Finance Research Letters 83 (2025) 107672

Contents lists available at ScienceDirect

Finance Research Letters

journal homepage: www.elsevier.com/locate/frl

Can ESG disclosure predict carbon risk? Evidence from machine and deep learning models

Md Abubakar Siddique [a], Sitara Karim [b,c,*]

[a] College of Business, Abu Dhabi University, United Arab Emirates
[b] School of Business, Faculty of Business, Design and Arts, Swinburne University of Technology, Sarawak Campus, Kuching, Malaysia
[c] Department of Business Administration, Faculty of Management Sciences, ILMA University, Karachi, Pakistan

ARTICLE INFO

Keywords:
AI predictability
ESG disclosure
Carbon risk
Machine learning

ABSTRACT

As ESG disclosure becomes increasingly shaped by shifting regulatory agendas, its role in predicting carbon risk warrants closer examination. This study employs machine learning, deep learning, and ensemble techniques to assess whether ESG and financial indicators can effectively predict carbon risk. Results demonstrate that advanced AI models significantly outperform traditional regressions by capturing complex, non-linear relationships often overlooked by conventional methods. SHAP analysis further identifies environmental disclosure as the most influential predictor. These findings offer a timely, data-driven framework for investors, firms, and policymakers navigating the uncertainties of evolving sustainability reporting standards.

1. Introduction

Recent shifts in ESG-related policies, particularly under the Trump administration,[1] have introduced uncertainty into corporate sustainability reporting. The U.S. Securities and Exchange Commission's (SEC) rollback of climate-focused rules, including restrictions on shareholder resolutions and reduced transparency requirements, has weakened mechanisms that investors traditionally use to evaluate corporate risk. Meanwhile, several major financial institutions have withdrawn from net-zero commitments, and funding for clean energy initiatives has been clawed back. These changes raise doubts about the reliability of ESG disclosures as tools for risk management. In response to evolving policy shifts and investor expectations, firms frequently adapt their ESG strategies, revising disclosure practices and sustainability commitments. These shifts highlight the non-static nature of ESG behavior, shaped by external pressures, market sentiment, and regulatory enforcement. In this context, understanding whether ESG disclosure can still predict carbon risk becomes critically important. In this context, understanding whether ESG disclosure can still predict carbon risk becomes critically important. This study fills a key research gap by examining the predictive value of ESG disclosures amid regulatory volatility and shifting corporate priorities.

The push for standardized ESG disclosures is intensifying, particularly in the United States, where regulatory bodies like the Securities and Exchange Commission (SEC) are moving toward mandatory climate-related disclosures.[2] These regulations require firms to report carbon emissions, climate risks, and sustainability governance, reflecting an increased demand for corporate transparency

* Corresponding author.
E-mail addresses: md.siddique@adu.ac.ae (M.A. Siddique), skarim@swinburne.edu.my (S. Karim).

[1] Please see: https://www.reuters.com/sustainability/sustainable-finance-reporting/esg-watch-new-sec-rules-under-trump-turn-up-heat-sustainable-investors-2025-03-17/

[2] Please see: https://www.sec.gov/newsroom/press-releases/2024-31

https://doi.org/10.1016/j.frl.2025.107672
Received 25 March 2025; Received in revised form 23 May 2025; Accepted 30 May 2025
Available online 30 May 2025
1544-6123/

and risk accountability (Esty and de Arriba-Sellier, 2023). However, while ESG disclosures provide retrospective insights into corporate sustainability performance, their predictive power in assessing future carbon risk remains largely unexplored (Feyisetan et al., 2025; Rani et al., 2025; Yadav et al., 2024). The ability to forecast carbon risk using ESG metrics could significantly enhance investment decision-making, regulatory compliance, and corporate sustainability strategies.

