Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Explainable Ensemble Forecasting of Corporate Green Innovation for Managerial Decision Support

Huan Chen

chenhuan48@mails.gdut.edu.cn

Guangdong University of Technology

Research Article

Keywords: ensemble forecasting, computational management, decision support, explainable machine learning, information value, green innovation

Posted Date: November 19th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11544795/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Explainable Ensemble Forecasting of Corporate Green Innovation for Managerial Decision Support

Huan Chen

Abstract

Combining heterogeneous firm information can improve prediction without necessarily changing managerial review priorities. We develop an explainable ensemble framework integrating innovation history, financial resources, artificial intelligence (AI) assets and inventions, and climate-related disclosures to forecast green invention applications among Chinese nonfinancial listed firms. Expanding historical windows separate training outcomes from later forecast cohorts, and historical out-of-fold predictions train a nonnegative ridge stacking layer. Across 10,447 test observations with outcome years 2021–2023, stacking achieves a coefficient of determination of 0.635 and root mean squared error of 0.7297, compared with 0.612 and 0.7525 for a financial ridge benchmark. Matched ensembles achieve root mean squared errors of 0.7422 with a common information core, 0.7381 after adding climate language, and 0.7356 after adding AI information. Each block contributes conditional predictive information. A ridge model using all 16 predictors attains root mean squared error of 0.7527. Feature attribution identifies innovation history and AI inventions as prominent predictors, while accumulated local effects reveal a rising and flattening software response. XGBoost slightly improves on stacking under squared error, whereas persistence achieves the lowest mean absolute error. Continuous forecasting gains produce only small changes in fixed-capacity review rankings. The study provides an empirical computational framework distinguishing information value, functional flexibility, explanation, and review objectives in managerial forecasting.

Keywords: ensemble forecasting; computational management; decision support; explainable machine learning; information value; green innovation

Mathematics Subject Classification (2020): 62M20; 62P20; 68T05

1. Introduction

Managers frequently combine heterogeneous corporate records to anticipate outcomes and select firms for further review. Additional information and flexible models may improve continuous predictions while leaving shortlist priorities almost unchanged. Corporate green invention provides a demanding setting for examining this distinction: inventive activity is persistent, many firms record no applications, and financial, technological, and disclosure signals have different meanings. We evaluate whether their integration improves later forecasts and whether those gains change priorities under limited review capacity.

Climate information matters for investment and risk assessment (Giglio et al., 2021; Krueger et al., 2020), while carbon exposure is associated with financing conditions (Bolton & Kacperczyk, 2021, 2023). Firms with similar climate concerns can nevertheless differ in research history, liquidity, and technology. Forecasting requires identifying useful combinations across these dimensions.

Green invention combines persistent activity among established innovators with many zero observations. Forecasts must represent accumulated capacity and distinguish subsequent activity. Innovation history is therefore an essential benchmark against which additional financial, technological, and disclosure signals should be evaluated under a common temporal protocol.

China's A-share firms provide linked financial statements, annual reports, technology investment records, and patents. Differences in inventive history, industries, and environmental conditions permit evaluation of multidimensional information across later forecast cohorts.

Financial statements describe resources, narratives describe attention, and patents record inventive knowledge. Their decision value depends on temporal alignment, functional representation, and the task being evaluated. Following Booker et al. (2024), we distinguish continuous forecasting from fixed-capacity review. This connects the construction of a computational model to the managerial use of its output. Green applications measure recorded inventive activity; the prioritization exercise evaluates its concentration among selected firms.

Climate language distinguishes managerial attention from independently observed exposure. Sautner et al. (2023) separate threats, regulatory exposure, and opportunities. Broad environmental narratives may signal strategic orientation, whereas hazard terms describe narrower concerns. We evaluate these representations separately to identify where predictive information resides.

Software and hardware describe recorded resources; AI inventions describe technical engagement; AI language describes disclosure. Capability research emphasizes complementary resources (Mikalef & Gupta, 2021; Enholm et al., 2022). Building on AI–green innovation studies (Chen et al., 2025; Dong et al., 2025; Zhong & Song, 2025; Zhong et al., 2025), we evaluate each dimension beyond existing inventive activity.

Three gaps motivate the analysis: associations among observed variables do not establish later forecasting performance; aggregate technology indicators conceal distinct information sources; and algorithm comparisons can confound richer inputs with greater flexibility. Moreover, lower continuous error may leave review rankings largely unchanged. Matched information sets and common-information benchmarks allow these questions to be evaluated separately under a chronological protocol.

We ask whether climate and AI information improve subsequent green invention forecasts separately and jointly, which dimensions dominate fitted predictions, and what nonlinear patterns characterize those predictions. Fixed-capacity review evaluates prioritization, while auxiliary panel regressions describe conditional associations.

The framework stacks XGBoost, CatBoost, and LightGBM using historical out-of-fold predictions and a nonnegative ridge meta-model. Matched information sets assess incremental value. A full-information ridge benchmark assesses functional flexibility. Feature attribution, accumulated local effects, and observed-group summaries interpret predictions; fixed review capacities evaluate prioritization.

The study contributes an empirical computational framework for managerial forecasting in three respects. First, matched information sets separate the conditional predictive value of climate and AI records from the shared financial and historical core. Second, a common-information ridge comparison distinguishes the tested nonlinear representation from the value of additional predictors. Weighted feature attribution reconciles individual ensemble predictions, while accumulated local effects describe supported nonlinear responses. Third, fixed-capacity review connects forecasting performance to a managerial prioritization task. The results show why information selection, model complexity, and explanation should be evaluated against both continuous losses and the decisions supported by forecast rankings.

Section 2 develops the foundations and questions. Section 3 presents data and computational methods. Sections 4 and 5 report forecasts, explanations, and robustness. Section 6 discusses managerial decision support; Sections 7 and 8 conclude and identify limitations. Appendix A reports auxiliary panel evidence.

2. Conceptual Foundations and Research Questions

2.1 Climate information and the task of forecasting green innovation

Physical threats may disrupt research, whereas environmental opportunities may encourage technological development. Differing climate–innovation relationships (Lee et al., 2025; Tian et al., 2024; Ren et al., 2024) motivate separate textual representations rather than a single signed expectation.

Annual reports combine attention, reporting conventions, and strategic communication. Report length and firm characteristics situate climate language within this context. Separating physical hazards from broader environmental terms tests whether predictive value concerns narrow risk language or wider strategic orientation.

RQ1. Do climate-related disclosures and multidimensional AI information improve subsequent green innovation forecasts separately and jointly under a common temporal evaluation protocol?

2.2 Financial position as predictive information

Financing-friction theory motivates financial predictors because uncertain research requires sustained expenditure. Financial constraints influence environmental allocation (Xu & Kim, 2022), but their relationship with green innovation varies across dimensions (Nguyen & Do, 2025). We therefore retain several indicators of available resources and near-term commitments.

Current liabilities less monetary funds, scaled by assets, measures near-term claims net of cash. Cash/assets and backward-looking operating cash-flow variability capture available resources and funding stability. Their predictive relevance is assessed conditional on inventive history and technology records.

Cash can represent preparedness, profitability, or uncommitted resources; liabilities can reflect strain or operating scale. Research on uncertainty and cash policy (Duong et al., 2020; Javadi et al., 2023) therefore motivates flexible prediction rather than a predetermined negative relationship.

2.3 Multidimensional AI information and accumulated inventive resources

The resource-based view motivates distinguishing technological assets, knowledge, and disclosure. Complementary resources underpin AI capability (Mikalef & Gupta, 2021; Enholm et al., 2022). Prediction retains the separate dimensions; the auxiliary panel uses an aggregate measure for a parsimonious comparison.

