University of Economics in Katowice

Journal of Economics and Management

Vol. 48 • 2026

Minh Duc Nguyen
https://orcid.org/0000-0003-4702-0888

Department of Economic Information Systems
University of Economics, Hue University
Hue, Vietnam
nguyenminhduc@hueuni.edu.vn

Toward responsible investment: A deep learning frame-work for ESG-differentiated portfolio optimization

Accepted by Editor Ewa W. Ziemba | Received: March 3, 2025 | Revised: August 26, 2025; October 17, 2025 | Accepted: December 4, 2025 | Published: January 22, 2026.

Abstract

Aim/purpose – This paper proposes a framework that integrates deep learning-based return forecasting with environmental, social, and governance (ESG)-differentiated optimization to align portfolio performance with financial and sustainability goals, enabling data-driven responsible investment decisions. The study hypothesizes that ESG risk dimensions influence portfolio performance differently: mitigating environmental risk imposes higher financial costs due to regulatory and operational pressures, whereas social and governance risks yield more balanced return-sustainability trade-offs.
Design/methodology/approach – This study employs the N-BEATS deep learning model to forecast one-day-ahead returns for S&P 100 stocks. The predicted returns serve as inputs to an enhanced Mean-Variance with Forecasting (MVF) model that integrates ESG risk as a penalty term. ESG factors are analyzed both collectively and across individual dimensions using a tunable risk-aversion parameter that reflects investor preferences. The dataset includes 99 S&P 100 stocks from January 2017 to December 2024, with distinct training, validation, and test sets for model development and evaluation.
Findings – The study reveals that incorporating ESG risk into portfolio optimization with forecasted returns produces distinct trade-offs across ESG dimensions. Mitigating environmental risk entails the greatest return cost, whereas social and governance risks allow more favorable balances between return and sustainability. The N-BEATS model achieves sufficient forecasting accuracy to inform investment decisions. Moreover, the elbow point method offers a practical means for selecting optimal ESG sensitivity levels, enabling investors to effectively balance performance and sustainability objectives.

Cite as: Nguyen, M. D. (2026). Toward responsible investment: A deep learning framework for ESG-differentiated portfolio optimization. *Journal of Economics and Management*, *48*, 1-16. https://doi.org/10.22367/jem.2026.48.01

Research implications/limitations – This research demonstrates that combining deep learning-based forecasting with ESG-differentiated optimization enables more nuanced and responsible investment strategies, offering practical tools for aligning financial and sustainability goals. However, the study is limited by its focus on a single market (S&P 100) and does not account for real-world factors such as transaction costs or dynamic rebalancing, which could affect practical applicability.
Originality/value/contribution – This study is among the first to integrate N-BEATS-based return forecasting with ESG-differentiated portfolio optimization in a unified framework. It offers a novel approach that enables investors to explicitly manage trade-offs between financial returns and individual ESG risk dimensions, providing both methodological innovation and practical guidance for responsible investing.

Keywords: responsible investment, ESG risk, portfolio optimization, deep learning, N-BEATS.
JEL Classification: G1, C5, Q5, M1.

1. Introduction

In today's increasingly dynamic and uncertain financial landscape, investors face the dual challenge of achieving strong portfolio performance while simultaneously managing long-term risk exposures and aligning with broader sustainability goals. Global capital markets are influenced not only by traditional macroeconomic indicators but also by geopolitical instability, climate-related disruptions, technological innovation, and regulatory changes (Abuzov, 2023; Hodula et al., 2024; Krueger et al., 2020). In this context, two fundamental questions persist at the core of modern portfolio management: How can future stock returns be predicted with sufficient accuracy? Furthermore, how can these predictions be effectively translated into optimal investment decisions?

For decades, the classical framework of mean-variance optimization, introduced by Markowitz (1952), has served as a cornerstone of modern portfolio theory. The approach constructs an efficient frontier by selecting portfolios that minimize variance for a given level of expected return, based on historical estimates of asset mean returns and covariances. While Markowitz's model does not explicitly assume that asset returns follow stationary distributions, it implicitly relies on the stability of these statistical parameters over time to ensure meaningful and effective optimization. In practice, this assumption has been increasingly questioned, as empirical research has shown that financial time series often exhibit non-stationarity, structural breaks, and regime-switching behavior (Ang & Bekaert, 2015; Cont, 2001). Additionally, the traditional mean-variance framework does not incorporate forward-looking or predictive information, such as analyst revisions, momentum indicators, or signals generated by machine learning models. As a result, portfolios constructed solely on historical data may fail to respond to rapid market shifts or exploit short-term return predictability. This limitation becomes particularly salient during periods of financial turbulence (such as the 2008 global financial crisis or the COVID-19 pandemic) when the statistical properties of asset returns may deviate significantly from historical norms (Naik & Reddy, 2021; Yahya, 2023).