Traditional approaches to carbon risk assessment have predominantly relied on linear regression models (Shu and Tan, 2023) and correlation-based analyses (Chen et al., 2025), assuming a fixed, linear relationship between ESG disclosures and carbon exposure. However, ESG data is multi-dimensional and highly dynamic, shaped by intricate interactions between environmental, social, governance, and financial factors. Given this complexity, conventional models struggle to capture non-linear dependencies and hidden risk patterns, resulting in suboptimal forecasting accuracy. The growing adoption of AI-driven predictive analytics in financial markets highlights the need for more sophisticated modeling techniques that can effectively analyze how ESG disclosures impact future carbon risk trajectories.

Existing research largely examines descriptive relationships rather than predictability, leaving a gap in whether ESG disclosures can serve as reliable risk indicators. Addressing this, our study applies machine learning and deep learning models to assess carbon risk predictability, integrating techniques such as Random Forest, Gradient Boosting, XGBoost, MLP, and GRU. Additionally, time-series forecasting (TFT) and ensemble learning (stacking, blending) are employed to determine if model combinations enhance predictive performance. This AI-driven approach aims to provide a more accurate, data-driven framework for investors, regulators, and policymakers managing ESG-related financial risks.

A key aspect of this study is the use of SHAP (Shapley Additive Explanations) analysis to improve model interpretability. While AI models enhance predictive accuracy, their black-box nature limits usability in financial decision-making. SHAP analysis identifies key ESG factors driving carbon risk, providing actionable insights for investors, regulators, and corporate leaders. This ensures AI-driven risk modeling aligns with policy objectives, including those set by the SEC reinforcing its practical relevance in sustainable finance.

This study's contribution is twofold. First, it introduces a novel AI-driven framework for ESG-based carbon risk prediction, bridging the current gap between disclosure transparency and forward-looking risk assessment. Second, it demonstrates the added value of deep learning and ensemble methods in enhancing predictive accuracy and model robustness compared to traditional approaches. The results offer critical insights for financial analysts, policymakers, and corporate sustainability officers, reinforcing the importance of predictive modeling in ESG risk management.

The remainder of this paper is structured as follows: Section 2 describes the Data and Methods employed. Section 3 presents predictive performance results. Section 4 provides key insights and implications in the Conclusion Section.

2. Methodology

2.1. Data and variables

Our study incorporates ESG scores, including EScore, SScore, and GScore, along with financial metrics such as Market Capitalization, Profit, Total Debt, Total Equity, and Capital Expenditure as key features. The target variable in this study is Carbon Risk (CR). To assess the predictability of ESG disclosure in relation to carbon risk, we analyzed data from 478 U.S. companies from 2013–2022. The data were obtained from Bloomberg Database.

2.2. Predictive models

2.2.1. Baseline models (Benchmarking approaches)

At first, we employed linear regression, ridge regression, and lasso regression as baseline models.

Linear regressions models the relationship between ESG disclosure variables and Carbon Risk:

$$CR = \beta_0 + \beta_1 E_{Score} + \beta_2 S_{Score} + \beta_3 G_{Score} + \beta_4 MKTCap + \beta_5 Profit + \varepsilon \tag{1}$$

Ridge regression introduces L2 regularization to penalize large coefficients:

$$CR = \text{argmin} \sum_{i=1}^{n} (y_i - \beta X_i)^2 + \lambda \sum_{j=1}^{p} \beta_j^2 \tag{2}$$

Lasso regression uses L1 regularization to enforce sparsity:

$$CR = \text{argmin} \sum_{i=1}^{n} (y_i - \beta X_i)^2 + \lambda \sum_{j=1}^{p} |\beta_j| \tag{3}$$

2.2.2. Traditional machine learning models

Among traditional machine learning models, we applied decision tree, random forest, gradient boosting (GBM), XGBoost, and Support Vector Machine.

Decision Tree is a non-linear model that partitions the feature space using the Gini Impurity:

$$Gini = 1 - \sum_{i=1}^{c} p_i^2 \tag{4}$$

Random Forest is an ensemble of multiple decision trees trained on bootstrapped samples.

$$\hat{y} = \frac{1}{n} \sum_{i=1}^{n} \hat{y}_i \tag{5}$$

Gradient Boosting (GBM) is an additive model that minimizes loss iteratively:

$$F_m(x) = F_{m-1}(x) + \gamma \, h_m(x) \tag{6}$$

XGBoost is an optimized GBM with L1 & L2 regularization:

$$L\ (\theta) = \sum_{i=1}^{n} (y_i - \hat{y}_i)^2 + \lambda \sum \beta_j^2 + \alpha \sum |\beta_j| \tag{7}$$

Support Vector Machines (SVM) is a supervised learning model that finds the optimal hyperplane to separate data points in a high-dimensional space. It maximizes the margin between different classes, ensuring better generalization.