AI inventions can indicate technical experience and knowledge-combination routines relevant to subsequent green invention. Digital innovation and AI-assisted search research support this interpretation (Nambisan et al., 2019; Muhlroth & Grottke, 2022). We evaluate their predictive content alongside assets and disclosure.

Dynamic capabilities concern sensing, seizing, and reconfiguration (Teece, 2018). Digital transformation also involves organizational change (Vial, 2019). These perspectives motivate flexible prediction of resource configurations: similar financial or technological amounts may accompany different inventive outcomes depending on complementary characteristics.

RQ2. What predictive importance do liquidity indicators and individual AI information dimensions have relative to prior green inventive activity, conventional firm characteristics, and climate-related disclosures?

2.4 Nonlinear configurations and heterogeneous prediction patterns

Software may differentiate firms establishing basic infrastructure more strongly than firms expanding extensive resources. Its predictive role may also depend on inventive history and financial capacity. Nonlinear learners accommodate such configurations without specifying every interaction beforehand.

SHAP attributes fitted predictions to features (Lundberg et al., 2020); accumulated local effects summarize supported model responses (Apley & Zhu, 2020). Observed-group comparisons retain differences in firm composition. These tools describe model behavior and configurations rather than structural interaction effects.

RQ3. What nonlinear response patterns and differences across observed climate-attention and software-intensity groups characterize the fitted green innovation forecasts?

Figure 1 links financial-resource and capability perspectives to feature selection, temporal evaluation, interpretation, and managerial review. Continuous forecasts and fixed-capacity rankings evaluate different uses of the same model. Arrows describe information processing and evaluation stages rather than causal relationships among corporate outcomes.

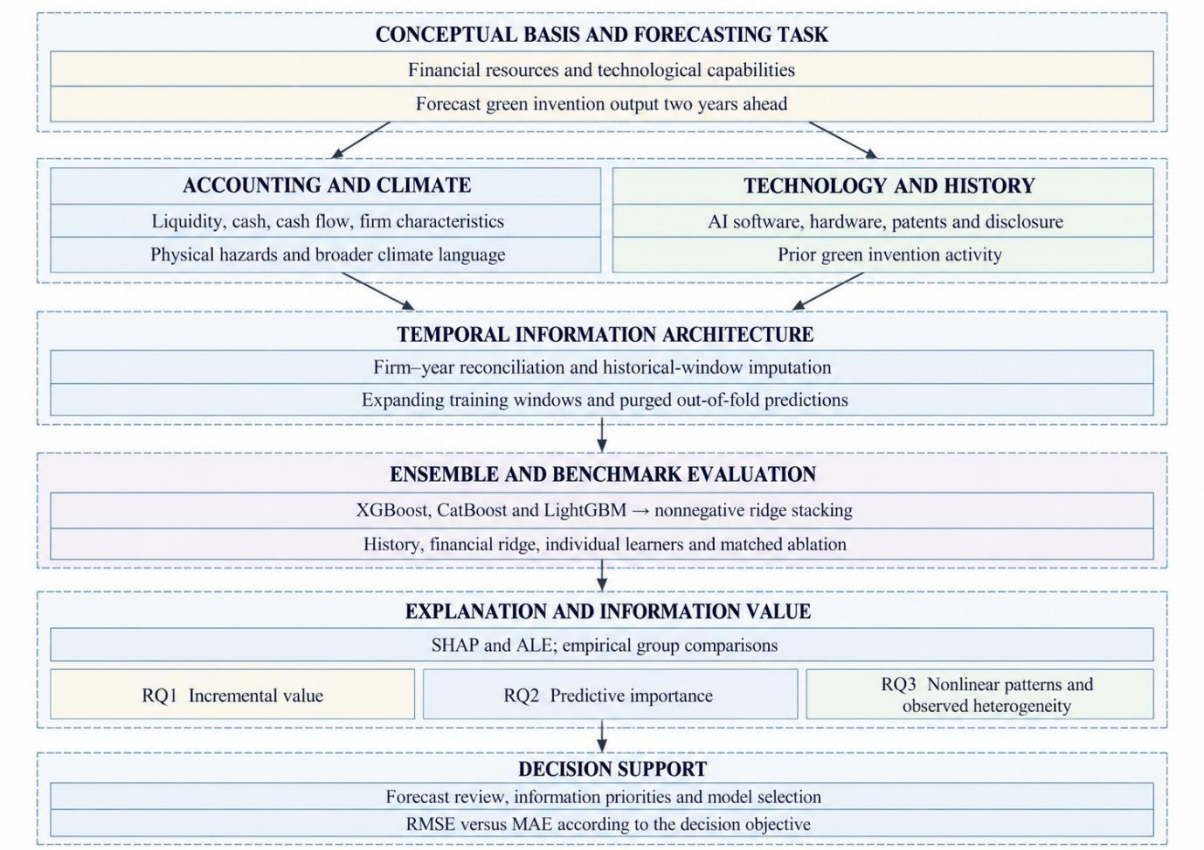


Figure 1. Conceptual framework for explainable green innovation forecasting. Financial-resource and capability perspectives motivate the information dimensions. Firm–year integration and temporal separation connect these inputs to the forecasting task. The ensemble and explanation layers support assessment of information value and fitted heterogeneity. Arrows represent information processing and evaluation relationships. The framework was prepared programmatically with assistance from ChatGPT (GPT-6; OpenAI).

3. Research Design and Predictive Framework

3.1 Data sources and panel construction

Financial statements and AI investment records are obtained from CSMAR; AI patents originate from CNRDS. The panel links balance-sheet resources and liabilities, revenue, profit, research expenditure, operating cash flow, annual-report keywords, and green applications through consistent firm–year identifiers.

Green applications use raw group-total, post-listing records. Independent and joint invention applications are added; utility models are constructed similarly. Parent records are not added to group totals, and each source must have a unique firm–year key.

AI investment retains accounting category 1, with software and hardware separate; category 2 is excluded to preserve measurement consistency. Unreported investments remain missing. AI inventions

combine independent and joint group-total, post-listing applications. AI word frequency remains a separate disclosure measure.

Climate records provide report identifiers, dates, word totals, and term counts. Security codes, fiscal years, and dates are parsed separately; the earliest dated firm–year record is retained. Duplicate dictionary columns are excluded, and the hazard subset remains fixed. Evaluation uses archival extracts.

The sample retains consolidated nonfinancial statements with positive assets and excludes available ST/PT names. Forecasting requires future innovation, prior innovation, and hazard attention; other missing predictors are handled within training windows. Auxiliary estimation additionally requires complete later liquidity and covariates. Documented patent zeros remain distinct from missing records.

Raw green applications contain 76,250 records (1990–2025), climate text 67,944 (2000–2024), and selected AI investment 57,340 (2007–2025). Financial–climate merging yields 59,921 records (2005–2024). Forecast eligibility retains 34,618 observations from 4,097 firms (feature years 2008–2021). Evaluation covers feature years 2019–2021 and outcomes 2021–2023. The separate complete-case panel contains 10,484 observations from 1,648 firms, with outcomes 2010–2023, after singleton removal.

3.2 Temporal alignment and outcome definition

Features from t or earlier predict green inventions in t+2. AI resources and prior invention refer to t−1; forecasting liquidity refers to t. Auxiliary t+1 liquidity never enters prediction. Explicit firm–year matching requires exactly one calendar year for annual lags, preventing gaps from being treated as annual transitions.