To overcome these limitations, researchers have explored enhancements that integrate return prediction into the portfolio optimization process (Ma et al., 2021; Yu et al., 2020). Notably, the rise of deep learning has provided powerful tools for modeling the nonlinear and temporal structures that characterize financial time series. Models such as convolutional neural networks (CNNs) and long short-term memory (LSTM) networks have shown promise in this space, with CNNs excelling at extracting localized temporal features and LSTMs capturing long-term dependencies (Kim & Kim, 2019; Mehtab & Sen, 2020). In addition to these recurrent and convolutional models, more recent advances have introduced feedforward architectures such as N-BEATS (Neural Basis Expansion Analysis for Time Series), which rely on residual and interpretable basis blocks to achieve state-of-the-art performance in univariate time-series forecasting (Oreshkin et al., 2019). N-BEATS has proven particularly effective for short-horizon financial forecasting, as it can model complex patterns without exogenous features or recurrent dependencies. These developments have collectively encouraged the integration of forecasted returns into quantitative strategies, resulting in more adaptive and forward-looking portfolio models.

However, financial return alone is no longer the sole concern for modern investors. The growing emphasis on environmental, social, and governance (ESG) criteria reflects a paradigm shift in investment priorities. ESG integration seeks to capture the long-term risks and opportunities associated with corporate behavior, regulatory exposure, and reputational outcomes. Numerous studies have shown that ESG-aware portfolios may offer improved resilience and reduced tail risks, particularly in turbulent market environments (Fatemi et al., 2015; Friede et al., 2015). Consequently, integrating ESG factors into portfolio optimization has become a significant area of research, with various computational approaches being explored.

Machine learning and deep learning techniques are increasingly employed to handle the complexities involved in creating ESG-aware portfolios. Recent literature showcases diverse methodologies: some studies utilize Multivariate Bidirectional LSTMs (MBi-LSTMs) coupled with reinforcement learning or genetic algorithms (NSGA) to predict returns and dynamically optimize ESG portfolios (Gaurav et al., 2025; Vo et al., 2019). Others combine clustering techniques, such as Agglomerative Hierarchical Clustering (AHC) or Robust Expectation Maximization (REM), with Multi-Criteria Decision-Making (MCDM) methods like MARCOS or VIKOR for ESG-integrated stock selection, often followed by return forecasting using CNNs or ensemble models (e.g., GRU--BiLSTM-XGBoost) and metaheuristic optimization like Jaya-TLBO (Jain et al., 2025; Sajadi et al., 2025). Furthermore, hybrid approaches, such as combining Multi-Objective Particle Swarm Optimizers (MOPSO) with LSTMs, have been developed to tackle bi-objective problems involving ESG volatility and portfolio diversification (Aprea et al., 2025).

While these studies demonstrate the successful application of various machine learning techniques to ESG-aware portfolio construction, there remains a need for frameworks that specifically leverage newer deep learning architectures primarily designed for time-series forecasting, such as N-BEATS, and integrate these predictions into optimization models, enabling fine-grained control over individual ESG risk dimensions. Existing approaches often integrate ESG as a single score or constraint, potentially masking distinct trade-offs associated with each environmental, social, and governance dimension. Building on these motivations, this study proposes the following hypotheses:

H1: Integrating ESG risk into portfolio optimization leads to measurable differences in return-risk trade-offs across ESG dimensions.

H2: Among these dimensions, mitigating environmental risk results in a greater reduction in portfolio returns than mitigating social or governance risks, due to higher exposure to physical and regulatory constraints.

These hypotheses guided the empirical analysis and provided a clear foundation for comparing the financial implications of sustainability-oriented investment decisions.

To examine these hypotheses, the study proposed a novel ESG-differentiated portfolio optimization framework that integrated deep learning forecasts into a modified Mean-Variance with Forecasting (MVF) model. Specifically, we adopted the N-BEATS model, a feedforward deep learning architecture designed explicitly for univariate time-series forecasting. N-BEATS offers a high level of flexibility and interpretability, making it well-suited for daily asset return prediction without requiring complex input features. The predicted returns generated by N-BEATS were then incorporated into an extended MVF formulation, which also included a reward for forecasting accuracy and a penalty term for ESG risk exposure. This ESG term, modeled either as total ESG risk or as its environmental, social, or governance components, is governed by a tunable risk-aversion parameter that reflects investors' sustainability preferences. To support balanced decision-making, we applied the Kneedle algorithm to identify the elbow point on the return-risk curve, offering a practical, data-driven method for selecting the optimal trade-off between financial performance and ESG exposure. Through this integrated approach, the paper contributes both methodological innovation and practical guidance for investors aiming to align return-driven strategies with sustainability objectives.

