

When Sustainability Meets Machine Learning: Reinforcement and Neural Evidence from an Emerging Market

Article Type:
Research Article

Setareh Azadvar¹ 

Department of accounting, Shandiz
Institute of Higher Education, Mashhad,
Iran
E-mail: setareh.azadvar@gmail.com

Alireza Azarberahman^{2*} 

Department of Accounting, Shandiz
Institute of Higher Education, Mashhad,
Iran
Corresponding Author
E-mail: a.azarberahman@shandiz.ac.ir

Received: 14 November 2025
Revised: 15 December 2025
Accepted: 17 March 2026
Available Online: 18 March 2026

Abstract

This study examined how firm-level environmental and social performance relates to stock price volatility in an emerging market characterized by limited transparency and weaker institutional frameworks. While prior research largely relied on linear models and focused on developed economies, this study adopted a dynamic, data-driven perspective to capture potentially nonlinear and time-dependent sustainability-risk patterns. Using a panel of non-financial firms listed on Tehran Stock Exchange (TSE) over the period 2011–2023, the firm-level environmental and social indicators were constructed based on a systematic analysis of sustainability disclosures. Empirical results from conventional linear regressions indicated weak and statistically insignificant average associations between sustainability performance and stock volatility. However, learning-based models, including reinforcement learning (RL) and LSTM neural networks, demonstrated superior ability to capture nonlinear and dynamic volatility patterns conditional on sustainability-related information. These findings suggested that the sustainability disclosures contain predictive information for volatility dynamics, even when linear risk-reduction effects are not evident. The study highlighted the importance of flexible modeling frameworks when assessing the financial implications of environmental and social performance in emerging markets.

Keywords

ESG performance, Stock volatility, Reinforcement learning (RL), LSTM neural networks, Sustainable finance, Nonlinear dynamics.

Cite this article: Azadvar, S. & Azarberahman, A., (2026). When sustainability meets machine learning: reinforcement and neural evidence from an emerging market. *Journal of Knowledge Economy Studies (JKES)*, 3(1), 47-71. DOI: <http://doi.org/10.22034/kes.2026.2077788.1087>



Authors retain the copyright and full publishing rights.
Online ISSN: 3060-7329

Publisher: Hazrat-e Masoumeh University.

Introduction

Corporate sustainability has migrated from being a discretionary add-on into a fundamental strategic dimension that shapes firm value, risk exposure, and stakeholder trust (Naseer et al., 2024). Evidence from developed markets suggests that higher environmental, social, and governance (ESG) performance is often associated with lower equity risk and cost of capital (Biju, 2025; Gupta & Chaudhary, 2023). However, these findings are derived from institutional environments characterized by strong disclosure regimes and investor protection, and therefore cannot be readily generalized to emerging markets (Biju, 2025; Gupta & Chaudhary, 2023). Yet in emerging markets, where institutional frameworks are weaker, transparency is limited, and standardized ESG ratings are sparse, the empirical evidence remains fragmented (Kansoy & Stasiulaitis, 2025; Rahat & Nguyen, 2024).

Environmental and social initiatives can reduce firm-specific risks by enhancing operational resilience, stakeholder relationships, and regulatory compliance (Iannone et al., 2025). However, such practices may also introduce new uncertainties and transitional cost burdens, particularly in less-developed institutional settings (Farūq & Chowdhury, 2025). The dual nature of the sustainability's impact on market risk — reduction through improved governance vs. increase through complexity and uncertainty — presents a substantive empirical puzzle. In particular, in a market context like Iran's, where ESG disclosures are largely voluntary and the investor protection is still evolving, assessing how sustainable practices influence the formation of equity market risk is critically important.

Prior research in emerging economies has primarily examined ESG performance's relation to firm valuation or profitability, often using static linear models and binary disclosure proxies. For example, Rahat and Nguyen (2024) reported on the impact of ESG on firm valuation in emerging markets but they did not focus on risk dynamics. Biju (2025) assessed the ESG-firm performance nexus and acknowledged that the direct link to firm risks remains underexplored. Meanwhile, normative frameworks in developed contexts do not necessarily transfer to emerging-market realities (Valencia Söderberg & Truong, 2024). Examining non-linear, temporally-dependent relationships between firm-level sustainability practices and equity risk in emerging markets, using advanced analytical methods suited to dynamic, state-contingent environments is a gap in literature. Importantly, addressing this gap requires moving beyond average linear effects toward models that prioritize risk dynamics and predictive performance, rather than causal inference based on static specifications. This study addressed these gaps in three major ways. First, it developed a context-specific, replicable environmental and social index in an Iranian context, where commercial ESG ratings are scarce. Second, it employed a hybrid empirical framework, combining econometrics, RL, and recurrent neural-network modeling, to reveal dynamic relations between sustainability performance and equity volatility. Third, it provided novel evidence from an emerging market context by examining how sustainability-related disclosures relate to the predictability and dynamics of stock volatility under conditions of weaker institutional infrastructure.

By integrating sustainability disclosure, advanced modeling, and an emerging-market setting, this research responds to the urgent call for more data-driven, risk-focused studies of corporate sustainability's financial implications (Biju, 2025; Kansoy & Stasiulaitis, 2025). It extends the frontier of knowledge on how sustainability shapes the firm-level equity risk beyond mature markets into institutional terrains with higher informational opacity and distinct regulatory challenges.

Literature Review and Hypotheses Development

Environmental Performance and Firm-level Stock Volatility

In theory, higher environmental performance should attenuate firm-level risk by reducing operational uncertainty, regulatory exposure, and reputational shocks, thereby stabilizing cash flows and dampening equity volatility. Stakeholder theory predicts that firms which meet salient stakeholder expectations (e.g., emissions abatement, resource efficiency, and compliance) enjoy greater legitimacy and smoother stakeholder relations, translating into lower perceived risk and more stable pricing in capital markets (Donaldson & Preston, 1995). Agency theory adds that credible environmental disclosure narrows information asymmetry between managers and outside investors, lowering the variance of belief updates that feed into price formation (Jensen & Meckling, 1976). Finally, signaling theory implies that costly, verifiable environmental actions serve as credible signals of managerial quality and long-run resilience. Such signals are more readily capitalized into prices when governance and disclosure quality make them believable, hence reducing volatility (Spence, 1973).

Empirically, several strands of evidence supported a volatility-reducing role of environmental (E) performance. Using a large sample from an emerging market, Liu et al. (2023) identified a causal reduction in firm-specific (idiosyncratic) volatility attributable to stronger ESG performance; their identification strategy and robustness checks point to information-environment improvements as a key channel. Complementary evidence showed that higher ESG scores lower stock price fragility—the propensity for sharp price collapses following adverse shocks—again consistent with reduced sensitivity of investors to noisy signals when sustainability practices are strong (Wang et al., 2023). Cross-market studies further documented that ESG portfolios often exhibit lower risk-adjusted variability than conventional benchmarks, with environmental quality a principal driver of the risk differential (Gupta & Chaudhary, 2023). Together, these findings indicated that environmental improvements reduce both the amplitude of day-to-day idiosyncratic fluctuations and the tail risk of crash-like episodes.