In Equation (1), GInv is the group-total green invention application count. The target ln(1 + GInv) retains zeros and limits prolific applicants' influence. Supplementary panel outcomes combine inventions and utility models or indicate any invention, distinguishing activity intensity from participation.

$$\mathrm{GI}_{i,t+2} = \ln\left(1 + \mathrm{GInv}_{i,t+2}\right) \qquad (1)$$

3.3 Climate-language and financial-position features

In Equation (2), n is term k's count, W total words, and P the fixed hazard dictionary: flooding, drought, storms, typhoons, freezing, extreme precipitation, and water scarcity. Counts per 10,000 words enter forecasting as ln(1 + frequency). Auxiliary equations standardize separately. Broader environmental and transition terms provide an alternative disclosure representation.

$$\mathrm{R}_{it} = \ln\left(1 + 10{,}000\,\frac{\sum_{k\in\mathcal{P}} n_{kit}}{W_{it}}\right) \qquad (2)$$

Equation (3) defines liquidity pressure as (current liabilities − monetary funds)/total assets. CL, Cash, and A denote those components. Higher values indicate greater near-term claims net of recorded

cash. Cash/assets and net current assets/assets describe complementary resources; these accounting measures do not directly measure external credit access.

$$\mathrm{LP}_{it} = \frac{\mathrm{CL}_{it} - \mathrm{Cash}_{it}}{A_{it}} \qquad (3)$$

Operating cash-flow variability is the sample standard deviation of operating cash flow/assets over the current and preceding two observed records. Three observations are required, but consecutive calendar years are not. Annual lags and future outcomes, by contrast, require explicit calendar-year matching.

3.4 Multidimensional AI information and firm characteristics

AI software and hardware intensities equal ln(1 + 100 × category 1 amount/assets). AI invention activity equals ln(1 + independent and joint applications); AI disclosure uses source-panel lagged word frequency. These remain separate predictors. Appendix A's equal-weight standardized composite is confined to auxiliary correlations.

Table 1 defines the 16 predictors and their timing. Other characteristics cover size, leverage, profitability, cash flow, tangibility, report length, and invention history. Auxiliary full-panel winsorization and standardization never enter forecasting, which uses historical-window preparation and preserves documented zeros.

Table 1. Forecasting variables and feature-year timing

Variable	Definition	Timing
Green invention target	ln(1 + independent and joint green invention applications), listed group after listing	t+2
Prior green invention	Same logarithmic application measure in the preceding year	t−1
Physical-hazard attention	ln(1 + physical-hazard keyword frequency per 10,000 report words)	t
Broad climate language	ln(1 + broad environmental and transition keyword frequency per 10,000 words)	t
AI software intensity	ln(1 + 100 × category 1 software amount / total assets)	t−1
AI hardware intensity	ln(1 + 100 × category 1 hardware amount / total assets)	t−1
AI invention activity	ln(1 + independent and joint AI invention applications)	t−1
AI disclosure	AI word-frequency measure from the annual-report source panel	t−1
Liquidity pressure	(Current liabilities − monetary funds) / total assets	t
Cash	Monetary funds / total assets	t
Operating cash flow	Net operating cash flow / total assets	t
Cash-flow variability	Three-observation sample SD of operating cash flow / total assets	Through t
Firm size	ln(total assets)	t
Leverage	Total liabilities / total assets	t
Profitability	Net profit / total assets	t
Asset tangibility	Net fixed assets / total assets	t

Variable	Definition	Timing
Report length	ln(total report words)	t

Note. The model uses 16 predictors and one target. The standardized AI composite and t+1 liquidity are confined to auxiliary panel analyses. Missing source investments are not recoded as documented zeros. Fiscal feature-year alignment does not establish original publication-vintage availability.

3.5 Temporal evaluation and ensemble construction

The feature vector combines prior invention; hazard and broad climate language; lagged AI software, hardware, patents, and disclosure; liquidity pressure, cash, operating cash flow and its variability; size, leverage, profitability, tangibility, and report length. Firm identifiers link records and define resampling clusters but are excluded as predictors.

At feature-year origin v, training feature years end at v−3, so training targets end at v−1. The same separation governs meta-model historical predictions. Checking both feature and target dates prevents overlapping outcomes from leaking into evaluation (Kapoor & Narayanan, 2023).

Expanding windows produce out-of-fold cohorts for feature years 2012–2021. Earlier cohorts train the meta-model; evaluation uses 2019–2021 with outcomes 2021–2023. Maximum training feature years are respectively 2016, 2017, and 2018. Later training admits earlier test records only after the same temporal separation.

Each ensemble's median imputer is fitted within the historical training window; financial ridge uses training means. Missing investments remain missing in source records. Auxiliary full-panel transformations are excluded, preserving origin-specific preparation.

XGBoost uses 180 trees, depth 3, learning rate 0.04, subsample/column fractions 0.85, and L2 penalty 10. CatBoost uses 220 iterations, depth 4, rate 0.04, and L2 10. LightGBM uses 180 trees, 15 leaves, minimum leaf count 60, rate 0.04, and L2 10. All use seed 42; settings are fixed without hyperparameter search.

Equation (4) combines origin-specific learners f using the 16-feature vector Z. The ridge meta-model has an intercept, penalty 10, and nonnegative coefficients fitted on admissible historical out-of-fold predictions. Weights are regularized coefficients, not probabilities, and need not sum to one.

$$\widehat{\mathrm{GI}}_{i,t+2} = \beta_0 + \sum_{m=1}^{3} \omega_m \hat{f}_m^{past}(Z_{it}), \quad \omega_m \geq 0 \qquad (4)$$

Computations used Python 3.12.14 with scikit-learn 1.8.0, XGBoost 3.4.1, CatBoost 1.2.10 and LightGBM 4.7.0. ChatGPT (GPT-6; OpenAI) assisted with development of the Python analysis and plotting code. Model estimation, numerical validation and data visualizations were produced by executing the specified computational procedures on the assembled data. The reproducibility package records the software environment, fixed settings and validation checks.

3.6 Benchmarks and measures of information value

Benchmarks comprise the historical training mean, prior-invention persistence, and financial ridge. Ridge combines invention history with size, leverage, profitability, cash flow, tangibility, and liquidity pressure, using penalty 10 and historical mean imputation.

The ten-feature common core excludes two climate and four AI predictors. Matched models add climate, AI, or both, holding sample, learner settings, ordering, imputation, historical cutoffs, and stacking protocol constant.

Ablations remove AI patents or narrow hazard language. Full-information ridge uses all 16 predictors, median imputation, penalty 10, and identical cutoffs, isolating functional representation. Seven-feature financial ridge supplies the conventional accounting benchmark.

Annual review capacities select the highest-predicted 10% or 20% of firms, rounded upward; firm codes resolve ties. Performance is the selected share of realized invention applications, recovered from observed logarithmic outcomes.

Equation (5) defines root mean squared error (RMSE), mean absolute error (MAE), and pooled test R^2 for N observations and residuals e. Errors use logarithmic outcomes. The observed test mean defines the descriptive R^2 denominator; it is not an available forecast.

$$\mathrm{RMSE} = \sqrt{N^{-1}\sum_i e_i^2}, \quad \mathrm{MAE} = N^{-1}\sum_i |e_i|, \quad R^2 = 1 - \frac{\sum_i e_i^2}{\sum_i (y_i - \bar{y})^2} \qquad (5)$$

Annual RMSE evaluates cohort consistency. Paired firm-block bootstraps with 1,999 draws retain each firm's test observations together. Percentile intervals are pointwise, unadjusted for multiple comparisons, and conditional on fixed fitted models and predictions rather than repeated estimation.

3.7 Global attribution and nonlinear model interpretation

Interpretation uses 800 randomly selected feature-year 2021 test observations (seed 42). Base-learner TreeSHAP values are weighted by ridge coefficients; the intercept enters the baseline. Equation (6), with contributions phi and p predictors, reconstructs the additive ensemble prediction.

$$\widehat{\mathrm{GI}}_i = \phi_{0i} + \sum_{j=1}^{p} \phi_{ij}, \quad I_j = N_s^{-1}\sum_{i=1}^{N_s} \left|\phi_{ij}\right| \qquad (6)$$

Following Rudin (2019), explanations accompany transparent benchmarks. Mean absolute SHAP summarizes model reliance. A local case is selected mechanically at median absolute forecast error, avoiding an unusually successful illustration. Local contributions explain one prediction; global importance summarizes the interpretation sample.