For a linear SVM, the objective function is given by:

$$max_{w,b} \frac{1}{\| w \|} \tag{8}$$

Subject to:

$$y_i = (w^T x_i + b) \geq 1, \ \forall_i \tag{9}$$

For non-linear problems, SVM employs kernel functions such as radial basis function (RBF) or polynomial kernels to transform the data into a higher-dimensional space, making it linearly separable.

2.2.3. Deep learning models

In deep learning models, we employed multi-layer perception (MLP) and gated recurrent unit (GRU).

Multi-Layer Perceptron (MLP) is a feedforward neural network as shown below.

$$h^l = \sigma \left(W^l h^{l-1} + b^l \right) \tag{10}$$

Gated Recurrent Unit (GRU) is a simplified LSTM variant as shown below.

$$r_t = \sigma \ (W_r [h_{t-1}, \ x_t] + \ b_r \tag{11}$$

$$z_t = \sigma \ (W_z [h_{t-1}, \ x_t] + \ b_z \tag{12}$$

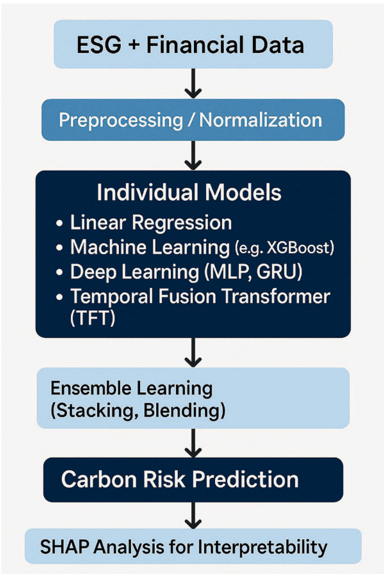

Fig. 1. Model Architecture Flow for ESG-Based Carbon Risk Prediction.

$$h_t = (1 - z_t) * h_{t-1} + z_t * \widetilde{h}_t \tag{13}$$

2.2.4. Advanced transformer-based model

We employed Temporal Fusion Transformer (TFT) which uses multi-head self-attention:

$$Attention\ (Q,\ K,\ V) = softmax\left(\frac{QK^T}{\sqrt{d_k}}\right)V \tag{14}$$

2.2.5. Ensemble learning techniques

In ensemble learning techniques, we employed Stacking and Blending. Stacking combines multiple base learners using meta-learning:

$$\widehat{y} = g\ (f_1(x),\ f_2(x),\ \ldots,\ f_n(x)) \tag{15}$$

While Blending is similar to stacking but uses a validation set.

For the stacking ensemble, we combined the top-performing models from each category: Random Forest, Gradient Boosting (GBM), XGBoost, and Multi-Layer Perceptron (MLP). These served as the base learners. A linear regression model was employed as the meta-learner to aggregate predictions from the base models and optimize overall predictive performance. For the blending ensemble, the same base learners were used, with their outputs combined using a weighted average approach on a validation set to generate final predictions.

Fig. 1 illustrates the end-to-end modeling pipeline used in the study. ESG and financial data undergo preprocessing before being fed into multiple predictive frameworks, including machine learning models (Random Forest, XGBoost, SVM), deep learning models (MLP, GRU), and the Transformer-based Temporal Fusion Transformer (TFT). Ensemble methods (stacking and blending) combine model outputs to improve predictive performance, while SHAP analysis is used to interpret the contribution of individual features to carbon risk predictions.