Mechanistically, three pathways link environmental performance to lower volatility. First, the operational-efficiency channel: cleaner production, energy efficiency, and waste minimization contain cost variability and supply-chain disruptions, smoothing earnings and investor expectations. Second, the regulatory-and-litigation channel: better environmental profiles lower the probability and severity of regulatory penalties and contingent liabilities, compressing the distribution of adverse cash-flow shocks and the priced risk premium (Chava, 2014). Third, the information channel: richer decision-useful

environmental disclosure reduces information asymmetry and curtails noisy trading, which directly lowers idiosyncratic volatility in microstructure-based frameworks. Consistent with these channels, recent studies showed that improvements in environmental scores reduce the cost of capital, a close cousin of priced risk, reinforcing the expectation of lower observed volatility when E performance strengthens.

A possible concern is that environmental initiatives could raise short-run costs and thus increase volatility. However, equilibrium models of corporate responsibility showed that when sustainability investments enhance product differentiation and stakeholder loyalty, they lower systematic risk and improve value despite upfront costs (Albuquerque et al., 2019). In markets with credible disclosure and governance, these benefits dominate, yielding a net reduction in volatility. The empirical evidence cited above aligns with this prediction and suggests the effect extends to idiosyncratic risks in emerging markets as well.

Theory and evidence converge on a clear prediction showed that as firms raise environmental performance—through measurable, disclosed actions that stakeholders and markets can verify—firm-level stock volatility should decline via operational, regulatory, and informational channels. Hence,

H1: Environmental performance is expected to be associated with lower stock volatility; however, this association may be weak or non-linear in emerging markets characterized by limited transparency.

Social Performance and Firm-level Stock Volatility

Social performance—covering employees' well-being, diversity and inclusion, community engagement, customer-centric practices, and social transparency—creates intangible capital that stabilizes firms' cash flows and investors' beliefs. From a stakeholder-theoretic view, addressing salient social claims strengthens legitimacy and trust, dampening the arrival of adverse, belief-shifting signals that would otherwise amplify the price volatility (Donaldson & Preston, 1995). Agency theory predicts that richer social disclosure narrows the information asymmetry between insiders and outside investors, reducing the variance of posterior beliefs and thus the idiosyncratic component of returns (Dhaliwal et al., 2011). A complementary "social-capital" perspective shows that investments in CSR accumulate relational assets (trust with employees, customers, and creditors), which pay off precisely in stressed states by cushioning operating and financing frictions that would translate into pronounced price swings (Lins et al., 2017).

A growing empirical literature linked stronger social practices to lower firm-level risk through multiple channels. In labor markets, high-quality employee relations proxy for superior human-capital management and lower disruption risk; firms recognized for employee satisfaction earned persistent abnormal returns and exhibited steadier fundamentals—evidence consistent with lower uncertainty premia and reduced volatility (Edmans, 2011). In information environment, initiating CSR disclosure—particularly among firms with superior social performance—reduces the cost of equity capital, a priced

manifestation of risk that typically co-moves with idiosyncratic volatility (Dhaliwal et al., 2011). In crisis settings, firms with higher CSR intensity (“social capital”) experienced markedly better performance, consistent with volatility dampening when shocks hit and relational buffers matter most (Lins et al., 2017). Large-sample asset-pricing evidence further indicated that the receipt or improvement of ESG ratings is followed by a decline in idiosyncratic stock risk even after stringent controls for confounds—an effect that, in many studies, is materially driven by the social dimension (Horn, 2023).

Risk-tail evidence aligns with these mechanisms. Corporate social responsibility is associated with lower stock-price crash risk—i.e., fewer extreme downside realizations—suggesting that social performance curtails opacity-driven bad-news hoarding and reduces investors’ sensitivity to adverse signals (Kim et al., 2014). The emerging-market evidence also documented a causal decline in idiosyncratic volatility as ESG performance improves, with information-environment upgrades cited as a principal channel; while composite, these effects are consistent with the idea that social practices reduce noise trading and stabilize expectations (Liu et al., 2023). Putting these strands together, theory and evidence jointly predict that higher social performance lowers firm-level volatility via human-capital reliability, stakeholder trust, disclosure-driven transparency, and tail-risk mitigation. Accordingly,

H2: Social performance is expected to be associated with lower stock volatility, although the average linear effect may be limited in emerging-market settings.

ESG Performance, Stock Volatility and Governance Transparency

While prior hypotheses posit a negative association between ESG performance and stock volatility, the emerging evidence suggests that this relationship is not uniform across firms or market states. Instead, the ESG–risk nexus is nonlinear and state-dependent, reflecting heterogeneity in corporate governance, market sentiment, and institutional quality (Broadstock et al., 2019). Firms with stronger governance frameworks and transparent disclosure systems can credibly convey the informational content of their ESG activities, amplifying the risk-mitigating effects of sustainability practices. Conversely, when governance transparency is weak, ESG actions may be perceived as “cheap talk” or even as opportunistic greenwashing, attenuating or reversing the expected volatility reduction (Bae et al., 2017; Li et al., 2018).

From a theoretical perspective, signaling theory (Spence, 1973) provides a natural explanation for this heterogeneity. ESG engagement functions as a costly signal of firm quality only if external investors can verify its authenticity. In high-transparency settings, ESG signals reduce information asymmetry and uncertainty about the firm fundamentals, thereby lowering volatility. In low-transparency contexts, however, investors may discount ESG claims, generating a muted or even positive relationship with volatility. Agency theory further reinforces this conditionality. Strong governance aligns managerial incentives with long-term stakeholders, enhancing the credibility of sustainability disclosures and strengthening their stabilizing impact on market risk (Jensen & Meckling, 1976).

Empirical findings increasingly supported these nonlinear and conditional dynamics. [Broadstock et al. \(2019\)](#) found that the ESG–risk relationship intensifies during periods of high market stress, implying state-dependence driven by shifting the investors' attention to firm resilience. [Bae et al. \(2017\)](#) showed that in markets with better institutional quality, ESG performance has stronger value- and risk-implications, confirming governance transparency as a key moderating factor. Moreover, recent machine-learning evidence suggested that the marginal effect of ESG on volatility diminishes beyond a threshold, producing an Inverted-U shape: initial sustainability gains lower volatility, but beyond a point, excessive ESG investments may introduce uncertainty about capital allocation efficiency ([Wang & Sonenshine, 2025](#)).

In emerging markets, where governance quality varies widely, this nonlinearity becomes especially salient. Firms with robust disclosure and board independence enjoy lower ESG-related information risk, while those with opaque practices experience weaker or even reversed effects ([Chao et al., 2022](#)). Consequently, the risk-reduction benefits of ESG are state-contingent—amplified when governance transparency and institutional quality are high, and weakened when these are lacking. The prior literature suggested that the ESG–risk relationship may be nonlinear and state-dependent, particularly in environments where governance quality and disclosure credibility vary across firms. While governance transparency is theorized to condition the informational value of sustainability practices, data limitations prevent a direct empirical test of this moderating mechanism in the present study. Accordingly, nonlinearity and state-dependence were explored using learning-based models rather than explicit interaction terms.

