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.

From Digital Claims to Verifiable Outcomes: A Deployment-Oriented AI and IoT Disclosure Index Linked to Emissions Intensity and Safety Operations

Elham Jahanbakhsh

elhamjahanbakhsh2014@gmail.com

University of the West of Scotland https://orcid.org/0009-0006-0442-3635

Research Article

Keywords: Sustainability reporting, text-as-data, AI/IoT deployment, emissions intensity, workplace safety, downside risk, accountability

Posted Date: February 25th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-8960368/v1

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

Additional Declarations: The authors declare no competing interests.

From Digital Claims to Verifiable Outcomes: A Deployment-Oriented AI and IoT Disclosure Index Linked to Emissions Intensity and Safety Operations

Abstract

Amid climate regulation and the low-carbon transition, firms increasingly disclose operational digitalization for cleaner production, yet existing studies relying on broad digital proxies or environmental, social, and governance (ESG) scores limit inference about what these claims signal for environmental management and workplace safety. A deployment-oriented text-as-data index of disclosed Artificial Intelligence (AI) and Internet of Things (IoT) deployment (AIDI) is developed and validated from U.S. annual filings and sustainability reports (2018–2024; 8,605 firms; 36,514 firm-years) using human coding, innovation benchmarks, and a placebo "digital buzz" measure. The index is then linked to operational outcomes—the U.S. EPA Greenhouse Gas Reporting Program (GHGRP) emissions intensity and workplace injury rates (total recordable incident rate; TRIR)—and tested in fixed-effects models with multiplicity adjustment, yielding limited within-firm links; facility- and process-level deployment constrains firm-year verification**.** The study further offers an open tool for assessing AI/IoT disclosures, identifies limits to disclosure-based accountability, and shows a robust safety–downside risk link, informing sustainability assessment and stakeholder assurance.

Keywords: Sustainability reporting, text-as-data, AI/IoT deployment, emissions intensity, workplace safety, downside risk, accountability.

1. Introduction

Firms are investing heavily in artificial intelligence (AI), the Internet of Things (IoT), and related digital infrastructures to support environmental management, regulatory compliance, and decarbonization targets (Gartner, 2024; KPMG, 2025). These technologies are frequently

presented as tools to monitor emissions in real time, detect leaks and pollution, improve process control and operational efficiency, strengthen environmental reporting, and respond to intensifying regulatory and stakeholder scrutiny (Khan et al., 2025; Popescu et al., 2024). At the same time, a growing volume of corporate disclosures describing AI/IoT initiatives raises a central accountability question for corporate social responsibility (CSR) and environmental management research: what do AI/IoT deployment disclosures credibly signal about environmental and occupational health and safety (OHS) performance, and what are the limits of inference from disclosure-based measures (Elsayed, 2025; Sneideriene & Legenzova, 2025)?

Existing evidence is difficult to interpret because broad digital proxies and composite environmental, social, and governance (ESG) measures blur deployment, disclosure, and performance, leaving it unclear whether disclosed AI/IoT deployments track measurable emissions and safety improvements or support monitoring and assurance (Benuzzi et al., 2025; Kathan et al., 2025; Zhu & Li, 2025). Digital transformation and Industry 4.0 research suggest that analytics and IoT can improve monitoring and efficiency, but empirical tests often use coarse indicators and do not link specific deployments to externally verifiable environmental outcomes (Ologeanu-Taddei et al., 2025; Zhu & Li, 2025). In parallel, ESG and finance research frequently treats disclosure-based measures as performance signals, even though disclosure and operational outcomes may diverge (Kathan et al., 2025; Lublóy et al., 2025).

This study addresses these gaps by examining three questions: (1) Are disclosed AI/IoT deployments associated with subsequent changes in environmental and safety operations, measured using GHG emissions intensity and workplace injury rates (total recordable incident rate; TRIR)? (2) Do these associations vary across industry contexts, particularly where carbon exposure and regulatory attention are higher? (3) Are operational outcomes linked to firm risk

outcomes, and do they help explain any association between AI/IoT deployment disclosures and downside risk? The focus is disclosure validity for operational AI/IoT deployment and its accountability implications.