3. Predictive performance results

The empirical findings from Table 1 reveal that baseline regression models, Linear, Ridge, and Lasso Regression, offer only limited explanatory power, with R^2 scores hovering around 0.27. These models fail to capture more than a quarter of the variance in carbon risk, highlighting the inadequacy of linear approaches in modeling ESG-related risk. Given that ESG disclosure is multidimensional and interacts non-linearly with financial metrics, such traditional models are insufficient for producing reliable forecasts. This is especially concerning in light of recent regulatory shifts, such as the SEC's controversial amendments, which have made ESG disclosure both more politicized and less predictable in content and format highlighting the need for more adaptive, data-driven tools to anticipate ESG-related financial exposures.

Advancing to Table 2, the performance of machine learning models demonstrates a notable improvement in predictive accuracy. XGBoost achieves the highest R^2 (0.659), followed closely by Gradient Boosting (0.648) and Random Forest (0.632), indicating that ML methods are better suited for modeling the non-linear interactions inherent in ESG and carbon risk. While Decision Trees and Support Vector Machines also surpass baseline regressions, their lower accuracy suggests limited generalization ability. Feature importance rankings from these models are provided in Appendix A. These results imply that as regulatory enforcement becomes inconsistent, particularly in the U.S., investors and firms must rely more on robust, flexible forecasting models to mitigate carbon risk independently of shifting policy narratives.

Results from deep learning models in Tables 3 and 4 further support this transition. The MLP model delivers a test MAE of 1.50, while the GRU model achieves even lower test error (MAE = 1.45), signaling that neural networks can capture hidden patterns that static models cannot. These models are particularly useful when ESG disclosures are noisy, inconsistent, or delayed, allowing for more comprehensive forecasting. Detailed feature contributions for MLP and GRU are shown in Appendix B and Appendix C, respectively. Moreover, GRU's performance demonstrates the advantage of temporal modeling, relevant in a landscape where disclosure practices and sustainability commitments are continuously evolving.

Table 5 shows that the Transformer-Based Model (TFT) outperforms all previous models in terms of both train and test performance, with the lowest test MAE (1.38) and RMSE (1.89). This model excels at modeling sequential dependencies and long-term trends, making it particularly relevant in a time when disclosure quality and consistency are under political threat. Feature impact

Table 1

Baseline regressions.

Model	R2 Score	MAE	MSE	RMSE
Linear Regression	0.273953	2.042264	6.607788	2.570562
Ridge Regression	0.273791	2.041887	6.609259	2.570848
Lasso Regression	0.264059	2.040852	6.697826	2.588016

Note: This table reports the performance of baseline linear models, Linear, Ridge, and Lasso Regression, using ESG disclosure scores and financial indicators to predict carbon risk. The relatively low R^2 scores (approximately 27 %) indicate that these models fail to capture the complex, non-linear structure inherent in ESG–carbon risk relationships, underscoring the need for more advanced predictive techniques.

Table 2
Machine learning results.

Model	R^2 Score	MAE	MSE	RMSE
Decision Tree	0.561	1.56	4.21	2.05
Random Forest	0.632	1.37	3.56	1.89
Gradient Boosting (GBM)	0.648	1.31	3.4	1.84
XGBoost	0.659	1.28	3.29	1.81
Support Vector Machine	0.521	1.71	4.65	2.16

Note: This table gives the performance of machine learning models in predicting carbon risk using ESG and financial variables. Tree-based models, particularly XGBoost and Gradient Boosting, demonstrate markedly higher predictive accuracy (R^2 up to 0.659) compared to linear regressions. These results highlight the advantage of capturing non-linear interactions and complex feature relationships in ESG-driven risk modeling.

Table 3
MLP Neural Network Performance.

Metric	Value
Training Loss (MSE)	83.52: 3.90
Validation Loss (MSE)	18.04: 3.90
Test Loss (MSE)	3.9
Mean Absolute Error (MAE)	1.5
First 10 Predictions	10.54, 7.83, 11.56, 14.80, 9.20, 8.79, 11.46, 9.70, 9.21, 10.33

Note: This table presents the performance of the Multi-Layer Perceptron (MLP) neural network in predicting carbon risk. The model achieves a test MAE of 1.5, demonstrating improved accuracy over traditional regression approaches. The MLP effectively captures non-linear relationships in the ESG and financial feature space, validating the utility of deep learning for sustainability risk forecasting.