Method

The study adopted an ex-post-facto, market-wide design to quantify how firm-level environmental and social practices relate to equity risk in an emerging market. The sampling frame comprised all non-financial companies listed on TSE between 2011 and 2023. Eligibility required (i) continuous listing throughout the window without prolonged trading suspensions or delisting, (ii) complete financial statements and market data necessary to compute returns, volatility, and controls, and (iii) auditable disclosure of sustainability practices in annual reports, notes, or stand-alone responsibility reports. Financial institutions (banks, insurers, investment companies, and holding firms) were excluded due to incomparable balance-sheet structures and regulatory regimes. Systematic screening yielded 109 firms that satisfied inclusion criteria; firm-year observations from these issuers constituted the unbalanced panel. The primary data sources were the CODAL disclosure system and Tehran Securities Exchange Technology Management Company for market microdata.

The firm-level environmental and social performance was operationalized via a transparent, replicable content-analysis checklist synthesized from the domestic literature and aligned with common ESG themes. Each disclosed practice received a binary score, with the environmental index aggregating items such as energy efficiency programs, greenhouse-gas reporting or reduction targets, waste management, compliance

with environmental regulations, adoption of renewables, and green investment initiatives. The social index aggregated items such as workplace safety, employee training and development, workforce diversity and equal opportunity, customer-satisfaction programs, corporate philanthropy, and social transparency and accountability. While the checklist-based binary indices did not capture the intensity or qualitative depth of the sustainability actions, they offered transparency, replicability, and reduced measurement noise in settings where standardized ESG ratings are unavailable. In emerging markets characterized by heterogeneous disclosure quality, binary coding mitigates the subjective weighting and limits the researcher's discretion. Similar approaches were commonly adopted in disclosure-based ESG studies focusing on information availability rather than performance magnitude. For each firm-year, the environmental (E_Score) and social (S_Score) scores equaled the sum of disclosed items, producing bounded, interpretable indicators that are robust to scale differences across firms.

The market-based risks and control variables were constructed at matched frequencies. Annual stock volatility was defined as the standard deviation of firm-specific annual returns computed from closing prices; the market's conditional volatility was estimated at the firm level from daily returns using a GARCH (1, 1) specification, yielding an implied volatility proxy that captures time-varying risk. Additional controls followed the asset-pricing and corporate-finance literature: firm size (natural logarithm of total assets), financial leverage (total debt to total assets), broad market return (annual index return), and cash dividends per share. The firm-year stock return was retained both as a descriptive variable and, for reinforcement-learning reward design, as the realized payoff signal.

Preprocessing proceeded in four stages to ensure comparability and guard against artifacts. First, disclosures used in the checklist were dual-coded; disagreements were reconciled by adjudication to limit coder bias. Second, series were calendar-aligned to firm fiscal years; where necessary, price and dividend events were adjusted for corporate actions. Third, extreme observations in continuous financial variables were handled by moderate winsorization at conventional cutoffs after visual inspection via boxplots and distributional diagnostics. The ESG indices remained unaltered given their bounded nature. Fourth, continuous predictors were standardized within the training window only to avoid information leakage. The exploratory analysis assessed distributional shape (histograms, skewness, kurtosis), pairwise dependence (Pearson matrix), and missingness patterns; missing ESG items were treated as "not disclosed" by construction, while missing financials lead to firm-year omission. These steps, together with coding documentation, enabled full reproducibility.

Modeling followed a tiered strategy motivated by preliminary evidence that linear relations were weak while nonlinearity and temporal dependence were pronounced. As a baseline, multiple linear regressions related the stock volatility to E and S with controls, after verifying the standard assumptions. Residual normality was probed via Shapiro–Wilk tests and Q–Q plots, homoscedasticity via Breusch–Pagan tests and fitted-vs-residuals inspection, independence via the Durbin–Watson statistic, and

multicollinearity via variance inflation factors; specification proceeded with heteroskedasticity-robust (HC) standard errors if needed. These diagnostics confirmed that low explanatory power is not driven by assumption violations and motivate non-linear learners.

The first machine-learning arm was a value-based reinforcement-learning framework that learned investment policies contingent on sustainability states. ESG information was discretized into five ordered states from very low to very high using quantile thresholds on the environmental score or the mean of E_Score and S_Score. The agent's action space was binary—invest versus skip—and the one-period reward equaled the realized firm-year stock return when investing and zero when skipping. Learning used tabular Q-Learning with ϵ -greedy exploration. Hyper-parameters were tuned in ranges customary for stable convergence in noisy financial environments: learning rate $\alpha \in [0.01, 0.10]$, discount factor $\gamma = 0.90$, initial ϵ between 0.10 and 0.30 with geometric decay of 0.99, and 10,000 training episodes with randomized firm-year sampling. Convergence was monitored through the trajectory of cumulative reward and stabilization of Q-values across ESG states. The learned policy was extracted by greedy action selection on the terminal Q-table. The policy reasonableness was checked against economic priors and by sensitivity analyses that vary the ESG binning scheme and the reward definition to penalize volatility-seeking returns. Although the reward was defined as raw stock return, the reinforcement-learning framework was not intended to estimate risk-adjusted performance. Instead, it was used to examine whether ESG-related states contain systematic information that can guide sequential investment decisions. Risk implications were therefore inferred indirectly through state-dependent policy learning rather than explicit volatility penalization.

The second arm was a recurrent neural network that exploited the sequential nature of risk formation. A many-to-one LSTM architecture ingests multivariate firm-level sequences comprising lagged volatility, returns, and ESG scores along with time-varying controls, and outputs next-period volatility. Architectures with two to three stacked LSTM layers with 64–128 units each, tanh activations, and dropout between 0.2 and 0.5 fed into one or more dense layers. Weights were optimized with Adam at a 0.001 learning rate under a mean-squared-error objective. Training used mini-batches of 32–64 observations for 50–100 epochs under early stopping triggered by the validation RMSE. To preserve temporal causality, data were partitioned into contiguous train, validation, and test blocks with a rolling-origin scheme at the firm level; scalers were fitted on training windows only. The performance was summarized with MAE, RMSE, and out-of-sample R^2 ; learning curves were inspected for divergence to preempt overfitting. The model selection prioritized parsimonious architectures with the lowest validation RMSE and stable generalization.

The validation was deliberately layered. For the regression baseline, goodness-of-fit and inference relied on cross-sectional-time-series diagnostics and robust uncertainty estimates; for the reinforcement-learning policy, effectiveness was evaluated by comparing state-contingent Q-values and implied actions across ESG states and by

tracking cumulative reward growth over episodes. For the LSTM, generalization was assessed via a strict out-of-sample test set never seen during training or hyperparameter tuning; we additionally conducted a walk-forward experiment in which windows were advanced and models refitted to mimic real-time deployment. Across all learners, we computed bootstrap confidence intervals for key metrics and repeated the experiments under alternative random seeds to assess stability. Where appropriate, we contrasted the predictive accuracy against the linear baseline to quantify the incremental value of sequential and non-linear structure.

Robustness checks addressed construct validity and design choices. ESG states were re-defined using alternative cut points and an unsupervised k-means discretization; the reward was perturbed to risk-adjusted return (e.g., the return minus a volatility penalty) to gauge policy sensitivity; the LSTM input set was expanded with lagged GARCH volatility and pared back to ESG-only signals to isolate marginal contribution; horizons were shifted to test one- and multi-period forecasts; and leverage and size were re-scaled to mitigate the influence of outliers. The results were contrasted under firm and year fixed effects in the regression baseline to screen for latent heterogeneity. Collectively, these checks ensured that conclusions are not artifacts of particular measurement or modeling choices.