To address these questions, a deployment-oriented text-as-data index of disclosed AI/IoT deployment (AIDI) is constructed from sentence-level disclosures in U.S. annual filings and sustainability reports (2018–2024; 8,605 firms; 36,514 firm-years). The index targets concrete operational deployments rather than generic digital rhetoric and is validated using sentence-level human coding, external innovation benchmarks, and a placebo "digital buzz" measure; the dictionary and code are released openly. AIDI is then linked to externally sourced operational outcomes (GHG emissions intensity from the U.S. EPA GHGRP and workplace injury rates) and to downside risk, using firm and year fixed-effects models with lagged regressors.

The study contributes to sustainability assessment relevant to cleaner production, CSR, and environmental management research in three ways. First, it introduces and openly releases a validated, deployment-oriented measure of AI/IoT disclosure content that distinguishes operational deployment from generic digital rhetoric. Second, it shows how data linkage, limited overlap in linked samples, and concentrated within-firm variation constrain accountability claims in disclosure-linked operational panels, clarifying what public reporting can and cannot reveal about responsible digitalization. Third, it documents a robust within-firm relationship between workplace safety (TRIR) and downside risk, reinforcing the material importance of OHS performance for risk management and stakeholder evaluation. For practice and policy, the results caution against treating AI/IoT deployment disclosure as a direct indicator of operational improvement and highlight the value of stronger operational measurement, verification, and research designs that better support assurance and accountability. The remainder of the article

proceeds as follows. Section 2 develops the conceptual framework and hypotheses; Section 3 describes the data, measurement, and empirical design; Section 4 presents the results and robustness analyses; Section 5 discusses implications for environmental and safety management, disclosure quality, and assurance and verification; and Section 6 concludes.

2. Theory and Hypotheses

This framework adopts an association-based accountability perspective on responsible digitalization in environmental and safety management. Digital technologies may strengthen monitoring and control of emissions and workplace risks, but disclosure-based indicators can also reflect symbolic reporting rather than operational change (Huber et al., 2025; Popescu et al., 2024; Roszkowska-Menkes et al., 2024). The framework therefore links disclosed AI/IoT deployment capability to environmental and OHS performance, and then to downside risk, while emphasizing boundary conditions that shape verifiability, stakeholder attention, and monitoring and assurance (Figure 1).

2.1 Disclosure-Based AI/IoT Deployment Capability as an Operational Management Capability

AI/IoT deployment capability is conceptualized as a firm-level bundle of data assets, analytics routines, technical infrastructure, and governance arrangements through which AI/IoT are integrated into operations (Chen et al., 2025; Phan et al., 2026). The focal construct is the joint deployment bundle—IoT-enabled sensing coupled with AI-enabled analytics and workflow integration—because many environmental and safety applications combine these components (e.g., sensors feeding predictive models that trigger operational adjustments; Khadam et al., 2025; Ramadan et al., 2024). In environmental and safety domains, formal intellectual property protection is often limited; performance differences may instead depend on data–process

complementarities, operational integration, and governance that supports reliable use, documentation, and accountability (Bernardo et al., 2024; Papagiannidis et al., 2025; World Intellectual Property Organization [WIPO], 2024).

Measurement boundary. The focal construct is treated as an observable disclosure footprint capturing implemented AI/IoT deployment and orchestration (including integration of vendor solutions), while underlying skills and governance remain latent. Accordingly, estimates are interpreted as disciplined within-firm associations consistent with the proposed logic, not definitive causal effects.

