG data providers (e.g., Refinitiv, Sustainalytics, MSCI) and accounting for transaction costs and turnover will enhance the practical applicability of these findings for long-term, sustainability-oriented investors.

References

Abate, G., Di Giacomo, M., & Ferraris, A. (2021). Downside risk and ESG investing: A multi-objective portfolio optimization approach. Journal of Sustainable Finance & Investment, 11(3), 213–230.

Acerbi, C., & Tasche, D. (2002). On the coherence of expected shortfall. Journal of Banking & Finance, 26(7), 1487–1503.

Alessandrini, F., & Jondeau, E. (2020). ESG investing: From sin stocks to smart beta. Journal of Portfolio Management, 46(3), 59–74.

Chen, L., & Mussalli, G. (2020). An integrated approach to ESG investing: ESG integration in equity portfolios. Journal of Investment Management, 18(2), 23–40.

Chen, L., Mussalli, G., & Wang, J. (2021). ESG investing through a factor lens. Journal of Portfolio Management, 47(6), 89–103.

Chen, Y., Yang, Y., & Zhang, Y. (2023). ESG performance evaluation and portfolio optimization: A cross-efficiency DEA approach. European Journal of Operational Research, 304(2), 567–580.

Coqueret, G. (2022). ESG investing and performance: A review of the empirical literature. Finance Research Letters, 45, 102133.

Gasser, S. M., Rammerstorfer, M., & Weinmayer, K. (2017). Markowitz revisited: Social portfolio engineering. European Journal of Operational Research, 258(3), 1181–1190.

Hirschberger, M., Steuer, R. E., Utz, S., & Wimmer, M. (2013). A portfolio optimization approach for socially responsible investment. European Journal of Operational Research, 234(2), 545–559.

Markowitz, H. (1952). Portfolio selection. The Journal of Finance, 7(1), 77–91.

Pedersen, L. H., Fitzgibbons, S., & Pomorski, L. (2021). Responsible investing: The ESG-efficient frontier. Journal of Financial Economics, 142(2), 572–597.

Rockafellar, R. T., & Uryasev, S. (2000). Optimization of conditional value-at-risk. Journal of Risk, 2(3), 21–41.

Sortino, F. A., & Price, L. N. (1994). Performance measurement in a downside risk framework. Journal of Investing, 3(3), 59–64.

Steuer, R. E., & Utz, S. (2023). Mapping the ESG-efficient frontier: A three-dimensional approach. European Journal of Operational Research, 304(1), 1–14.

Utz, S., Wimmer, M., & Hirschberger, M. (2014). Tri-criterion modeling for constructing socially responsible portfolios. European Journal of Operational Research, 234(2), 518–527.

Xidonas, P., & Essner, R. (2022). Multi-objective ESG portfolio optimization: A goal programming approach. Annals of Operations Research, 312(1), 123–145.