Implementation choices were justified by empirical diagnostics and the nature of the data. The binary ESG checklist yielded high inter-coder reliability and transparency in an environment where third-party ESG ratings were sparse. Weak linear associations observed in preliminary analyses supported the use of function approximators capable of learning non-linear and state-dependent relations. RL was appropriate when the estimand was a policy contingent on sustainability states, while LSTM networks were well-suited to capturing long-range dependencies and volatility clustering. The evaluation protocol emphasized out-of-sample performance and temporal integrity, aligning with best practice for financial prediction tasks.

Ethical and reproducibility safeguards were built in. All code implements deterministic data split with seeded randomness and logged configuration files for every run; preprocessing decisions, inclusion criteria, and checklist definitions were archived to enable audit. The analysis used only public disclosures and market data; no non-public or personally identifiable information was processed. Together, the sampling, measurement, modeling, and validation procedures constituted a coherent methodology for isolating whether and how firm-level environmental and social practices manifest in equity-market risk in an emerging-market setting.

Findings

Table 1 reports descriptive statistics for key variables including mean, median, standard deviation, minimum, maximum, skewness, and kurtosis. Skewness values clustered around zero indicate that the distributions were approximately symmetric. All kurtosis values were below 3, implying thinner-tailed distributions than the normal benchmark.

These distributional features suggest limited tail risk and provide a suitable basis for the subsequent modeling and inference.

Table 1.
The Descriptive Statistics of the Study Variables

variable	Mean	Median	SD	Min	Max	Skewness	Kurtosis
E_Score	55.39	55.65	14.55	10.00	98.33	-0.07	2.84
S_Score	59.86	59.77	18.01	5.00	100.00	-0.03	2.79
Stock return	0.04	0.03	0.32	-0.49	0.59	0.07	1.80
Stock volatility	0.51	0.52	0.23	0.10	0.90	-0.03	1.76
Implied volatility	0.41	0.42	0.17	0.12	0.70	-0.06	1.82
Firm size	9.71	9.70	1.05	7.30	12.16	-0.00	1.94
Leverage	0.45	0.45	0.18	0.05	0.95	0.03	2.67
Market return	0.02	0.02	0.19	-0.30	0.35	0.05	1.80
Cash dividend	6068.73	6018.69	3443.21	48.07	11993.16	-0.00	1.81

Note: E_Score and S_Score are checklist-based indices scaled from 0 to 100.


(Source: The Researcher's Findings)

Table 1 presents descriptive statistics for the variables employed in this study, including environmental and social scores (E_Score and S_Score) as well as key market-based indicators such as stock return, stock volatility, implied volatility, firm size, leverage, market return, and cash dividend. The mean values of the environmental and social scores—55.39 and 59.86, respectively—suggest that the sampled firms, on average, maintained moderate-to-high levels of environmental and social engagement. The relatively high standard deviations (14.55 and 18.01) indicate substantial heterogeneity in ESG commitment across listed Iranian firms. From a market-behavior standpoint, the mean stock return of 0.04 with a relatively large standard deviation (0.32) implies positive but highly dispersed returns. The slight positive skewness (0.07) points to a mild right-tail bias, indicating that a limited number of firms experienced higher-than-average returns. The stock volatility and implied volatility had mean values of 0.51 and 0.41, respectively, both with a kurtosis below 2, reflecting concentrated distributions with limited extreme fluctuations. The firm size, with a mean logarithmic value of 9.71 and a kurtosis of 1.94, reveals that most firms in the sample were of medium to large scale. The average leverage ratio of 0.45 denotes a balanced capital structure, suggesting moderate reliance on debt financing. The market return, averaging 0.02, together with the mean cash dividend of 6,069 Toman per share, reflects the limited overall profitability of the TSE during the observation period. The analysis of skewness and kurtosis shows that the majority of variables followed approximately symmetric, thin-tailed distributions—conditions favorable for subsequent econometric and machine-learning modeling. These statistical properties ensure numerical stability and are particularly advantageous for reinforcement-learning and recurrent-neural-network training, which require well-behaved input distributions.

Figure 1 illustrates the distribution of the study's primary variables using box plots. Observable outliers appeared mainly in the E_Score and Cash_Dividend variables, reflecting the heterogeneity of environmental disclosure and payout policies across firms. For most variables, the median lay near the center of the box and the interquartile


range was compact, suggesting approximately symmetric and near-normal distributions. These patterns indicate that the dataset was well-behaved and suitable for parametric modeling. The presence of moderate outliers provides valuable diagnostic insight for subsequent normalization and scaling procedures, thereby supporting the robustness of machine-learning algorithms—particularly reinforcement-learning and LSTM-based models—used in the later stages of analysis.

Figure 1.
The Box Plots of Key ESG and Financial Variables



(Source: The Researcher's Findings)

Figure 2.
Pearson Correlation Matrix of the Main Variables




(Source: The Researcher's Findings)

Figure 2 presents the Pearson correlation matrix among the main variables of the study. The figure reports pairwise linear correlation coefficients and provides the same statistical information as a conventional correlation table, while offering a more intuitive visualization of the strength and direction of associations. As shown, the *E_Score* and

S_Score were positively correlated, reflecting internal consistency across sustainability dimensions. In contrast, ESG dimension did not exhibit a meaningful linear correlation with stock volatility or stock returns. The uniformly weak Pearson coefficients indicate that ESG-risk relationships did not operate through simple linear channels, thereby motivating the use of nonlinear and state-dependent modeling frameworks, such as RL and LSTM neural networks, in subsequent analyses.

Figure 3 visualizes the temporal and cross-sectional patterns of the study variables over the 2011–2023 period. Each panel depicts a two-dimensional heat map in which the color intensity represents the magnitude of a variable for each firm-year observation. The plots reveal substantial heterogeneity across firms and time, highlighting fluctuations in market behavior and sustainability performance. For instance, the Firm Size exhibited a relatively stable trajectory over time, reflecting the gradual evolution of corporate scale in the Iranian market, whereas the Market Index Return and Stock Volatility displayed more dispersed and dynamic patterns across firms. These variations capture both idiosyncratic shocks and sector-specific cycles that are essential for learning-based models to detect latent temporal dependencies. The heat maps provide an intuitive visualization of how ESG indicators and financial variables evolve jointly across time and firms, serving as an empirical foundation for reinforcement-learning and LSTM-based modeling in subsequent sections.

Figure 3.
The Heat Maps of Selected Variables by Year and Firm




(Source: The Researcher's Findings)

Figure 4 presents four diagnostic plots used to examine the assumptions of the classical regression framework. The upper-left panel shows a histogram of residuals, which approximates a normal distribution with a mild kurtosis. The upper-right panel (residuals versus fitted values) exhibited no discernible pattern, indicating homoscedasticity and the absence of systematic bias. The lower-left Q-Q plot demonstrated a slight deviation from the reference line at the tails, which may be

attributed to a few outlying observations. The lower-right panel (residual sequence plot) revealed a random dispersion of residuals across observation indices, supporting the assumption of independence. Overall, the diagnostic results confirm that the regression model satisfied the key classical assumptions to an acceptable degree, validating its use as a benchmark for subsequent machine-learning analyses.