2.2 From Disclosed Digital Deployment to Environmental and Safety Operations

Operationally deployed AI/IoT may influence environmental and safety outcomes through three routine classes: sensing (continuous monitoring of emissions, leaks, equipment condition, near-miss events), prediction (forecasting failures and incidents), and control and reconfiguration (process adjustments, preventive maintenance, and adaptive safety responses; Huber et al., 2025; Khadam et al., 2025; Popescu et al., 2024). When embedded in practice and supported by governance, these routines can improve environmental performance (lower emissions intensity) and workplace safety (lower injury rates; El-Helaly, 2024; Papagiannidis et al., 2025; Zhang et al., 2025a). However, disclosure–outcome links may be difficult to detect at the consolidated firm-year level. Deployments often occur at the facility or process level; disclosures and outcomes are aggregated; and both disclosure measures and external outcomes can be noisy. Some disclosures also involve partners or supply-chain actors, whereas outcomes are measured at the consolidated firm level (Khadam et al., 2025; Roszkowska-Menkes et al., 2024). Thus, AIDI is best viewed as a firm-centric indicator of operational deployment and orchestration that may not align cleanly with firm-year outcome perimeters.

Hypothesis 1 (H1). Higher disclosed, deployment-oriented AI/IoT capability (AIDI) is expected to be associated with improved environmental and safety operational outcomes, reflected in lower emissions intensity and lower injury rates; however, firm-year aggregation and measurement noise may attenuate within-firm associations.

2.3 Environmental and Safety Operations and Downside Risk

Environmental and safety performance can influence firm outcomes through (i) operational and compliance exposure (reduced inefficiencies, downtime, penalties, and liability) and (ii) incident risk (lower probability and severity of accidents, shutdowns, enforcement actions, and litigation), which should reduce downside risk (Tran et al., 2024; Yang et al., 2024; Zhang et al., 2025b). The operational outcomes used here—emissions intensity and recorded workplace injuries—are intentionally narrow and do not capture broader social impacts of digitalization (e.g., surveillance, bias). These measures may also proxy for broader organizational attributes such as operational discipline, reporting systems, and risk management. Consequently, they are treated as informative operational-channel indicators rather than formal causal mediators.

Hypothesis 2a (H2a: Operations → profitability). Emissions intensity and workplace safety (TRIR) are associated with subsequent profitability (ROA).

Hypothesis 2b (H2b: Operations → downside risk). Workplace safety is associated with subsequent downside risk: lower TRIR is expected to predict lower downside risk. Lower emissions intensity may also be associated with lower downside risk.

Hypothesis 2c (H2c: Joint-model patterns). In joint models including lagged AIDI and lagged operational outcomes, operational outcomes are expected to retain explanatory content for firm outcomes, yielding patterns consistent with an operations channel.

Mechanism-consistent interpretation. Joint models that include operational outcomes alongside AIDI are used to assess whether estimated AIDI–outcome associations change in ways consistent with an operations channel, recognizing that measurement noise and limited within-firm variation may attenuate such patterns.

2.4 Boundary Conditions: Carbon Salience as a Context for Monitoring and Verification

Whether disclosed AI/IoT deployment aligns with operational improvement—and whether improvements are recognized—depends on contexts that affect materiality, visibility, and verification (European Securities and Markets Authority [ESMA], 2025; Pizzi et al., 2024). A central contextual dimension is industry carbon salience, which proxies for the extent to which greenhouse gas (GHG) emissions are economically material and subject to heightened regulatory and stakeholder scrutiny (Financial Stability Board, 2023; IFRS Foundation, 2024). In high–carbon-salience industries, AI/IoT deployments that generate verifiable improvements in emissions control and safety management may be more consequential and more likely to be monitored and rewarded through compliance outcomes, contracting, and stakeholder evaluation (Pizzi et al., 2024; Seera et al., 2025). Carbon salience is a coarse proxy, however, and may also reflect correlated features such as disclosure norms, enforcement intensity, and political attention; moderation tests are therefore interpreted as evidence of heterogeneity in associations rather than identification of a single mechanism (Shui et al., 2024; von Hippel & Schuetze, 2025).

Hypothesis 3 (H3). Industry carbon salience is expected to condition the association between AIDI and subsequent environmental and safety operational outcomes; however, aggregation and verification constraints may attenuate detectable differences in within-firm associations across salience contexts.