2. Deep learning-based portfolio optimization framework

2.1. Forecasting stock return with N-BEATS

To forecast daily stock returns, we adopt N-BEATS, a deep feedforward architecture introduced by Oreshkin et al. (2019) for univariate time-series forecasting. Separate models are trained for each stock to capture firm-specific temporal dynamics while maintaining flexibility and interpretability. Although multivariate frameworks could exploit cross-asset dependencies, they often require extensive data preprocessing and risk overfitting when applied to large equity sets such as the S&P 100. Moreover, cross-stock correlations are explicitly handled later in the portfolio optimization stage through the covariance structure of the MVF model. N-BEATS operates without requiring feature engineering, domain--specific knowledge, or exogenous variables, making it highly suitable for financial forecasting tasks. N-BEATS operates without feature engineering, domain--specific knowledge, or exogenous variables, making it well-suited for financial forecasting tasks. Its strong generalization capabilities also make it well-suited for modeling noisy, nonlinear return sequences in equity markets. The overall architecture of N-BEATS is illustrated in Figure 1.

The model consists of multiple stacks, each containing a series of blocks. Every block receives a lookback window of past returns. It generates two outputs: a backcast, which estimates and removes components already captured from the input, and a forecast, which contributes to the overall prediction. Internally, each block includes a fully connected (FC) stack of dense layers followed by basis expansion functions that output interpretable components. The residuals from each block are propagated downstream to the next block, allowing successive layers to refine the forecast. Finally, outputs from all blocks across all stacks are aggregated to form the global forecast. This residual and additive structure enables the model to hierarchically learn time-dependent features.

Figure 1. N-BEATS architecture for time-series forecasting

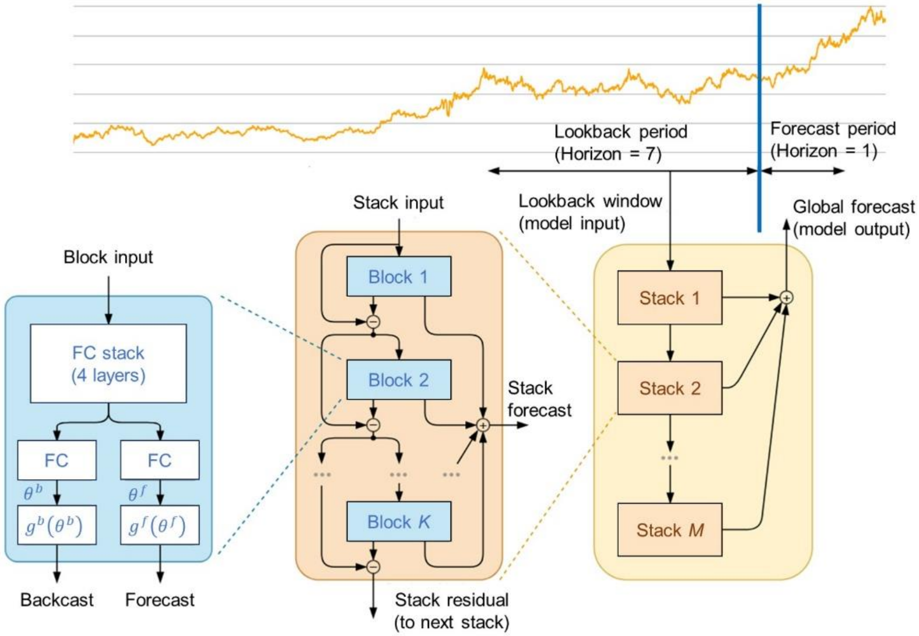

Source: Adapted from Oreshkin et al. (2019).

The modular and interpretable design of N-BEATS enables it to effectively decompose and learn both long-term trends and short-term fluctuations in financial time series. In this study, the model is configured in its generic form to use a sliding window of the previous seven trading days as input to predict the return for the following day. This horizon was chosen after testing multiple lookback lengths (5, 7, 10, and 14 days), with the seven-day window providing the best balance between capturing short-term temporal dependencies and preventing overfitting. The one-day-ahead forecasting horizon was retained based on empirical validation and prior evidence that such short-term horizons yield superior predictive accuracy in financial time-series forecasting (Yu et al., 2020). A separate N-BEATS model is trained for each individual stock, enabling the model to capture stock-specific temporal dynamics without interference from other series. Each model is trained for 100 epochs, providing sufficient learning capacity while avoiding overfitting. These one-day-ahead forecasts serve as critical inputs to the ESG-aware portfolio optimization process described in Section 2.2.