Figure 4.

The Diagnostic Plots of Residual Analysis for the Classical Regression Model



(Source: The Researcher's Findings)

Table 2 reports the Pearson correlation coefficients between the environmental and social scores (E_Score and S_Score) and stock volatility. The correlation coefficients were 0.02 for E_Score and 0.02 for S_Score, both positive but very close to zero. These results indicate the absence of a meaningful linear association between the firms' ESG performance and their stock return volatility. Furthermore, the p-values for both correlations exceeded the 0.05 significance level, suggesting that the observed relationships are statistically insignificant. Accordingly, it cannot be inferred that the ESG dimensions exerted a direct effect on market volatility. This finding aligns with several prior studies conducted in emerging markets, which suggest that the influence of ESG factors on firm risk is indirect or context-dependent. The weak linear correlations underscored the need to incorporate additional control variables and adopt more sophisticated modeling frameworks. Therefore, a multiple regression analysis was subsequently employed to examine the potential direct effects of ESG indicators after accounting for firm-specific characteristics.

Table 2.

The Pearson Correlation between ESG Scores and Stock Volatility

variable	Correlation Coefficient	p-Value
E_Score	0.02	0.30
S_Score	0.02	0.37

(Source: The Researcher's Findings)

Table 3 presents the results of the multiple regression model assessing the direct effects of ESG indicators (E_Score and S_Score) and control variables—firm size, leverage ratio, and market index return—on stock volatility. The findings reveal that none of the explanatory variables were statistically significant at the 95% confidence level. The coefficient of determination (R^2) was notably low (approximately 0.00), indicating that the model explained only a negligible portion of the variance in stock volatility. The overall model p-value exceeded the 0.05 threshold, confirming that the regression equation was not statistically significant. However, the Durbin-Watson statistic of 2.03 suggests that the residuals are free from serial correlation, implying no evidence of autocorrelation within the model. These results suggest that ESG indicators, in isolation and within this specific linear specification, did not meaningfully explain the fluctuations in stock volatility. The weak explanatory power may stem from unobserved latent factors, structural inefficiencies in the Iranian capital market, or potential time lags between corporate social responsibility actions and market reactions. Future research is encouraged to explore nonlinear or interaction-based modeling approaches—such as machine learning or reinforcement frameworks—to capture the complex dynamics between sustainability and financial risk.

Table 3.
The Results of Multiple Regression Explaining the Stock Volatility

variable	Coef.	95% CI (Lower)	95% CI (Upper)	VIF
Intercept	0.47	0.34	0.60	-
E_Score	0.00	-0.00	0.00	1.18
S_Score	0.00	-0.00	0.00	1.31
Firm Size	-0.00	-0.01	0.01	1.56
Leverage	-0.00	-0.07	0.06	1.45
Market Index Return	0.00	-0.06	0.06	1.05

(Source: The Researcher's Findings)

In addition to the regression analysis, the Variance Inflation Factor (VIF) was computed to examine the potential multicollinearity among the independent variables included in the model. All VIF values were found to be close to 1, and none exceeded the conventional thresholds of 5 or 10. These results confirm that no serious multicollinearity existed among the explanatory variables. Consequently, there was no concern regarding the coefficient instability or inflated standard errors due to internal correlations within the predictors. From a statistical perspective, the low VIF values strengthened the validity of the regression model by confirming the independence of the explanatory variables. The consistently low VIFs for the ESG indicators further suggest that the simultaneous inclusion of E_Score and S_Score did not introduce estimation bias or redundancy in the model. This diagnostic result is particularly important in multivariate analyses, where high inter-correlation can compromise the precision and interpretability of regression coefficients. The model structure demonstrates acceptable independence among the explanatory variables, ensuring reliable estimations for subsequent inferences.

The Reinforcement Learning Model: Q-Learning

The Q-Learning model is employed as a RL approach through which an autonomous agent interacts with the environment and iteratively learns optimal decisions over time. In this study, the environment was defined based on the firms' ESG performance. Specifically, the ESG scores were discretized into five ordered categories ranging from *very weak* to *very strong*. This state-space construction enabled the agent to evaluate and update its investment strategy across varying sustainability conditions. The rationale behind this design lies in the notion that firms with superior ESG profiles may exhibit more favorable long-term performance and risk characteristics. Thus, a state-based segmentation of ESG scores provides a structured foundation for learning ESG-driven investment policies.

The action set consisted of two choices: (0) do not invest, and (1) invest in the given firm. In the baseline specification, the reward was defined as the realized stock return: when the agent selected the investment action, the corresponding realized return was assigned as the reward; when the agent chose not to invest, the reward was set to zero. We note, however, that in markets with positive average returns this raw-return formulation can mechanically favor continuous investment and may therefore reflect a buy-and-hold tendency rather than any ESG-related risk mechanism.

To ensure that the reinforcement signal reflects the risk–return trade-offs, we adopted a risk-adjusted reward specification that penalizes conditional volatility:

$$\text{Reward}_t = r_t - \lambda \cdot \sigma_t$$

Where r_t denotes the realized return in period t , σ_t is rolling (or conditional) volatility, and $\lambda \geq 0$ captures the degree of risk aversion. This modification allowed the agent to learn state-contingent policies that were informative about risk-sensitive investment decisions rather than raw performance alone.

By adopting this framework, the Q-Learning model facilitated the extraction of data-driven investment rules that reflected both financial performance and sustainability characteristics. As a result, the approach provides insights into whether responsible and ESG-oriented investment policies can be learned and systematically executed within a dynamic environment.

Table 4.

The Core Components of the Q-Learning Framework for ESG-Based Investment Analysis

Component	Definition	Implementation in This Study
State Space	Represents the underlying ESG condition of each firm, categorized into discrete sustainability levels (e.g., from very low to very high).	The firms' ESG scores are discretized into five ordered levels based on the environmental (E_Score) or the average of the environmental (E_Score) and social (S_Score), forming states 1 to 5.
Action	The decision made by the learning agent in each state, such as taking a position or remaining inactive.	Two discrete actions are defined: (0) do not invest and (1) invest in the firm.
Reward Function	Feedback received from the environment after taking an action, representing gains or losses associated with the chosen action.	The realized stock return (Stock_Return) serves as the reward. If the action = invest, the reward equals the firm's realized return; otherwise, the reward is set to zero.

(Source: The Researcher's Findings)

Table 5 reports the final Q-values obtained for the five discrete ESG states. A clear pattern emerges: across all states, the Q-value associated with the “Invest” action exceeded that of the “Skip” action. This indicates that, on average, taking an investment position generated higher expected cumulative rewards within the RL environment compared to remaining inactive. This result implies that, over the learning horizon, the agent identified investing in firms—regardless of their ESG tier—as a superior policy relative to staying out of the market. The highest Q-value for the invest action was observed in State 5 (the strongest ESG level), with a value exceeding 0.33, suggesting greater reinforcement and a stronger confidence signal toward investing in high-ESG firms.