Accumulated local effects (ALE) follow Apley and Zhu (2020). Quantile boundaries are deduplicated; predictions at occupied intervals' bounds retain other observed features. Mean differences are accumulated and centered by interval counts. Sparse support determines the displayed interval count.

Observed software records are split at the preceding training-window median and grouped by hazard-attention bins. Mean fitted outcomes retain actual firm composition. These comparisons describe configurations rather than an isolated software–climate interaction.

4. Forecasting Results and Model Interpretation

4.1 Forecast cohorts and the importance of historical benchmarks

Evaluation includes 10,447 observations from 3,919 firms (Table 2). Mean logarithmic invention output is 0.8779; 51.35% have no applications. This zero mass and skewed positive tail motivate reporting both squared and absolute errors.

Table 2. Rolling forecast cohorts and training information boundaries

Feature year	Target year	Last training feature year	Last training target year	Base training N	Meta training N	Test N
2019	2021	2016	2018	18,057	12,005	3,283
2020	2022	2017	2019	20,916	14,864	3,381
2021	2023	2018	2020	24,171	18,119	3,783

Note. Base training uses an expanding historical window. Meta training uses earlier out-of-fold predictions subject to the same three-year feature separation. Counts refer to the eligible observations in each forecast cohort.

4.2 Comparative forecasting performance

Table 3 shows persistence R^2 of 0.555 and financial ridge R^2 of 0.612 with RMSE 0.7525. Prior invention and financial characteristics thus provide strong comparisons beyond the weak historical-mean forecast.

Stacking achieves R^2 0.635 and RMSE 0.7297, approximately 3.0% below financial ridge. The bootstrapped MSE gain is 0.03377 [0.02483, 0.04326]. Annual stacking RMSE is 0.7409, 0.7253, and 0.7238, versus financial ridge's 0.7586, 0.7548, and 0.7450.

XGBoost slightly improves pooled RMSE to 0.7286 (R^2 0.636). Persistence has the lowest MAE, 0.5005, versus stacking's 0.5601. With many zero outcomes, a model emphasizing larger deviations can rank differently under absolute loss. Selection therefore depends on which errors matter to the monitoring task.

The common-core ensemble achieves R^2 0.623 and RMSE 0.7422. Adding climate and AI jointly raises R^2 by 0.0127 and lowers RMSE by approximately 1.7%. Section 5 separates information increments from functional flexibility; SHAP describes fitted attribution.

Table 3. Out-of-sample forecasting performance

Model	RMSE	MAE	R^2
Historical mean	1.2544	0.8916	-0.0779
Patent persistence	0.8057	0.5005	0.5553
Financial ridge	0.7525	0.5566	0.6121
XGBoost	0.7286	0.5580	0.6364
CatBoost	0.7315	0.5575	0.6335
LightGBM	0.7317	0.5573	0.6332
Temporal stacking	0.7297	0.5601	0.6352
Stacking without climate and AI	0.7422	0.5802	0.6226

Note. N = 10,447. Feature years are 2019–2021 and outcome years 2021–2023. The reduced ensemble removes all climate and AI features. The same test records are used for every comparison. Lower RMSE and MAE indicate better forecasts.

4.3 Predictive importance of financial and technological information

Figure 2 and Table 4 show final-origin mean absolute SHAP importance: prior invention 0.4846, AI inventions 0.1444, size 0.1182, and broad environmental language 0.0758. Profitability, report length, tangibility, and AI disclosure follow. Broader language may indicate disclosed technological orientation rather than physical exposure.

Liquidity pressure contributes 0.0185, software 0.0155, and cash 0.0122. Hardware, operating cash flow, variability, and hazard language contribute less; hazard importance is 0.000488. The fitted model thus relies more on inventive knowledge than recorded hardware or narrow hazard attention.

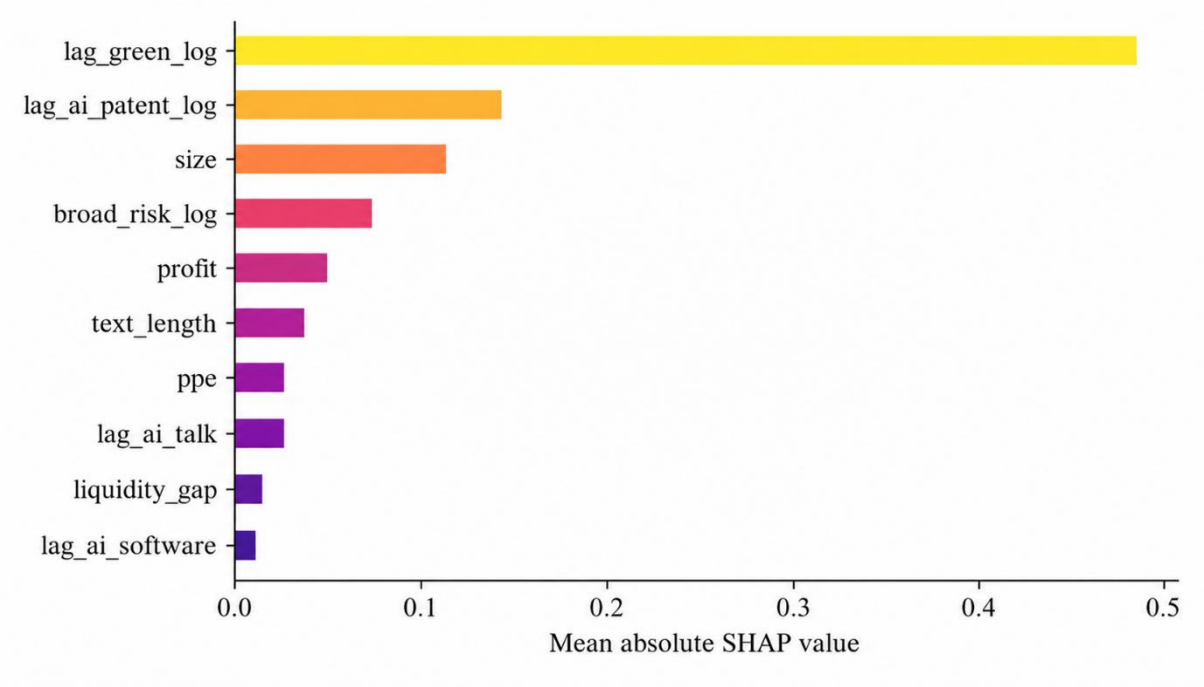


Figure 2. Global SHAP importance for the final temporal stacking model. The display shows the ten largest mean absolute contributions among 16 predictors, calculated on 800 randomly selected observations in the feature-year 2021 test cohort. Contributions from the three base learners are weighted by the ridge meta-model coefficients. Values are on the logarithmic forecast scale. Table 4 reports the complete ranking.

Table 4. Complete final-origin predictive importance ranking

Rank	Feature	Mean absolute SHAP
1	Prior green invention	0.484634
2	AI invention activity	0.144364
3	Firm size	0.118222
4	Broad climate language	0.075830
5	Profitability	0.053012
6	Report length	0.041093
7	Asset tangibility	0.032561
8	AI disclosure	0.031882
9	Liquidity pressure	0.018529
10	AI software intensity	0.015483
11	Cash	0.012158
12	Leverage	0.006870
13	AI hardware intensity	0.004275
14	Operating cash flow	0.004253
15	Cash-flow variability	0.003651
16	Physical-hazard attention	0.000488

Note. Final-origin ensemble, fixed sample of 800 observations. Values describe average absolute contributions to logarithmic predictions, not causal effects or isolated feature-removal gains.