Table 4
GRU results.

Metric	Train Value (GRU)	Test Value (GRU)
MAE	1.08	1.45
MSE	3.02	3.85
RMSE	1.74	1.96

Note: This table shows the performance of the Gated Recurrent Unit (GRU) model, which incorporates temporal dynamics in ESG and financial data. The GRU achieves a lower test MAE than the MLP, indicating its strength in modeling time-dependent patterns relevant to carbon risk. These results underscore the value of sequence-aware deep learning models in sustainability forecasting.

scores for the TFT model are summarized in Appendix D. As regulatory landscapes become more fragmented and disclosure standards more ambiguous, models like TFT offer the flexibility and precision needed to extract meaning from incomplete or evolving ESG signals.

Finally, Table 6 presents result from ensemble learning, where stacking ($R^2 = 0.678$) marginally outperforms all other models suggesting that blending multiple predictive techniques, each capturing different facets of ESG-carbon risk relationships, yields the most robust results. Feature contribution insights from ensemble models are available in Appendix E.

The improved performance of the stacking model can be attributed to its ability to integrate the complementary strengths of diverse base learners. Tree-based models such as Random Forest and XGBoost capture complex feature interactions and handle non-linearity well, while deep learning models like MLP and GRU excel at modeling hierarchical structures and sequential patterns in the data. By combining these distinct learning paradigms, stacking reduces the individual biases of each model and leverages their unique generalization capabilities. In this way, ensemble approach enables a more balanced and robust prediction of carbon risk than any

Table 5
Transformer-Based model (TFT) results.

Metric	Train Value (TFT)	Test Value (TFT)
MAE	0.98	1.38
MSE	2.75	3.6
RMSE	1.66	1.89

Note: This table presents the performance of the Temporal Fusion Transformer (TFT), which captures long-term dependencies and variable interactions in ESG data. With the lowest test MAE and RMSE among all models, the TFT demonstrates superior predictive capability, making it a highly effective tool for forecasting carbon risk in a dynamic ESG reporting environment.

Table 6
Ensemble learning.

Model	R^2 Score	MAE	MSE	RMSE
Optimized Stacking Regression	0.678	1.22	3.1	1.76
Blending Regression	0.662	1.26	3.27	1.81

Note: This table reports the performance of ensemble learning approaches, including stacking and blending, which combine predictions from multiple base models. The stacking regression achieves the highest overall R^2 (0.678) and lowest error metrics, indicating that integrating diverse model architectures improves predictive accuracy for carbon risk compared to individual models.

single model alone.

As summarized in Table 7, across all models, Environmental Disclosure Score (E_Score) consistently emerges as the most significant predictor of carbon risk, followed by Market Capitalization (MKTCap) and Social Score (S_Score). The weak predictive power of traditional financial metrics like Profit, Equity, and Capital Expenditure signals a shift: carbon risk is increasingly driven by sustainability metrics, not conventional accounting indicators. This finding is vital for investors navigating post-regulatory ESG uncertainty, as it highlights the importance of credible environmental disclosure as a forward-looking risk signal, regardless of political will.

While advanced models such as GRU, TFT, and ensemble methods offer superior predictive performance, they inherently function as black-box systems, making their internal decision processes difficult to interpret. This poses a challenge for regulators and ESG auditors who require transparency and accountability in model-driven evaluations.

To enhance interpretability of the predictive models, we employed SHAP (Shapley Additive Explanations), which quantify the contribution of each input feature to the predicted carbon risk. SHAP values offer a unified measure of feature importance across models, allowing us to identify the most influential ESG and financial indicators. Across all models, environmental disclosure (E_Score) consistently emerged as the top predictor, followed by firm size (MKTCap) and social transparency (S_Score). These findings are especially relevant in the context of current regulatory debates, where transparency in environmental reporting is increasingly being mandated. SHAP-based insights provide a data-driven rationale for emphasizing environmental disclosure in ESG standards, thereby supporting more informed regulatory frameworks and corporate accountability mechanisms.