2.2. ESG-differentiated MVF framework

To support responsible investment while leveraging predictive modeling, we extend the MVF model (Yu et al., 2020) by incorporating both forecast-based return modeling and ESG risk considerations. This leads to an enhanced MVF formulation, which simultaneously accounts for expected return, predictive reliability, and sustainability exposure.

Let x_i denote the portfolio weight of asset i, σ_{ij} denote the covariance between assets i and j, $\hat{r}_i$ denote the predicted return, $\bar{\varepsilon}_i$ denote the average forecasting error, and ESG_i denote the ESG risk score of asset i with γ controlling the degree of ESG concern. The optimization objective is defined as:

$$min\left(\sum_{i=1}^{n}\sum_{j=1}^{n} x_i x_j \sigma_{ij} - \sum_{i=1}^{n} x_i \hat{r}_i - \sum_{i=1}^{n} x_i \bar{\varepsilon}_i + \gamma \sum_{i=1}^{n} x_i ESG_i\right) \tag{1}$$

constraint to:

$$\sum_{i=1}^{n} x_i = 1 \; and \; 0 \le x_i \le 1 \; for\ all\ i \tag{2}$$

The ESG risk component can represent either the total ESG risk score or any of its individual components: environmental (E), social (S), or governance (G). These values are normalized prior to optimization to maintain scale consistency across dimensions and assets. By varying the risk-aversion parameter γ, investors can explore a continuum of trade-offs among financial return, ESG exposure, and forecast reliability. When $\gamma = 0$, the model reduces to a forecasting-based MVF formulation without ESG concerns. As γ increases, the optimization increasingly favors assets with lower ESG risk scores, facilitating ESG--compliant portfolio construction. This flexible formulation supports personalized investment strategies aligned with diverse sustainability priorities.

3. Experiment

This section presents the experimental framework used to evaluate the proposed deep learning-based ESG-differentiated portfolio optimization approach. We describe the dataset construction, preprocessing steps, forecasting procedures using the N-BEATS model, and the integration of predicted returns into the extended MVF optimization model. The experiments are designed to assess both the deep learning model's forecasting accuracy and the portfolio outcomes across different ESG risk dimensions. The following provides detailed explanations of the experiment setup and the empirical results.

3.1. Experiment setup

To empirically evaluate the proposed framework, we constructed a dataset based on stocks from the S&P 100 index, spanning the period from January 2017 to December 2024. Daily closing prices and corresponding ESG risk scores were sourced from Yahoo Finance, which aggregates ESG data from providers such as Sustainalytics. The initial universe comprised 101 stocks; however, two were excluded during preprocessing (one due to missing price data and another due to the absence of ESG scores), resulting in a final sample of 99 stocks for analysis.

The ESG risk scores obtained represent numeric evaluations of a firm's exposure to environmental, social, and governance-related risks. These scores range from 0 to 100, with lower values indicating stronger ESG performance and reduced sustainability-related risk. For interpretation, ESG scores are typically categorized as follows: 0-10: negligible risk, 10-20: low risk, 20-30: medium risk, 30-40: high risk, and 40+: severe risk. These scores were directly integrated into the portfolio optimization process as risk-penalty terms to evaluate trade-offs between financial return and sustainability exposure. Table 1 provides a statistical summary of the ESG risk scores across the 99 selected stocks used in the analysis. The average total ESG risk score is 21.65, placing most firms in the medium-risk category, with some variation reflected by a standard deviation of 6.89. Environmental scores have the lowest mean (5.04), suggesting generally low environmental risk, while social and governance scores average 10.39 and 6.22, respectively. These descriptive statistics highlight the differing risk profiles across ESG dimensions, supporting the rationale for ESG-differentiated portfolio strategies.

Table 1. Descriptive statistics of ESG risk scores for selected stocks

Statistic	Total ESG score	Environmental score	Social score	Governance score
Count	99	99	99	99
Mean	21.65	5.04	10.39	6.22
Std	6.89	4.81	3.64	2.74
Min	8.52	0.07	1.40	2.09
25%	17.01	1.60	8.42	4.66
50%	20.42	3.04	10.13	5.65
75%	25.14	7.80	12.82	6.98
Max	43.66	25.09	22.24	19.02

Source: Author's own study.

To prepare the dataset for return forecasting, daily price data were first converted into return series, capturing the relative change between consecutive trading days. To enhance training stability and ensure comparability across stocks, each return series was standardized using the formula:

$$\hat{d}_i = \frac{d_i - \mu}{\sigma} \tag{3}$$

where d_i is the return at time i, and μ, σ represent the mean and standard deviation of that stock's return series during the training period.