4.4 Local explanation of an individual forecast

The median-error example is firm 603725 in feature year 2021: observed logarithmic outcome 0, forecast 0.4293, baseline 0.5621. Prior invention contributes −0.2100, AI patents −0.0572, and broad language +0.1680. Other financial and disclosure signals complete the prediction.

The positive forecast despite a zero outcome illustrates competing signals. Explanations guide examination of contributing records but do not convert the output into a patenting probability or expected count. Weighted SHAP reconstructs predictions with maximum discrepancy below 0.000003.

4.5 Nonlinear responses and heterogeneous prediction patterns

Physical-hazard ALE values are approximately 0.00023, −0.00021, and −0.00115 (Figure 3). Duplicate boundaries leave three supported intervals. This narrow range agrees with low SHAP importance and limited final-origin model reliance.

Software ALE rises from approximately −0.0315 to +0.00965, then flattens. Differentiation is strongest at low intensity, consistent with a basic technology resource interpretation. These fitted responses do not identify investment returns or a stable threshold.

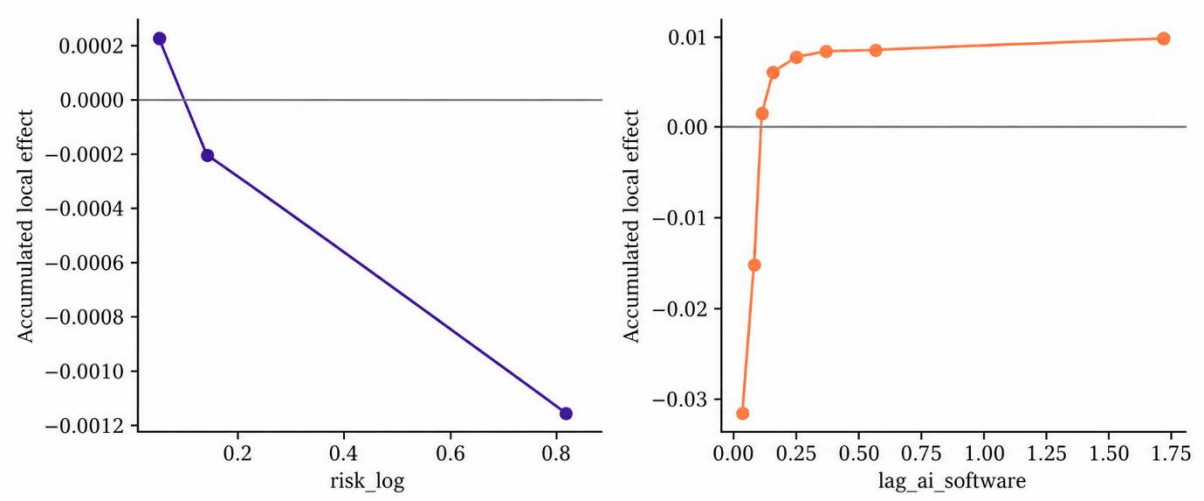


Figure 3. Accumulated local effects for physical-hazard attention and lagged AI software intensity. The final-origin model is evaluated on the fixed interpretation sample. Duplicate quantile boundaries are collapsed, yielding three occupied physical-hazard intervals and eight software intervals. Software intensity is ln(1 + 100 × software amount / total assets). Curves summarize fitted local responses; no cross-origin threshold test or threshold confidence interval is reported.

Figure 4 and Table 5 report fitted group means. Higher/lower software groups score 1.0178/0.8134 in the lower-attention bin and 1.1681/1.0292 in the higher bin. The software gap narrows at higher attention; actual group composition remains represented.

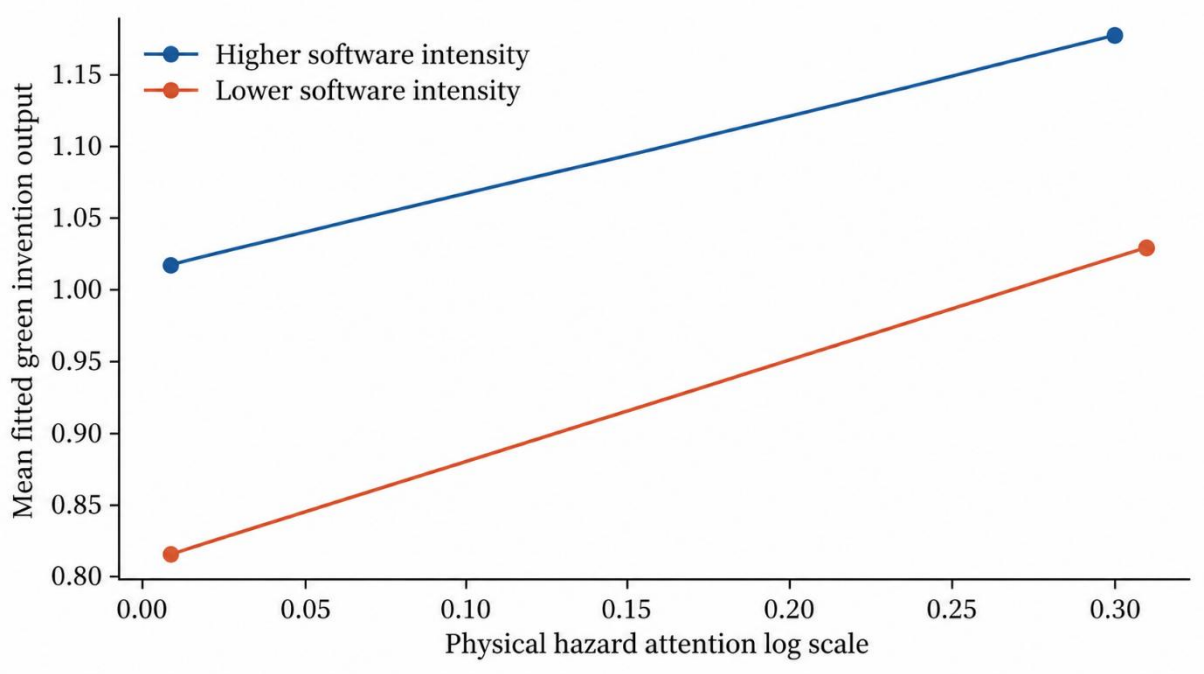


Figure 4. Conditional dependence of fitted green invention output across physical-hazard attention bins and observed software-intensity groups. The software split uses the preceding training-window median. Only records with observed software enter the display. Points are group averages of fitted predictions, retaining each group's other characteristics. The lines connect descriptive group summaries and do not isolate a climate–software interaction holding the remaining features constant.

Table 5. Observed climate-attention and software groups

Attention bin	Software group	Mean attention	Mean fitted output	N
Lower	Higher	0.0089	1.0178	2,007
Lower	Lower	0.0102	0.8134	656
Higher	Higher	0.2858	1.1681	471
Higher	Lower	0.3149	1.0292	195

Note. The final-origin descriptive display includes 3,329 observations with observed software intensity. Higher and lower software are defined by the training-window median. Attention uses the logarithmic physical-hazard frequency. Group composition is retained.

ALE assesses local model sensitivity; group averages compare firms with different histories and resources. Higher group predictions can therefore coexist with minimal hazard sensitivity. This distinction guides review toward a specific input or toward broader firm characteristics.

5. Predictive Robustness and Review Prioritization

5.1 Conditional information value and functional flexibility

Table 6 gives RMSE 0.7422 for the core, 0.7381 with climate, 0.7356 with AI, and 0.7297 with both. Each block improves squared-error prediction in every cohort. Conditional MSE gains are 0.00859 [0.00547, 0.01166] for climate and 0.01241 [0.00694, 0.01788] for AI (Table 7).