4. Conclusion

We examined the predictability of carbon risk using ESG disclosure data by applying a wide range of machine learning, deep learning, and ensemble learning techniques. Traditional regression models provided limited explanatory power, indicating that linear approaches fail to capture the complex relationships between ESG scores and carbon risk. In contrast, advanced models, particularly XGBoost, GRU, and TFT, demonstrated superior performance, highlighting the importance of non-linear and temporal modeling in sustainability risk prediction. The consistently strong performance of Environmental Disclosure Score (E_Score) across all models affirms its role as the most critical driver of carbon risk.

From a methodological perspective, the inclusion of deep learning and transformer-based models allowed for the modeling of time-dependent ESG patterns, which traditional models overlook. Furthermore, ensemble learning techniques such as stacking and blending yielded the highest accuracy, confirming that combining multiple algorithms can enhance model robustness. SHAP-based feature importance analysis reinforced the dominance of ESG dimensions, particularly E_Score and S_Score, over traditional financial metrics like profit and capital expenditure, emphasizing the shift from financial to sustainability-driven risk factors in modern finance.

These findings carry significant implications for investors, regulators, and corporate decision-makers, particularly in an environment of shifting ESG regulations and reporting standards. For investors, the study offers a predictive framework to proactively identify firms with high carbon risk exposure, enabling more resilient portfolio management. For regulators, the evidence underscores the need to mandate standardized, high-quality environmental disclosures, especially given the central role of E_Score in predicting carbon risk. For corporate leaders, the results emphasize the strategic importance of transparent ESG reporting, not just for compliance, but as a tool to mitigate risk and align with investor expectations. By bridging ESG transparency with forward-looking risk modeling, this research supports informed and sustainable decision-making across the financial and corporate landscape.

This study, while comprehensive in scope, has certain limitations that open opportunities for future research. Although real-time ESG disclosure data was utilized, the analysis is confined to U.S. firms, which may limit the generalizability of the findings to other regulatory and reporting environments. Additionally, the models do not yet incorporate alternative data sources such as news sentiment, climate events, or unstructured disclosure narratives that could capture broader dimensions of carbon risk. Future studies could extend this work by comparing ESG predictability across countries, exploring sector-specific effects, or integrating richer data streams to enhance both model accuracy and policy relevance in rapidly evolving sustainability landscapes.

CRediT authorship contribution statement

Md Abubakar Siddique: Conceptualization, Data curation, Methodology, Software, Formal analysis, Visualization, Writing – original draft, Writing – review & editing, Supervision. **Sitara Karim:** Conceptualization, Data curation, Methodology, Software, Formal analysis, Visualization, Writing – original draft, Writing – review & editing, Supervision.

Table 7
Summary across all models (machine learning, deep learning, and ensemble learning).

Model	Top Predictor	Secondary Factors	Weak Predictors
Machine Learning	**E_Score**	MKTCap, S_Score, Tdebt	Profit, Tequity, CapitalExp
MLP (Deep Learning)	**E_Score**	MKTCap, S_Score, G_Score	Profit, Tequity, CapitalExp
GRU (Deep Learning)	**E_Score**	MKTCap, S_Score, Tdebt	Profit, Tequity, CapitalExp
TFT (Transformer-Based)	**E_Score**	MKTCap, S_Score, Tdebt	CapitalExp
Ensemble Learning (Stacking & Blending)	**E_Score**	MKTCap, S_Score, Tdebt, G_Score	Profit, Tequity, CapitalExp

Note: This table summarizes feature importance findings across all model categories. E_Score consistently emerges as the top predictor of carbon risk, while financial indicators such as Profit, Tequity, and CapitalExp show limited influence. The consistency across models reinforces the central role of environmental disclosure in sustainability-focused risk prediction.

Acknowledgments

The authors gratefully acknowledge the support provided by Abu Dhabi University (Grant No. 19300819), which enabled the successful completion of this research.