For temporal modeling, we employed a sliding window approach, transforming each stock's return sequence into overlapping input-output pairs. Each input window comprised seven consecutive trading days and was used to predict the return on the following day. This setup allows the model to learn short-term temporal dependencies relevant for daily return forecasting.

The dataset was split chronologically into three non-overlapping subsets:

- Training set: January 2017 – December 2021,
- Validation set: January 2022 – December 2022,
- Test set: January 2023 – December 2024.

The N-BEATS model, introduced earlier, was trained on the return sequences from the training set, with hyperparameter tuning performed using the validation set. Predictions from the test set were then used as input to the MVF model, in which ESG risks, including the total score and individual environmental, social, and governance components, were incorporated as additive penalty terms. This structure enabled analysis of ESG-aware optimization under realistic, data-driven forecasts. All preprocessing, modeling, and optimization tasks were implemented in Python, with the N-BEATS model developed using the Keras framework.

3.2. Results and discussion

3.2.1. Forecasting performance

To assess the predictive accuracy of the N-BEATS model, we evaluated its performance using a set of both magnitude-based and direction-based error metrics. The magnitude-based metrics included Mean Absolute Error (MAE), Mean Squared Error (MSE), and Root Mean Squared Error (RMSE), which quantify the average discrepancy between the actual and predicted stock returns. In addition, directional accuracy metrics were employed to examine the model's ability to predict the correct movement direction of returns. These included the overall hit ratio H_R, the upward movement hit ratio H_{R+}, and the downward movement hit ratio H_{R-}. Table 2 summarizes the formulas, mean values, and standard deviations for each metric across the test set, providing a comprehensive evaluation of the N-BEATS model's forecasting performance.

Table 2. Evaluating the forecasting performance of the N-BEATS model

Metric	Formula	Mean	Std.
MAE	$\frac{1}{n}\sum_{i=1}^{n}\lvert y_i - \hat{y}_i\rvert$	1.65×10^{-2}	5.10×10^{-3}
MSE	$\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2$	5.64×10^{-4}	3.94×10^{-4}
RMSE	$\sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}$	2.27×10^{-2}	7.08×10^{-3}
H_R	$\frac{count(y_i \cdot \hat{y}_i > 0)}{count(y_i \cdot \hat{y}_i \neq 0)}$	50.85%	2.36×10^{-2}
H_{R+}	$\frac{count(y_i > 0 \; and \; \hat{y}_i > 0)}{count(\hat{y}_i > 0)}$	52.84%	3.40×10^{-2}
H_{R-}	$\frac{count(y_i < 0 \; and \; \hat{y}_i < 0)}{count(\hat{y}_i < 0)}$	48.01%	3.64×10^{-2}

Note: y_i is the actual return, $\hat{y}_i$ is the predicted return, and n is the number of predictions.

Source: Author's own study.

Specifically, the magnitude-based metrics indicated low average errors, with an MAE of 1.65×10^{-2} and an RMSE of 2.27×10^{-2}, suggesting that the model produces relatively accurate return forecasts. The direction-based metrics showed that the model correctly predicts the movement direction of returns approximately 50.85% of the time overall. Specifically, it achieved a 52.84% hit rate for upward movements and a 48.01% hit rate for downward movements. These results indicated that the N-BEATS model achieved a level of forecasting accuracy sufficient to support its use in portfolio optimization. Moreover, this performance was consistent with results reported in prior studies on financial time-series forecasting using N-BEATS and related deep learning architectures (Kim & Kim, 2019; Ma et al., 2021), confirming the model's reliability as a forecasting backbone for the subsequent portfolio optimization process. Accordingly, we proceeded to evaluate how these forecasts influence asset allocation under different ESG risk types in the next section.

3.2.2. Portfolio optimization performance under different ESG risk types

To examine the practical value of our forecasting-integrated portfolio optimization framework, we assess how portfolios perform when different ESG risk dimensions are incorporated into the objective function. Using the predicted stock returns from the N-BEATS model, we applied the MVF model, in which ESG risk is treated as an additional penalty term controlled by the risk-aversion parameter γ. The optimization was performed under the assumption of no transaction fees and with a rebalancing frequency of every 20 trading days, following the empirical setup of Yu et al. (2020). This setting allowed us to isolate the effects of ESG risk aversion on portfolio outcomes while maintaining a realistic trading horizon that balances responsiveness and transaction stability. Furthermore, to capture an effective trade-off between portfolio return and different types of risk, we identified the elbow point on each risk curve using the KneeLocator algorithm (Satopaa et al., 2011). This algorithm detects the point of maximum curvature – that is, where the marginal improvement in risk reduction becomes offset by a disproportionate decline in return. This point represents the stage beyond which further reductions in environmental, social, governance, or total ESG risk led to disproportionately greater sacrifices in return, offering limited overall benefit to portfolio quality.