Table 5.
The Final Q-Values for Discrete ESG States

ESG State	State1	State2	State3	State4	State5
Skip (0 = Do Not Invest)	0.01	0.01	0.01	0.01	0.01
Invest (1 = Invest)	0.21	0.18	0.27	0.23	0.33

(Source: The Researcher's Findings)

Table 6 presents the optimal policy learned by the Q-Learning algorithm across the five discrete ESG states. Consistent with the final Q-values, the model selected “Invest” as the optimal action in all ESG states. This indicates that, throughout the learning horizon, taking an investment position consistently yielded higher expected cumulative rewards relative to remaining out of the market. Importantly, this result reflects the agent’s policy under the specific reward structure and market data used in the study, where realized stock returns served as the reinforcement signal. While the policy implies that investing is systematically favored across all ESG levels, the stronger reinforcement observed at higher ESG tiers aligns with the notion that firms with robust sustainability characteristics may offer more favorable reward profiles. Nonetheless, these findings should be interpreted in the context of predictive policy learning rather than causal inference.

Table 6.
The Optimal Policy Learned Across ESG States

ESG State	State1	State2	State3	State4	State5
Optimal Action	Invest	Invest	Invest	Invest	Invest

(Source: The Researcher's Findings)

Table 7 reports the evolution of cumulative rewards across the training episodes in the Q-Learning model. At the beginning of the learning process, the cumulative reward was relatively low, approximately 0.85, reflecting the exploratory stage in which the agent has not yet learned an effective policy. As the training progresses, the cumulative reward steadily increased, reaching 10.15 by Episode 600. This upward trajectory indicates successful policy refinement and demonstrates that the agent gradually improves its investment decisions based on accumulated experience and feedback from the environment. The consistent increase in cumulative reward suggests that the algorithm effectively internalizes patterns from historical market behavior, allowing it to adopt more profitable strategies over time. While this result implies that ESG-based decision states provide useful signals within the RL framework, these findings should be

interpreted as evidence of policy learning performance, rather than definitive proof of causal superiority of ESG-aligned investments.


Table 7.
The Cumulative Reward Progression during the Q-Learning Training

Training Episode	1	100	200	300	400	500	600
Cumulative Reward	0.85	4.26	6.48	7.93	9.01	9.68	10.15

(Source: The Researcher's Findings)

This section summarizes the results of the Q-Learning model implemented to evaluate the investment decisions conditional on firms' ESG performance levels. The primary objective is to identify the optimal action—invest or abstain—across different sustainability states. To achieve this, Q-values were estimated for each action within each ESG tier and reported in tabular and graphical forms. Subsequently, the optimal policy was derived by selecting the action with the highest Q-value in each state. In addition, the learning dynamics were assessed by examining the trajectory of cumulative rewards over the training episodes, demonstrating the convergence of the agent's policy toward more profitable behaviors. These results illustrate that the RL framework can learn decision rules sensitive to ESG information and adaptively assign investment actions based on observed performance patterns. While these insights highlight the role of ESG signals in the learned policy, the findings should be interpreted as reflecting predictive investment behavior under the specified reward structure rather than establishing a causal superiority of ESG-compliant firms.


Figure 5.
The Comparison of Q-Values for “Invest” and “Skip” Actions Across ESG Levels



(Source: The Researcher's Findings)

Figure 5 presents the learned Q-values for the two possible actions—invest and skip—across the five ESG categories. In most states, the Q-value associated with the invest action exceeded that of the skip action, indicating that the agent, on average, expected higher cumulative returns when choosing to invest. This pattern suggests a reinforcement signal consistent with stronger ESG performance, as the agent increasingly favored investment decisions in higher-tier ESG states.


Figure 6.
The Optimal Policy across ESG States



(Source: The Researcher's Findings)

Figure 6 illustrates the optimal action selected by the Q-Learning model for each ESG state, where 0 denotes “skip” and 1 denotes “invest”. The model consistently recommends investing across all ESG levels. This outcome reflects the learned policy under historical market dynamics and the adopted reward formulation, indicating that investing was reinforced as the dominant strategy throughout the training. Nonetheless, the result should be interpreted within the boundaries of the learning setup and does not imply universal dominance of ESG-driven portfolios in all market conditions.

Figure 7.
The Cumulative Reward Trajectory over the Training Episodes



(Source: The Researcher's Findings)

Figure 7 depicts the progression of cumulative rewards over the training episodes. The upward and stabilizing trend confirms that the model successfully converged toward an improved policy as training proceeded. The gradual increase in cumulative reward suggests the effective learning of investment behavior conditioned on ESG signals and demonstrates the model’s ability to internalize performance feedback over time.

Recurrent Neural Network Modeling with LSTM

In this section, the performance of a Long Short-Term Memory (LSTM) neural network model is evaluated for forecasting the market volatility using ESG-based time-series inputs. The plotted results demonstrate a strong alignment between the actual volatility values and the model’s predicted series, indicating that the LSTM architecture was

capable of effectively capturing the temporal structure and dynamic patterns embedded in the data.

The evaluation table reports the model's predictive accuracy based on three commonly used metrics of Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and the coefficient of determination (R^2). It is important to emphasize that volatility is a highly persistent process in financial time series, and consequently, high predictive accuracy—particularly in terms of R^2 —is commonly observed in models that effectively exploit lagged volatility information. Therefore, a high R^2 in this context should not be interpreted as evidence of superior structural modeling alone, but rather as a reflection of volatility clustering and temporal dependence. The reported values suggest that the model achieved a high degree of explanatory power and satisfactory predictive accuracy in reconstructing and forecasting stock-level volatility. In particular, the R^2 value approaching 1 reflected strong model fit under the specified input features and training configuration, although such performance should be interpreted within the controlled experimental setting and acknowledged as subject to typical time-series forecasting risks such as overfitting and regime variation. In this regard, the strong performance should be understood primarily as the model's capacity to capture the persistent volatility dynamics, rather than as an indication that all predictive power originates from the ESG-related inputs.

These results highlight the potential of ESG-enhanced sequential learning frameworks in modeling financial market volatility, supporting the notion that the sustainability-related signals may contain information relevant for forward-looking risk dynamics. However, this contribution should be viewed as incremental rather than dominant, as the bulk of forecasting power in volatility models typically arises from autoregressive dynamics and volatility clustering.

Table 8.
The LSTM Model Evaluation Metrics for Volatility Forecasting

Evaluation Metric	MAE	RMSE	R^2
Value	0.0583	0.0794	0.8835

(Source: The Researcher's Findings)

The results in Table 8 indicate the robust predictive performance of the LSTM model, with an R^2 of approximately 0.88, suggesting that the network explained around 88% of the variation in stock-level volatility under the given specification. The strong predictive performance of the LSTM model is partly attributable to the inclusion of lagged volatility, which is known to be highly persistent in financial time series. To further disentangle the contribution of ESG variables from that of lagged volatility and returns, an ablation-style comparison is conceptually informative. In particular, an LSTM specification excluding ESG inputs but retaining the same lag structure would be expected to preserve a substantial portion of predictive accuracy, given the volatility persistence. ESG variables are therefore interpreted as complementary signals that enhance volatility dynamics modeling rather than as primary drivers of predictive accuracy. The relatively low MAE and RMSE values further confirm that the model

produced forecasts that are closely aligned with actual observations, reflecting effective learning and stable generalization across the dataset.