Appendix A. Machine Learning Models Feature Importance

Rank	Feature	Importance Score	Impact on Carbon Risk
1	E_Score (Environmental Disclosure Score)	Highest	Strongest predictor, firms with detailed environmental reporting show clearer Carbon Risk trends.
2	MKTCap (Market Capitalization)	High	Larger firms tend to have structured ESG risk management strategies.
3	S_Score (Social Disclosure Score)	High	Social responsibility initiatives moderately influence Carbon Risk.
4	G_Score (Governance Disclosure Score)	Moderate	Governance transparency affects risk exposure but is less predictive than environmental/ social factors.
5	Tdebt (Total Debt)	Moderate	Higher debt levels increase financial vulnerability to ESG risks.
6	Profit	Low	Weak correlation with Carbon Risk, suggesting financial success alone doesn't indicate sustainability risk.
7	Tequity (Total Equity)	Low	Minimal impact, overshadowed by other financial variables.
8	CapitalExp (Capital Expenditure)	Lowest	Negligible impact, capital investments do not strongly influence Carbon Risk trends.

Note: This table presents SHAP-based feature importance rankings from machine learning models. E_Score is identified as the most influential predictor of carbon risk, followed by MKTCap and S_Score. Traditional financial metrics such as Profit and CapitalExpenditure contribute minimally, highlighting the dominance of ESG-related disclosures in model performance.

Appendix B. Multi-Layer Perceptron (MLP) Feature Importance

Rank	Feature	SHAP Score	Impact on Carbon Risk
1	E_Score (Environmental Disclosure Score)	Highest	MLP assigns significant weight to environmental disclosures, reinforcing their importance.
2	MKTCap (Market Capitalization)	High	Larger firms exhibit structured ESG risk frameworks, influencing carbon risk.
3	S_Score (Social Disclosure Score)	High	Strong social governance practices reduce sustainability risks.
4	Tdebt (Total Debt)	Moderate	Financial constraints affect ESG risk mitigation efforts.
5	G_Score (Governance Disclosure Score)	Moderate	Governance transparency is relevant but less impactful than E & S factors.
6	Profit	Low	Profitability alone does not significantly determine carbon risk.
7	Tequity (Total Equity)	Low	Minimal influence, consistent with machine learning findings.
8	CapitalExp (Capital Expenditure)	Lowest	Insignificant impact, indicating capital investments are not closely linked to ESG risk exposure.

Note: This table summarizes SHAP-based feature importance derived from the Multi-Layer Perceptron (MLP) model. E_Score remains the most significant predictor, consistent with machine learning results. MKTCap and S_Score also show strong contributions, while financial indicators such as Profit and CapitalExpenditure have minimal impact, reaffirming the relevance of ESG disclosures in carbon risk prediction.

Appendix C. Gated Recurrent Unit (GRU) Feature Importance

Rank	Feature	SHAP Score	Impact on Carbon Risk
1	E_Score (Environmental Disclosure Score)	Highest	Dominates time-series carbon risk predictions, confirming its predictive power.
2	MKTCap (Market Capitalization)	High	Larger firms exhibit structured carbon risk trends over time.
3	S_Score (Social Disclosure Score)	High	Social transparency influences risk but fluctuates over time.

(*continued on next page*)

(*continued*)

Rank	Feature	SHAP Score	Impact on Carbon Risk
4	Tdebt (Total Debt)	Moderate	Financial leverage affects risk patterns but is not the dominant factor.
5	G_Score (Governance Disclosure Score)	Moderate	Contributes to risk predictability but less than E & S factors.
6	Profit	Low	Weak temporal influence on carbon risk prediction.
7	Tequity (Total Equity)	Low	Insignificant effect in time-series forecasting.
8	CapitalExp (Capital Expenditure)	Lowest	No meaningful impact on carbon risk over time.

Note: This table displays SHAP-based feature importance from the GRU model, which captures temporal dependencies in the data. E_Score remains the dominant predictor of carbon risk across time, while MKTCap and S_Score show consistent but time-sensitive influence. Financial metrics such as Tequity and CapitalExpenditure contribute little, emphasizing the limited temporal relevance of traditional financial variables.