Figure 2 illustrates the impact of varying ESG risk aversion levels on portfolio return, ESG risk, maximum drawdown (MDD), and turnover rate across four scenarios: total ESG risk, environmental risk, social risk, and governance risk. In each subfigure, the left vertical axis represents portfolio return (%), while the right vertical axis represents the respective ESG risk value. The horizontal axis shows the risk aversion parameter γ, which ranges from 0 (no ESG concern) to 2.0 (high ESG concern). Blue curves represent portfolio returns, and orange curves represent ESG risks.

In Figure 2(a), which depicts the optimization results under total ESG risk, the portfolio returns decreases steadily as the ESG risk aversion parameter γ increases. At low levels of γ, returns remain relatively high but are associated with elevated ESG risk. As γ approaches 0.37 (the identified elbow point), there is a noticeable drop in ESG risk from above 20 to around 15, with only a moderate decline in return to 49.09%. At this optimal trade-off point, the portfolio exhibits a maximum drawdown of 65.48% and a turnover rate of 51.60%. Beyond this point, further reductions in ESG risk yield increasingly smaller benefits, while portfolio returns fall more sharply, indicating diminishing returns to sustainability improvement.

Figure 2. Performance of portfolio optimization under different ESG risk types

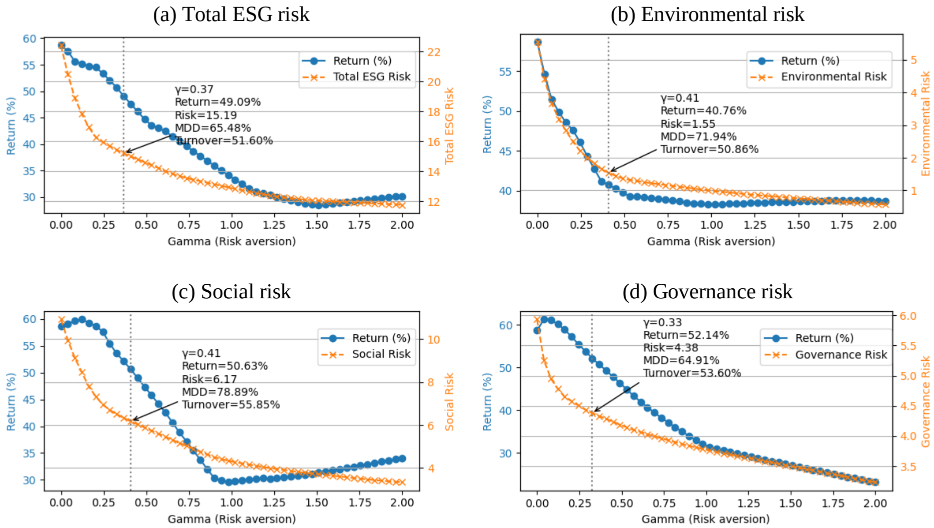

Source: Author's own study.

In Figure 2(b), under environmental risk, a similar trade-off pattern emerges, though with steeper consequences for return. The return curve drops more rapidly as γ increases, reaching 40.76% at the elbow point, $\gamma = 0.41$, where environmental risk is reduced to 1.55. At this point, the portfolio exhibits a maximum drawdown of 71.94% and a turnover rate of 50.86%. Notably, environmental risk is highly sensitive to the risk aversion parameter, showing significant improvement even at moderate γ levels. However, the financial cost of doing so is also more pronounced than in other ESG dimensions.

Figure 2(c) presents the trade-off under social risk. Here, portfolio returns remain relatively stable above 50% across a broader range of γ values. The elbow point is also found at $\gamma = 0.41$, where the return is 50.63% and social risk is reduced to 6.17. At this point, the portfolio experiences a maximum drawdown of 78.89% and a turnover rate of 55.85%. Compared to environmental risk, the return penalty for reducing social risk is milder, suggesting that social considerations may be incorporated into the optimization with less disruption to performance.

In Figure 2(d), focused on governance risk, the return curve is the flattest of all, peaking at 52.14% when $\gamma = 0.33$. Governance risk also shows a consistent downward trend as γ increases, dropping to 4.38 at the elbow. At this optimal point, the portfolio records a maximum drawdown of 64.91% and a turnover rate of 53.60%. This subfigure indicates that governance risk can be reduced efficiently without a significant trade-off in return, making it an appealing ESG dimension for investors seeking both performance and responsibility.