Table 6. Matched information sets and full-information ridge

Specification	RMSE	MAE	R^2
Common core	0.7422	0.5802	0.6226
Core plus climate	0.7381	0.5754	0.6267
Core plus AI	0.7356	0.5682	0.6294
Full stacking	0.7297	0.5601	0.6352
Full without AI patents	0.7334	0.5685	0.6315
Full without physical hazards	0.7297	0.5599	0.6352
Full-information ridge	0.7527	0.5690	0.6119

Note. N = 10,447 for every model. The core contains ten predictors; climate and AI add two and four predictors, respectively. Ensembles use identical learner specifications and temporal cutoffs. Full-information ridge uses all 16 predictors and training-window median imputation.

Table 7. Paired predictive increments under common temporal evaluation

Comparison	MSE gain	95% lower	95% upper
Climate added to core	0.00606	0.00352	0.00857
AI added to core	0.00988	0.00460	0.01522
AI added after climate	0.01241	0.00694	0.01788
Climate added after AI	0.00859	0.00547	0.01166
AI patents added to full-minus-patents	0.00543	0.00076	0.00976
Physical hazards added to full-minus-hazards	0.00002	-0.00079	0.00092
Stacking versus full-information ridge	0.03404	0.02765	0.04064

Note. Positive MSE gains favor the added information or stacking. Pointwise percentile intervals use 1,999 paired firm-block bootstrap draws on fixed test predictions. Intervals are not adjusted for multiple comparisons.

Removing AI patents yields RMSE 0.7334, although their final-cohort increment is negative. Removing hazard language yields 0.72970, with an MSE interval spanning zero. Full-information ridge achieves 0.7527; stacking's gain interval is [0.02765, 0.04064]. Broader language and nonlinear representation drive the evaluated improvements.

5.2 Forecasts under fixed review capacities

At 10% capacity, 1,047 selected observations account for 76.29% of applications under stacking, 75.64% under financial ridge, 75.04% under persistence, and 76.19% under the core (Table 8). The core difference is 84 applications or 0.10 percentage points. At 20%, stacking/core shares are 86.70%/86.63%.

Table 8. Retrospective concentration of innovation under fixed review capacities

Ranking model	10% application share	20% application share	10% positive share
Full stacking	76.29	86.70	97.13
Common core	76.19	86.63	97.13
Core plus climate	76.31	86.65	97.23
Core plus AI	76.43	86.70	97.04
Financial ridge	75.64	86.35	96.75
XGBoost	76.32	86.67	97.23
Patent persistence	75.04	85.92	96.28

Note. Application share is the percentage of 87,495 realized green invention applications attributable to selected firm–years. Positive share is the percentage of selected observations with at least one application. Capacities select 1,047 and 2,091 observations across the three cohorts after annual upward rounding. Results describe quantity concentration, not invention quality, classification calibration, review costs, or realized intervention benefits.

5.3 Auxiliary panel correlations and alternative accounting measures

Auxiliary firm/year-effect regressions yield an AI coefficient of 0.0732 (firm-clustered standard error 0.0226) and hazard coefficient 0.0041 (0.0081). Extended AI is 0.0648. Appendix A documents these complete-case conditional associations.

The conditional coefficient-product intervals all contain zero. They provide no support for assigning the ensemble's predictive gain to a sequential financial-transmission pathway. Financial indicators remain economically interpretable features whose predictive usefulness is evaluated directly in the forecasting design.

Auxiliary cash, working-capital, and variability comparisons retain weak hazard and hazard–AI associations, consistent with the forecasting diagnostics.

5.4 Alternative innovation outcomes and climate-language definitions

With inventions plus utility models, the auxiliary AI coefficient remains approximately 0.0604; hazard and interaction estimates remain weak. AI's participation coefficient is smaller and not conventionally distinguishable from zero. These regression checks distinguish intensity from participation without validating alternative forecasts.

Broad language yields auxiliary innovation and AI-interaction coefficients 0.0689 and 0.0720 ($p < 0.001$). Adding later liquidity yields 0.0689 and 0.0715; the broad-language liquidity coefficient is −0.00112 and weak. This supports separating wider environmental orientation from narrow hazard attention.

6. Discussion

6.1 Separating information value from computational representation

Climate and AI signals each contribute conditional information, while the nonlinear ensemble improves squared-error performance relative to ridge using the same predictors. These comparisons

separate input value from the tested functional representation. Fixed-capacity review changes much less. Computational evaluation therefore needs to distinguish population-wide forecast accuracy from the priorities induced by a model under a specified review rule.

Financial resources, technology records, and narratives support anticipatory assessment. Strategic foresight links anticipation to capability development (Haarhaus & Liening, 2020), while computational search supports identification of technological developments (Muhlroth & Grottke, 2022). Following task alignment (Booker et al., 2024), the framework makes information selection, chronological preparation, benchmark comparison, and explanation explicit. These components provide an auditable route from heterogeneous company records to managerial forecasts.

6.2 Why inventive knowledge predicts more strongly than some asset measures

AI inventions have greater fitted importance than hardware, consistent with capability research (Mikalef & Gupta, 2021; Enholm et al., 2022). Their removal worsens pooled prediction but not every cohort. Remaining technology records still contribute, supporting separate dimensions rather than equating attribution with stable incremental performance.

Digital transformation and AI-assisted search involve organizational knowledge and innovation processes (Nambisan et al., 2019; Vial, 2019; Muhlroth & Grottke, 2022). Shared technical capabilities may help explain AI inventions' predictive relevance, extending AI–green innovation evidence (Chen et al., 2025; Dong et al., 2025; Zhong & Song, 2025) toward subsequent forecasting.

Software's flattening response is compatible with greater differentiation during initial resource accumulation. Capabilities and foresight concern how resources are organized (Teece, 2018; Haarhaus & Liening, 2020). The response motivates configuration research rather than identifying a causal investment return.

6.3 Financial information and climate representations have different roles

Liquidity contributes less than inventive history and technological knowledge. Financial-constraint research (Xu & Kim, 2022; Liu & Xu, 2024; Nguyen & Do, 2025) gives these indicators economic meaning, while other predictors may already represent firms' capacity to sustain research. Economic relevance and conditional predictive importance can therefore differ.

Cash can reflect flexibility, profitability, or preparedness (Duong et al., 2020; Javadi et al., 2023; Lee et al., 2023). Flexible prediction accommodates this ambiguity. Auxiliary correlations do not establish a sequential liquidity pathway behind the forecast gains.

Broad language adds information while narrow hazard terms add little. This distinction accords with differentiated climate concepts (Sautner et al., 2023) and varied climate–innovation findings (Lee et al., 2025; Tian et al., 2024; Ren et al., 2024). Monitoring should preserve the meaning of each textual signal.

6.4 Model selection and explanations for managerial review

At 10% review capacity, full/core application shares differ only from 76.19% to 76.29%, despite improved continuous prediction. Organizational AI research emphasizes decision structures (Shrestha et al., 2019; Raisch & Krakowski, 2021). A manager seeking differentiated forecasts across all firms may therefore prefer a different model from one selecting a fixed shortlist. Additional data collection should be assessed against the objective and costs of the review process.

Global attribution, local decomposition, and nonlinear responses guide examination of unexpected forecasts. This connects anticipation and adaptation (Duchek, 2020; Singh & Goyal, 2023; Gupta et al., 2023) to observable company records. The weighted additive explanation also permits a numerical reconstruction check. Its managerial role is to direct further investigation; review rankings evaluate subsequent application concentration rather than intervention benefits.