Appendix D. TFT Model Feature Importance

Rank	Feature	SHAP Score	Impact on Carbon Risk
1	E_Score (Environmental Disclosure Score)	Highest	Strongest predictor of Carbon Risk. Firms with higher environmental transparency have more predictable Carbon Risk patterns.
2	MKTCap (Market Capitalization)	High	Larger firms often have structured ESG policies, influencing Carbon Risk.
3	S_Score (Social Disclosure Score)	High	Social transparency (e.g., employee treatment, diversity) influences sustainability risks.
4	G_Score (Governance Disclosure Score)	Moderate	Governance quality affects sustainability commitments but has a lower direct influence compared to Environmental & Social factors.
5	Tdebt (Total Debt)	Moderate	Highly leveraged firms may have more exposure to Carbon Risk due to financial constraints on sustainability investments.
6	Profit	Low	Profitability appears to have a weaker direct impact on Carbon Risk, indicating that financial performance alone does not predict sustainability risks.
7	Tequity (Total Equity)	Low	Lower influence on Carbon Risk compared to other financial metrics.
8	CapitalExp (Capital Expenditure)	Lowest	Insignificant impact on Carbon Risk, suggesting that capital investment decisions may not be directly linked to sustainability disclosure trends.

Note: This table presents SHAP-based feature importance derived from the Temporal Fusion Transformer (TFT) model. E_Score is consistently the most influential variable, underscoring the significance of environmental disclosure in time-series prediction of carbon risk. MKTCap and S_Score also contribute meaningfully, while traditional financial variables such as Profit and CapitalExpenditure show limited predictive value in the TFT framework.

Appendix E. Ensemble Learning (Stacking & Blending) Feature Importance

Rank	Feature	Ensemble Contribution	Interpretation
1	E_Score (Environmental Disclosure Score)	Highest	Consistently the strongest factor across all models.
2	MKTCap (Market Capitalization)	High	Financial size remains crucial in all risk assessments.
3	S_Score (Social Disclosure Score)	High	Social transparency is a significant but secondary predictor.
4	Tdebt (Total Debt)	Moderate	Financial leverage matters but doesn't dominate risk prediction.
5	G_Score (Governance Disclosure Score)	Moderate	Board transparency & governance policies show mixed predictive relevance.
6	Profit	Low	Weak predictor, confirming previous findings.
7	Tequity (Total Equity)	Low	Minimal role in risk forecasting.
8	CapitalExp (Capital Expenditure)	Lowest	Negligible influence across all models.

Note: This table summarizes feature importance from ensemble learning models using stacking and blending. E_Score remains the most influential predictor across combined models, followed by MKTCap and S_Score. The consistency of results across ensemble and individual models reinforces the central role of environmental and social disclosures in determining carbon risk, while traditional financial variables contribute marginally.

Data availability

Data will be made available on request.

References

Chen, Z.H., Ren, Q., Gao, X., Kaakeh, M., Koedijk, K.G., 2025. Can ESG performance shape dynamic risk spillovers? Evidence from Chinese carbon and equity markets. Finance Res. Lett. 72, 106547.

Esty, D.C., de Arriba-Sellier, N., 2023. Zeroing in on net-zero: from soft law to hard law in corporate climate change pledges. *U. Colo*. L. Rev 94, 635.

Feyisetan, O.O., Alkaraan, F., Le, C., 2025. The influence of ESG on mergers and acquisitions decisions and organisational performance in UK firms: comparison between financial and non-financial sectors. J. Appl. Account. Res.

Rani, R., Vasishta, P., Singla, A., Tanwar, N., 2025. Mapping ESG and CSR literature: a bibliometric study of research trends and emerging themes. Int. J. Law Manag.
Shu, H., Tan, W., 2023. Does carbon control policy risk affect corporate ESG performance? Econ. Model. 120, 106148.
Yadav, S., Samadhiya, A., Kumar, A., Luthra, S., Pandey, K.K., 2024. Environmental, social, and governance (ESG) reporting and missing (m) scores in the industry 5.0 era: broadening firms' and investors' decisions to achieve sustainable development goals. Sustain. Dev.