Overall, Figure 2 demonstrates that the impact of ESG risk aversion on portfolio performance is not uniform across ESG dimensions. Environmental risk reduction typically entails higher financial costs, whereas governance and social risks can be mitigated with more favorable trade-offs. To statistically validate these observations, Mann–Whitney tests were conducted between each ESG-integrated portfolio (evaluated at its respective elbow point – optimal γ) and the non-ESG baseline. All comparisons yielded significant results (p-value $\approx$ 0.004), confirming that ESG integration leads to statistically meaningful differences in portfolio outcomes. These insights reinforce the value of ESG--differentiated optimization in aligning investment strategies with sustainability priorities.

In terms of portfolio stability, the MDD analysis highlights notable differences among the ESG dimensions. The governance-focused portfolio demonstrates the lowest MDD (64.91%), reflecting stronger downside protection and greater resilience to market volatility. Conversely, the social-focused portfolio records the highest MDD (78.89%), indicating heightened vulnerability to adverse market movements despite delivering competitive returns. The environmental and total ESG portfolios experience moderate drawdowns, suggesting a balanced risk-return trade-off. Moreover, turnover rates remain relatively stable across all portfolios (ranging from 50% to 56%), implying that ESG integration maintains consistent portfolio management behavior. Overall, these findings suggest that ESG integration reshapes the portfolio's risk-return dynamics while maintaining management efficiency.

4. Conclusions

This study proposes a novel framework that integrates deep-learning-based stock-return forecasting with ESG-differentiated portfolio optimization to support responsible investment decision-making. By leveraging the predictive capabilities of the N-BEATS model, we generate return forecasts that serve as inputs to an MVF model, which incorporates ESG risks as penalty terms. Our empirical evaluation demonstrates that the impact of ESG risk aversion on portfolio performance is not uniform across ESG dimensions. In particular, reducing environmental risk tends to incur higher return sacrifices, while social and governance risks can be mitigated with more favorable trade-offs. These insights contribute to the literature on ESG-aware quantitative investing by providing empirical evidence of how deep learning can be aligned with sustainability objectives.

Compared with prior research that uses screening methods or aggregates ESG into a single score, this study advances the field by directly embedding differentiated environmental, social, and governance risk-aversion components into the optimization model's objective function. This integration enables a more granular, explicit control of individual sustainability objectives, offering a methodological bridge between data-driven forecasting and multi-objective decision-making. In practice, the results guide investors in identifying the “elbow points” on return-risk curves that balance profit objectives with ethical preferences, enabling flexible adaptation of ESG sensitivity in portfolio design. Institutional investors, in particular, can apply this framework to calibrate ESG constraints according to regional regulations, sustainability mandates, or market liquidity conditions.

Nonetheless, the study's assumptions impose several limitations that should be acknowledged. The analysis assumes a frictionless market environment (e.g., excluding transaction costs or liquidity constraints) and focuses exclusively on stocks from the S&P 100 index. These simplifications enhance tractability but limit the model's real-world applicability. Incorporating transaction costs and liquidity effects would likely reduce overall returns, particularly for portfolios with high turnover rates, while implementing dynamic rebalancing could improve adaptability to market fluctuations at the expense of higher operational costs. Similarly, restricting the analysis to a single U.S. index constrains the generalizability of the findings; performance outcomes may vary under different market structures, regulations, and ESG disclosure standards.

Future research could extend this work by (i) evaluating multi-market and cross-country datasets to test model generalizability, (ii) incorporating transaction costs and dynamic rebalancing to simulate realistic trading environments, (iii) experimenting with alternative deep learning architectures or hybrid ensemble forecasting models, and (iv) analyzing longer forecast horizons to capture dynamic ESG effects over time. Such extensions would enhance the practical relevance of this framework and further align quantitative portfolio optimization with the principles of sustainable finance.

Funding

This research received no external funding.

Conflict of interest

The author declares that there is no conflict of interest.

Ethics statement

This study did not involve human participants and did not require ethical approval.

Data availability statement

The data supporting the findings of this study are available from the corresponding author upon reasonable request.