International climate–finance evidence (Huang et al., 2018) motivates evaluating the framework in other firm settings. Transfer requires locally appropriate dictionaries, accounting definitions, patent coverage, and chronological tests. The reusable contribution is the computational comparison among information sets, representations, and review objectives. Decision costs, invention quality, and realized environmental outcomes would extend the present assessment.

7. Conclusion and Decision Implications

Across 10,447 test observations, stacking achieves R^2 0.635 and RMSE 0.7297. Matched core, climate-only, and AI-only RMSE values are 0.7422, 0.7381, and 0.7356. Both blocks contribute conditional information. Full-information ridge's RMSE of 0.7527 supports the relevance of nonlinear representation.

Inventive history and AI patents dominate fitted explanations; broader climate language adds more than narrow hazard terms. Software responses rise and flatten. XGBoost leads squared-error performance and persistence leads absolute-error performance. Small review-ranking gains show why managerial forecasting requires separate evaluation of information increments, functional flexibility, and the decision objective.

Managers can begin with invention history and financial characteristics, adding climate and AI records where they improve the intended forecast. A matched core ensemble remains competitive for fixed-capacity review. Explanations guide examination of unusual forecasts, while funding or intervention decisions require evidence on quality, adoption, outcomes, and decision costs. The resulting workflow links computational evaluation to the practical use of corporate predictions.

8. Research Limitations

First, climate language represents attention and disclosure rather than observed shocks. Broader narratives provide most incremental value; overlap, sectoral vocabulary, and strategic communication

may shape that signal. Location-specific hazards and independently coded concepts would clarify transferability.

Second, AI assets and patents represent selected resources, not deployment quality. AI and green patent classifications may overlap. Ablation varies across cohorts, and explanations concern the final origin. Deployment measures and repeated-origin interpretation would clarify resource configurations and temporal stability.

Third, three test cohorts use archival extracts rather than reconstructed vintages. Pointwise intervals condition on fixed models. Review metrics value patent quantity without decision costs or environmental benefits. Archived vintages, broader periods, and prospective evaluation would strengthen operational assessment.

Appendix A. Auxiliary Panel Correlations

A.1 Sample and descriptive information

The auxiliary complete-case panel contains 10,484 observations from 1,648 nonfinancial firms. Feature years span 2008–2021, liquidity t+1, and inventions t+2. Financial ratios are winsorized at 1%/99%; climate and composite AI are standardized. These transformations remain separate from forecasting.

Table A1. Descriptive statistics for the auxiliary complete-case panel

Variable	Mean	SD	P25	Median	P75
Future green invention output	0.7293	1.1054	0.0000	0.0000	1.0986
Next-year liquidity pressure	0.1814	0.2335	0.0333	0.1953	0.3410
Physical climate exposure	-0.0000	1.0000	-0.3652	-0.3652	-0.3652
AI capability	0.0000	1.0000	-0.7161	-0.2718	0.4370
Firm size	22.1777	1.2329	21.3138	22.0127	22.8792
Leverage	0.4312	0.2015	0.2713	0.4258	0.5815
Profitability	0.0341	0.0691	0.0141	0.0366	0.0659
Operating cash flow	0.0454	0.0713	0.0067	0.0445	0.0859
Asset tangibility	0.1919	0.1484	0.0777	0.1592	0.2745
Prior green invention output	0.5442	0.9574	0.0000	0.0000	0.6931

Note. N = 10,484. These statistics refer to the auxiliary complete-case panel. The main forecasting test sample has 10,447 observations and a different missing-data eligibility rule. The auxiliary outcome has a zero share of 57.80%.

A.2 Specifications and statistical interpretation

Equation (A1) re-standardizes the mean of standardized software, hardware, and AI inventions. This auxiliary composite is distinct from the forecasting model's separate dimensions.

$$\mathrm{AI}_{i,t-1} = z\left[\frac{z\left(\mathrm{Soft}_{i,t-1}\right) + z\left(\mathrm{Hard}_{i,t-1}\right) + z\left(\mathrm{AIPat}_{i,t-1}\right)}{3}\right] \qquad (A1)$$

Z contains standardized hazards, lagged AI, and their interaction. Equations (A2)–(A3) model liquidity and inventions using firm/year effects and covariates. Parsimonious innovation omits later liquidity and its interaction. Later liquidity enters auxiliary equations only. These full-period associations are not forecasting benchmarks.

$$\mathrm{LP}_{i,t+1} = \alpha_i + \lambda_t + a^{'} Z_{it} + \gamma^{'} X_{it} + \varepsilon_{it} \qquad (A2)$$

$$\mathrm{GI}_{i,t+2} = \mu_i + \tau_t + c^{'} Z_{it} + b_1 \mathrm{LP}_{i,t+1} + b_2 \mathrm{LP}_{i,t+1} \mathrm{AI}_{i,t-1} + \delta^{'} X_{it} + u_{it} \qquad (A3)$$

Alternating projections remove firm/year effects. Firm-clustered score covariances include a finite-cluster adjustment. Table A2 reports coefficients and standard errors; Table 3 separately evaluates later forecasts.

Table A2. Auxiliary conditional panel associations

Variable	Parsimonious innovation	Liquidity equation	Extended innovation
Climate exposure	0.00410 (0.00814)	-0.00114 (0.00172)	0.00408 (0.00814)
AI capability	0.07319 (0.02264)	0.01100 (0.00462)	0.06478 (0.02446)
Climate × AI	0.01245 (0.00916)	0.00028 (0.00153)	0.01262 (0.00919)
Liquidity pressure	—	—	-0.01267 (0.05892)
Liquidity × AI	—	—	0.05318 (0.06260)
Observations	10,484	10,484	10,484
Firm and year effects	Yes	Yes	Yes
Within R^2	0.0580	0.2573	0.0582

Note. Firm-clustered standard errors are in parentheses. The innovation equations include the financial covariates and lagged innovation and liquidity characteristics. All equations include firm and year effects. Later liquidity in the extended equation is a descriptive conditioning variable and is excluded from prediction inputs.

Equation (A4) defines conditional coefficient products at AI level w. A joint Gaussian firm-cluster multiplier bootstrap uses 1,999 draws and common firm multipliers across equations. Table A3 reports percentile intervals for these association summaries.

$$D(w) = (a_1 + a_3 w)(b_1 + b_2 w) \qquad (A4)$$

Table A3. Auxiliary conditional coefficient products

AI level	Coefficient product	95% lower	95% upper
-1 SD	0.0000932	-0.0003623	0.0007291
0 SD	0.0000144	-0.0002648	0.0002926
1 SD	-0.0000347	-0.0006447	0.0004475

Note. Every 95% interval includes zero. The high-minus-low product difference is −0.0001279, with interval [−0.0009698, 0.0005414]. Intervals are based on 1,999 joint firm-cluster multiplier bootstrap draws.

Statements and Declarations

Clinical trial number: not applicable.

Ethics, Consent to Participate, and Consent to Publish declarations: not applicable.

Funding

This research received no funding.

Declaration of competing interest

The authors declare no competing financial interests or personal relationships that could have influenced this work.

Data availability

The financial and patent datasets are subject to their providers' access and redistribution conditions. Analysis code and variable-construction documentation are available from the corresponding author upon reasonable request. Access to source-level and derived records is subject to the relevant provider permissions.