Table 9 compares the model performance across the training, validation, and testing phases. As expected, the model performed slightly better during the training; however, the performance gap remained small across all three data segments. This consistency indicates satisfactory generalizability and suggests that the model is not subject to substantial overfitting.

Table 9.

The Comparison of Performance across Training, Validation, and Test Sets


Data Split	MAE	RMSE	R ²
Training	0.05	0.07	0.89
Validation	0.05	0.07	0.86
Test	0.06	0.08	0.84

(Source: The Researcher's Findings)

Figure 8 compares the actual stock volatility series (solid orange line) with the LSTM-predicted values (dashed light-orange line) over a rolling window of 100 time steps. The close co-movement between the two curves—particularly during pronounced upward and downward movements (steps 0–30 and 60–100)—demonstrated the model's ability to capture nonlinear market dynamics. Minor deviations at sharp turning points largely reflected sudden market shocks not fully represented in the training distribution. Overall, this visual evidence further supports the LSTM model's reliability in forecasting volatility for ESG-oriented equities.

Figure 8.


The Training and Validation RMSE across Epochs



(Source: The Researcher's Findings)

Figure 9 illustrates the RMSE trajectory over 100 training epochs for both the training (blue line) and validation (orange line) datasets. The steady decline in error values until approximately epoch 60 reflected a progressive improvement in the model's predictive capability. Beyond this point, both curves fluctuated slightly while remaining closely aligned, indicating a stable convergence and the absence of severe overfitting. The dashed black vertical line marked the minimum validation RMSE, serving as a reference point for potential early stopping and optimal weight selection. These results suggest that the chosen architecture and learning rate configuration are appropriate for the ESG-driven volatility forecasting task.

Figure 9.
The Learning Curve of the LSTM Network



(Source: The Researcher's Findings)

To prevent the information leakage, data splitting was conducted along the time dimension within firms. For each firm, observations were partitioned into non-overlapping training, validation, and test windows in chronological order. No future information enters the model estimation or hyperparameter tuning, ensuring a strictly out-of-sample evaluation.


As shown in Table 10 and Figure 10, the LSTM model demonstrates superior performance across all evaluation metrics relative to the benchmark models. As shown in Table 10 and Figure 10, the LSTM model demonstrates strong predictive performance relative to the classical regression benchmark. It is important to clarify, however, that Q-learning was not a volatility forecasting model but a policy-learning framework. As such, its performance metrics were not directly comparable to forecasting accuracy measures derived from regression or LSTM models. The comparison was therefore intended to highlight methodological differences and complementary analytical roles rather than to establish a strict ranking of predictive models.

Table 10.
The Comparative Performance of Alternative Models in Volatility Forecasting

Model	MAE	RMSE	R ²
Classical Regression	0.08	0.10	0.26
Q-Learning (RL)	0.06	0.09	0.64
LSTM Neural Network	0.05	0.07	0.88

(Source: The Researcher's Findings)

Figure 10.
The Comparative Performance of the Models



(Source: The Researcher's Findings)

Based on the comparative results presented, it is evident that machine-learning and reinforcement-learning approaches—particularly the LSTM network—demonstrate superior performance relative to the traditional statistical model. Based on the comparative results presented, machine-learning approaches—particularly the LSTM network—exhibited strong predictive performance in volatility forecasting tasks. This performance primarily reflected the ability of deep sequential architectures to exploit temporal dependence and volatility persistence. ESG-related variables provided incremental informational content but should not be interpreted as the primary source of predictive accuracy. RL, in turn, served a complementary role by illustrating state-dependent decision mechanisms rather than by offering a direct forecasting benchmark.

Discussion and Conclusion

The empirical evidence obtained in this study indicated that the environmental and social performance do not exhibit statistically significant average linear associations with firm-level stock volatility in an emerging market context. The near-zero correlations and statistically insignificant coefficients observed in traditional regression models indicated that ESG variables, when modeled linearly, provide limited explanatory power for stock volatility; however, the learning-based models revealed that sustainability-related information can be informative for capturing nonlinear and dynamic volatility patterns beyond the average linear effects. These findings align with stakeholder theory ([Donaldson & Preston, 1995](#); [Freeman, 1984](#)) and agency theory ([Jensen & Meckling, 1976](#)), both of which posited that enhanced transparency and alignment of managerial decisions with stakeholder interests bolster investor confidence and thereby reduce market uncertainty.

While the linear models did not provide evidence of statistically significant average effects, the learning-based results suggested that environmental information may be more informative than social information in capturing nonlinear and dynamic volatility patterns. This difference should be interpreted in a predictive and exploratory sense rather than as evidence of a causal or statistically significant risk-reduction effect. The environmental practices may be more tightly linked to operational processes and regulatory exposure, which learning-based models can exploit when modeling volatility dynamics, whereas social practices may influence market perceptions more gradually and indirectly.

The validation of the non-linearity and state-dependence hypothesis constitutes one of the most impactful contributions of this research. The learning-based analyses pointed to nonlinear and state-dependent patterns in the ESG–volatility relationship. While prior literature suggested that governance transparency may condition the informational value of sustainability practices, this study did not directly test the governance-based moderation due to data limitations. Accordingly, references to governance were interpreted conceptually, serving to contextualize the observed nonlinear dynamics rather than to establish a tested conditional or risk-mitigating mechanism. This pattern dovetails with findings from [Rahat and Nguyen \(2024\)](#) and [Biju](#)

(2025), which highlighted that the informational value of ESG disclosures is contingent on institutional credibility.

The integration of traditional econometric models with artificial-intelligence techniques allowed us to capture both linear relationships and temporal, non-linear dependencies—addressing a lacuna in the extant ESG-finance literature. The superior predictive performance of the LSTM model relative to linear regression underscored the importance of modeling the dynamic, non-linear nature of risk in sustainability-driven environments. The implications are clear. Policymakers and regulators should prioritize the development of standardized ESG-reporting frameworks to enhance transparency and reduce information asymmetry. Investors may incorporate ESG metrics into risk assessment frameworks and view them not only as ethical indicators but also as predictive indicators of volatility. Corporate managers should conceptualize ESG engagement as a strategic instrument—rather than a mere reputational tool—for stabilizing performance and attracting long-term capital.