References

Abuzov, A. (2023). The role of technological innovations in institutional regulation of the financial capital market. *E3S Web of Conferences, 376*, 5047. https://doi.org/10.1051/e3sconf/202337605047

Ang, A., & Bekaert, G. (2015). International asset allocation with regime shifts. *The Review of Financial Studies, 15*(4), 1137-1187. https://doi.org/10.1093/rfs/15.4.1137

Aprea, I. L., Bosi, G., Sbaiz, G., & Scognamiglio, S. (2025). Bi-objective portfolio optimization under ESG volatility via a MOPSO-deep learning algorithm. *Mathematics, 13*(20), 3308. https://doi.org/10.3390/math13203308

Cont, R. (2001). Empirical properties of asset returns: Stylized facts and statistical issues. *Quantitative Finance, 1*(2), 223-236. https://doi.org/10.1088/1469-7688/1/2/304

Fatemi, A., Fooladi, I., & Tehranian, H. (2015). Valuation effects of corporate social responsibility. *Journal of Banking & Finance, 59*, 182-192. https://doi.org/10.1016/j.jbankfin.2015.04.028

Friede, G., Busch, T., & Bassen, A. (2015). ESG and financial performance: Aggregated evidence from more than 2000 empirical studies. *Journal of Sustainable Finance & Investment, 5*(4), 210-233. https://doi.org/10.1080/20430795.2015.1118917

Gaurav, A., Baishnab, K., & Singh, P. K. (2025). Intelligent ESG portfolio optimization: A multi-objective AI-driven framework for sustainable investments in the Indian stock market. *Sustainable Futures, 9*, 100832. https://doi.org/10.1016/j.sftr.2025.100832

Hodula, M., Janků, J., Malovaná, S., & Ngo, N. A. (2024). *Geopolitical risks and their impact on global macro-financial stability: Literature and measurements* (BOFIT Discussion Paper, No. 9). https://doi.org/10.2139/ssrn.4968346

Jain, V., Sahay, R. R., & Nupur. (2025). Leveraging ESG rating and advanced analytics for portfolio optimization. *International Journal of System Assurance Engineering and Management*. https://doi.org/10.1007/s13198-025-03008-8

Kim, T., & Kim, H. Y. (2019). Forecasting stock prices with a feature fusion LSTM-CNN model using different representations of the same data. *PLOS ONE, 14*(2), e0212320. https://doi.org/10.1371/journal.pone.0212320

Krueger, P., Sautner, Z., & Starks, L. T. (2020). The importance of climate risks for institutional investors. *The Review of Financial Studies, 33*(3), 1067-1111. https://doi.org/10.1093/rfs/hhz137

Ma, Y., Han, R., & Wang, W. (2021). Portfolio optimization with return prediction using deep learning and machine learning. *Expert Systems with Applications, 165*, 113973. https://doi.org/10.1016/j.eswa.2020.113973

Markowitz, H. (1952). Portfolio selection. *The Journal of Finance, 7*(1), 77-91. https://doi.org/10.2307/2975974

Mehtab, S., & Sen, J. (2020). Stock price prediction using CNN and LSTM-based deep learning models. In *2020 International Conference on Decision Aid Sciences and Application (DASA)* (pp. 447-453). IEEE. https://doi.org/10.1109/DASA51403.2020.9317207

Naik, P., & Reddy, Y. V. (2021). Stock market liquidity: A literature review. *SAGE Open, 11*(1), 2158244020985529. https://doi.org/10.1177/2158244020985529

Oreshkin, B. N., Carpov, D., Chapados, N., & Bengio, Y. (2019). N-BEATS: Neural basis expansion analysis for interpretable time series forecasting. In *Conference Paper at International Conference on Learning Representations* (ArXiv:1905.10437). Cornell University. https://doi.org/10.48550/arXiv.1905.10437

Sajadi, S. M. A., Fereydooni, A., Athari, S. A., & Farhadi, S. (2025). A sustainable portfolio construction model based on ESG and deep learning algorithms: Evidence from the U.S. market. *Computational Economics*. https://doi.org/10.1007/s10614-025-10976-3

Satopaa, V., Albrecht, J., Irwin, D., & Raghavan, B. (2011). Finding a “kneedle” in a haystack: Detecting knee points in system behavior. In *2011 31st International Conference on Distributed Computing Systems Workshops* (pp. 166-171). IEEE. https://doi.org/10.1109/ICDCSW.2011.20

Vo, N. N. Y., He, X., Liu, S., & Xu, G. (2019). Deep learning for decision making and the optimization of socially responsible investments and portfolio. *Decision Support Systems, 124*, 113097. https://doi.org/10.1016/j.dss.2019.113097

Yahya, H. (2023). The role of ESG performance in firms’ resilience during the COVID-19 pandemic: Evidence from Nordic firms. *Global Finance Journal, 58*, 100905. https://doi.org/10.1016/j.gfj.2023.100905

Yu, J.-R., Paul Chiou, W.-J., Lee, W.-Y., & Lin, S.-J. (2020). Portfolio models with return forecasting and transaction costs. *International Review of Economics & Finance, 66*, 118-130. https://doi.org/10.1016/j.iref.2019.11.002