References

Apley, D. W., & Zhu, J. (2020). Visualizing the effects of predictor variables in black box supervised learning models. *Journal of the Royal Statistical Society Series B: Statistical Methodology, 82*(4), 1059–1086. https://doi.org/10.1111/rssb.12377

Bolton, P., & Kacperczyk, M. (2021). Do investors care about carbon risk? *Journal of Financial Economics, 142*(2), 517–549. https://doi.org/10.1016/j.jfineco.2021.05.008

Bolton, P., & Kacperczyk, M. (2023). Global pricing of carbon-transition risk. *The Journal of Finance, 78*(6), 3677–3754. https://doi.org/10.1111/jofi.13272

Booker, A., Chiu, V., Groff, N., & Richardson, V. J. (2024). AIS research opportunities utilizing machine learning: From a meta-theory of accounting literature. *International Journal of Accounting Information Systems, 52*, 100661. https://doi.org/10.1016/j.accinf.2023.100661

Chen, L., Li, S., & She, Z. (2025). A study on the impact of artificial intelligence applications on corporate green technological innovation: A mechanism analysis from multiple perspectives. *International Review of Economics & Finance, 103*, 104490. https://doi.org/10.1016/j.iref.2025.104490

Dong, X., Zhou, N., Zhao, X., & Yang, S. (2025). The impact of artificial intelligence on corporate green innovation: Can "increasing quantity" and "improving quality" go hand in hand? *Journal of Environmental Management, 376*, 124439. https://doi.org/10.1016/j.jenvman.2025.124439

Duchek, S. (2020). Organizational resilience: A capability-based conceptualization. *Business Research, 13*(1), 215–246. https://doi.org/10.1007/s40685-019-0085-7

Duong, H. N., Nguyen, J. H., Nguyen, M., & Rhee, S. G. (2020). Navigating through economic policy uncertainty: The role of corporate cash holdings. *Journal of Corporate Finance, 62*, 101607. https://doi.org/10.1016/j.jcorpfin.2020.101607

Enholm, I. M., Papagiannidis, E., Mikalef, P., & Krogstie, J. (2022). Artificial intelligence and business value: A literature review. *Information Systems Frontiers, 24*(5), 1709–1734. https://doi.org/10.1007/s10796-021-10186-w

Giglio, S., Kelly, B., & Stroebel, J. (2021). Climate finance. *Annual Review of Financial Economics, 13*(1), 15–36. https://doi.org/10.1146/annurev-financial-102620-103311

Gupta, S., Modgil, S., Choi, T. M., Kumar, A., & Antony, J. (2023). Influences of artificial intelligence and blockchain technology on financial resilience of supply chains. *International Journal of Production Economics, 261*, 108868. https://doi.org/10.1016/j.ijpe.2023.108868

Haarhaus, T., & Liening, A. (2020). Building dynamic capabilities to cope with environmental uncertainty: The role of strategic foresight. *Technological Forecasting and Social Change, 155*, 120033. https://doi.org/10.1016/j.techfore.2020.120033

Huang, H. H., Kerstein, J., & Wang, C. (2018). The impact of climate risk on firm performance and financing choices: An international comparison. *Journal of International Business Studies, 49*(5), 633–656. https://doi.org/10.1057/s41267-017-0125-5

Javadi, S., Masum, A. A., Aram, M., & Rao, R. P. (2023). Climate change and corporate cash holdings: Global evidence. *Financial Management, 52*(2), 253–295. https://doi.org/10.1111/fima.12420

Kapoor, S., & Narayanan, A. (2023). Leakage and the reproducibility crisis in machine-learning-based science. *Patterns, 4*(9), 100804. https://doi.org/10.1016/j.patter.2023.100804

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

Lee, C. C., Li, M., & Zhang, J. (2025). How climate risk affects corporate green innovation: Fresh evidence from China's listed companies. *Emerging Markets Finance and Trade, 61*(8), 2302–2315. https://doi.org/10.1080/1540496x.2024.2449463

Lee, C. C., Wang, C. W., & Thinh, B. T. (2023). Green development, climate risks, and cash flow: International evidence. *Pacific-Basin Finance Journal, 79*, 102021. https://doi.org/10.1016/j.pacfin.2023.102021

Liu, Y., & Xu, R. (2024). Money talks: How financial constraint influence corporate environmental behavior prioritization. *International Review of Economics & Finance, 91*, 556–578. https://doi.org/10.1016/j.iref.2024.01.014

Lundberg, S. M., Erion, G., Chen, H., DeGrave, A., Prutkin, J. M., Nair, B., Katz, R., Himmelfarb, J., Bansal, N., & Lee, S. I. (2020). From local explanations to global understanding with explainable AI for trees. *Nature Machine Intelligence, 2*(1), 56–67. https://doi.org/10.1038/s42256-019-0138-9

Mikalef, P., & Gupta, M. (2021). Artificial intelligence capability: Conceptualization, measurement calibration, and empirical study on its impact on organizational creativity and firm performance. *Information & Management, 58*(3), 103434. https://doi.org/10.1016/j.im.2021.103434

Muhlroth, C., & Grottke, M. (2022). Artificial intelligence in innovation: How to spot emerging trends and technologies. *IEEE Transactions on Engineering Management, 69*(2), 493–510. https://doi.org/10.1109/tem.2020.2989214

Nambisan, S., Wright, M., & Feldman, M. (2019). The digital transformation of innovation and entrepreneurship: Progress, challenges and key themes. *Research Policy, 48*(8), 103773. https://doi.org/10.1016/j.respol.2019.03.018

Nguyen, Q. M. N., & Do, T. K. (2025). Do financial constraints hinder or spur green innovation? *Economics Letters, 257*, 112688. https://doi.org/10.1016/j.econlet.2025.112688

Raisch, S., & Krakowski, S. (2021). Artificial intelligence and management: The automation–augmentation paradox. *Academy of Management Review, 46*(1), 192–210. https://doi.org/10.5465/amr.2018.0072

Ren, X., Li, W., & Li, Y. (2024). Climate risk, digital transformation and corporate green innovation efficiency: Evidence from China. *Technological Forecasting and Social Change, 209*, 123777. https://doi.org/10.1016/j.techfore.2024.123777

Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence, 1*(5), 206–215. https://doi.org/10.1038/s42256-019-0048-x

Sautner, Z., van Lent, L., Vilkov, G., & Zhang, R. (2023). Firm-level climate change exposure. *The Journal of Finance, 78*(3), 1449–1498. https://doi.org/10.1111/jofi.13219

Shrestha, Y. R., Ben-Menahem, S. M., & von Krogh, G. (2019). Organizational decision-making structures in the age of artificial intelligence. *California Management Review, 61*(4), 66–83. https://doi.org/10.1177/0008125619862257

Singh, S., & Goyal, M. K. (2023). Enhancing climate resilience in businesses: The role of artificial intelligence. *Journal of Cleaner Production, 418*, 138228. https://doi.org/10.1016/j.jclepro.2023.138228

Teece, D. J. (2018). Dynamic capabilities as (workable) management systems theory. *Journal of Management & Organization, 24*(3), 359–368. https://doi.org/10.1017/jmo.2017.75

Tian, Y., Chen, S., & Dai, L. (2024). How climate risk drives corporate green innovation: Evidence from China. *Finance Research Letters, 59*, 104762. https://doi.org/10.1016/j.frl.2023.104762

Vial, G. (2019). Understanding digital transformation: A review and a research agenda. *The Journal of Strategic Information Systems, 28*(2), 118–144. https://doi.org/10.1016/j.jsis.2019.01.003

Xu, Q., & Kim, T. (2022). Financial constraints and corporate environmental policies. *The Review of Financial Studies, 35*(2), 576–635. https://doi.org/10.1093/rfs/hhab056

Zhong, K., & Song, L. (2025). Artificial intelligence adoption and corporate green innovation capability. *Finance Research Letters, 72*, 106480. https://doi.org/10.1016/j.frl.2024.106480

Zhong, Q., Zhang, Q., & Yang, J. (2025). Can artificial intelligence empower energy enterprises to cope with climate policy uncertainty? *Energy Economics, 141*, 108088. https://doi.org/10.1016/j.eneco.2024.108088