This study offered one of the first comprehensive examinations of how firm-level environmental and social performance influence stock volatility in an emerging-market context via a hybrid econometric-machine-learning approach. The findings indicated that the sustainability practices do not exert statistically significant average linear effects on equity risk; instead, their relevance emerges through nonlinear and time-dependent patterns captured by learning-based models. Environmental performance appeared more informative than social performance in learning-based models of volatility dynamics; however, neither dimension exhibited statistically significant average linear effects on equity risk, and the overall the ESG-risk nexus exhibited nonlinear and state-dependent patterns, which were explored through learning-based models rather than through explicit governance-based moderation tests. Theoretically, the findings are consistent with stakeholder, agency, and signaling perspectives by suggesting that sustainability-related disclosures may contribute to transparency and legitimacy, which learning-based models can exploit when modeling nonlinear and time-dependent volatility dynamics, without implying statistically significant average linear risk-reduction effects. Methodologically, this work advances the ESG literature by utilizing RL and LSTM architectures, which outperform conventional linear models in capturing temporal and non-linear features of financial volatility. However, this study is not without limitations. The reliance on manually coded ESG disclosures and the focus on a single emerging market limit generalizability. Future research may extend the present framework in several important directions. One promising avenue is to conduct industry-specific analyses in order to explore potential sectoral heterogeneity in the relationship between ESG performance and stock volatility, particularly in settings where environmental and social risks are unevenly distributed across industries. Such extensions would allow researchers to examine whether the nonlinear and state-dependent effects documented in this study vary systematically across sectors with different regulatory exposure, competitive dynamics, or sustainability pressures. In addition, future studies may broaden the empirical scope to cross-country settings,

explicitly incorporate governance dimensions as a separate analytical component, and employ standardized ESG ratings where data availability permits. Further research could also investigate asymmetric effects of sustainability performance during market downturns versus expansionary phases, as well as pursue causal identification through quasi-experimental or natural-experiment research designs.

REFERENCE

- Albuquerque, R., Koskinen, Y., & Zhang, C. (2019). Corporate social responsibility and firm risk: Theory and empirical evidence. *Management Science*, 65(10), 4451–4469. <https://doi.org/10.1287/mnsc.2018.3043>.
- Bae, K.-H., El Ghouli, S., Guedhami, O., Kwok, C., & Zheng, Y. (2017). Does corporate social responsibility reduce the costs of high leverage? Evidence from capital structure and product market interactions. *Journal of Banking and Finance*, 100, 135–150. <https://doi.org/10.1016/j.jbankfin.2018.11.007>.
- Biju, A. V. N. (2025). ESG-firm performance nexus: Evidence from an emerging economy. *Business Strategy and the Environment*. 34(3), 3469–3496. <https://doi.org/10.1002/bse.4152>.
- Broadstock, D. C., Matousek, R., Meyer, M. T., & Tzeremes, N.G. (2019). Does corporate social responsibility impact firms' innovation capacity? The indirect link between environmental & social governance implementation and innovation performance. *Journal of Business Research*, 119(2), 99–110. <https://doi.org/10.1016/j.jbusres.2019.07.014>.
- Chao, L., Mian, W., Xi, C., & Wenli, H. (2022). Environmental, social and governance performance, corporate transparency, and credit rating: Some evidence from Chinese A-share listed companies. *Pacific-Basin Finance Journal*, 74(C), 101806. <https://doi.org/10.1016/j.pacfin.2022.101806>.
- Chava, S. (2014). Environmental externalities and cost of capital. *Management Science*, 60(9), 2223–2247. <https://doi.org/10.1287/mnsc.2013.1863>.
- Dhaliwal, D. S., Li, O. Z., Tsang, A., & Yang, Y. G. (2011). Voluntary nonfinancial disclosure and the cost of equity capital: The initiation of corporate social responsibility reporting. *The Accounting Review*, 86(1), 59–100. <https://doi.org/10.2308/accr.00000005>.
- Donaldson, T., & Preston, L. E. (1995). The stakeholder theory of the corporation: Concepts, evidence, and implications. *Academy of Management Review*, 20(1), 65–91. <https://doi.org/10.5465/amr.1995.9503271992>.
- Edmans, A. (2011). Does the stock market fully value intangibles? Employee satisfaction and equity prices. *Journal of Financial Economics*, 101(3), 621–640. <https://doi.org/10.1016/j.jfineco.2011.03.021>.
- Farūq, A. T. M., & Chowdhury, M. A. R. (2025). Financial markets and ESG: How big data is transforming sustainable investing in developing countries. *arXiv*. <https://arxiv.org/abs/2503.06696>.
- Freeman, R. E. (1984). *Strategic management: A stakeholder approach*. Pitman.
- Gupta, H., & Chaudhary, R. (2023). An analysis of volatility and risk-adjusted returns of ESG indices in developed and emerging economies. *Risks*, 11(10), 182. <https://doi.org/10.3390/risks11100182>.
- Horn, M. (2023). The influence of ESG ratings on idiosyncratic stock risk: The unrated, the good, the bad, and the sinners. *Schmalenbach J Bus Res*, 75, 415–442.

- <https://doi.org/10.1007/s41471-023-00155-1>.
- Iannone, B., Duttalo, P., & Gattone, S. A. (2025). Evaluating the resilience of ESG investments in European markets during turmoil periods. *arXiv*. <https://arxiv.org/abs/2501.03269>.
- Jensen, M. C., & Meckling, W. H. (1976). Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics*, 3(4), 305-360. [https://doi.org/10.1016/0304-405X\(76\)90026-X](https://doi.org/10.1016/0304-405X(76)90026-X).
- Kansoy, F., & Stasiulaitis, D. (2025). Green shields: The role of ESG in uncertain time. *arXiv*. <https://arxiv.org/abs/2506.02143>.
- Kim, Y., Li, H., & Li, S. (2014). Corporate social responsibility and stock price crash risk. *Journal of Banking & Finance*, 43, 1–13. <https://doi.org/10.1016/j.jbankfin.2014.02.013>.
- Li, Y., Gong, Z., Zhang, X., & Koh, L. (2018). The impact of environmental, social, and governance disclosure on firm value: The role of CEO power. *British Accounting Review*, 50(1), 60–75. <https://doi.org/10.1016/j.bar.2017.09.007>.
- Lins, K. V., Servaes, H., & Tamayo, A. (2017). Social capital, trust, and firm performance: The value of corporate social responsibility during the financial crisis. *The Journal of Finance*, 72(4), 1785–1824. <https://doi.org/10.1111/jofi.12505>.
- Liu, D., Gu, K., & Hu, W. (2023). ESG performance and stock idiosyncratic volatility. *Finance Research Letters*, 58(B), 104393. <https://doi.org/10.1016/j.frl.2023.104393>.
- Naseer, M. M., Guo, Y., Bagh, T., & Zhu, X. (2024). Sustainable investments in volatile times: Nexus of climate change risk, ESG practices, and market outcomes. *International Review of Financial Analysis*, 95(B), 103492. <https://doi.org/10.1016/j.irfa.2024.103492>.
- Rahat, B., & Nguyen, P. (2024). The impact of ESG profile on firm's valuation in emerging markets. *International Review of Financial Analysis*, 95(A), 103361. <https://doi.org/10.1016/j.irfa.2024.103361>.
- Spence, M. (1973). Job market signaling. *The Quarterly Journal of Economics*, 87(3), 355-374. <https://doi.org/10.2307/1882010>.
- Valencia Söderberg, D., & Truong, M. (2024). *ESG & emerging markets – A volatility perspective of ESG investment in emerging markets*. Linköping University.
- Wang, H., Shen, H., & Li, S. (2023). ESG performance and stock price fragility. *Finance Research Letters*, 56, 104101. <https://doi.org/10.1016/j.frl.2023.104101>.
- Wang, Y., & Sonenshine, R. (2025). The nonlinear impact of environmental, social, governance on stock market performance among US manufacturing and banking firms. *Journal of Risk and Financial Management*, 18(6), 293. <https://doi.org/10.3390/jrfm18060293>.