2.5 Summary

The framework links disclosed AI/IoT deployment to environmental and OHS performance and then to downside risk, emphasizing boundary conditions related to carbon salience and the feasibility of monitoring and verification. The contribution is a disclosure-validity and operational-accountability structure—AIDI -> environmental and OHS operations -> firm outcomes—offering a deliberately bounded, association-based interpretation appropriate to disclosure-linked firm-year panel evidence. Section 3 describes the data sources, construction of AIDI and outcomes, and the empirical strategy used to test these hypotheses.

3. Data, Measures, and Empirical Strategy

3.1 Sample and Scope

U.S. listed firms with usable annual filings and/or sustainability reports from 2018–2024 yield 8,605 firms and 36,514 firm-year observations after reporting-quality and text-processing filters. The corpus skews toward larger, higher-disclosure issuers and industries with more visible environmental and safety exposures, reflecting data availability and linkage feasibility rather than sector targeting.

Firm identifiers and financial fundamentals are drawn from standard U.S. market and accounting databases and matched to SEC registrants using CIK and security-identifier crosswalks (Center for Research in Security Prices, 2025; Wharton Research Data Services [WRDS], 2026). Annual filings (primarily Form 10-K) are retrieved from the SEC's Electronic Data Gathering, Analysis, and Retrieval (EDGAR) system and, where available, supplemented with sustainability reports from firm websites and ESG-report archives (Corporate Register, 2026; U.S. Securities and Exchange Commission, 2025). Operational outcomes include TRIR from a commercial occupational-safety and ESG compilation (definitions aligned to OSHA guidance; Occupational Safety and Health Administration [OSHA], 2025) and facility-level greenhouse gas emissions

from the U.S. EPA Greenhouse Gas Reporting Program (GHGRP), linked to corporate parents and SEC registrants using established facility–issuer and parent–subsidiary mappings (see Supplementary Appendix S2; U.S. Environmental Protection Agency, 2025).

The text corpus is constructed independently of outcome availability, and outcomes are linked afterward to form outcome-specific panels (Figure 2). Regressions are estimated on outcome-specific locked samples (so nested specifications use identical rows within each outcome), with selection into linkage summarized in Supplementary Appendix Table SA1; outcomes and controls are winsorized at the 1st and 99th percentiles by year, observations failing denominator, unit, or text-quality screens are removed (Supplementary Appendix S1), and financial institutions are excluded or, where required, handled using sector fixed effects. The replication package (code and constructed indices, including AIDI) is available in a blinded OSF repository. Inferences are most directly applicable to relatively data-intensive, higher-disclosure U.S. issuers during 2018–2024; generalization beyond this setting should be cautious given selection into disclosure and linked operational datasets.

3.2 Text-as-Data Measure of AI and IoT Deployment Capability (AIDI)

AI/IoT deployment capability is proxied by AIDI, a deployment-oriented text-as-data measure constructed from annual filings and sustainability reports that captures disclosed, implemented operational uses (including environmental and safety applications) rather than generic digital transformation rhetoric.

Sentence-level deployment definition. Deployment sentences contain an AI/IoT term, reference a concrete operational application, and use an implementation verb indicating ongoing use. Intentions, pilots, and generic rhetoric are excluded; boilerplate and duplicated text are removed;

discontinued initiatives are coded as non-deployments unless ongoing use is explicit, so residual misclassification should attenuate estimated associations.

Dictionary scope and index construction. The dictionary covers AI/IoT terms (e.g., machine learning, sensors, and telemetry), and qualifying sentences are tagged as AI, IoT, or both and aggregated into AIDI to capture an integrated deployment bundle; disclosures commonly reference emissions monitoring, predictive maintenance, anomaly detection, safety analytics, traceability, and compliance support.

Timing and variants. The baseline measure is a deployment-share index, AIDI_share_{i,t} (deployment sentences divided by total sentences), entered with a one-year lag (AIDI_share_{i,t−1}). When report-length or disclosure-intensity controls create collinearity, robustness checks use an unscaled count variant, AIDI_index_{i,t} (deploy_relaxed), entered as AIDI_index_{i,t−1}; moderation models use the standardized lagged share, AIDI_share,z_{i,t−1}. Because AIDI captures deployments described as in use (not start dates), it is interpreted as a stock-like disclosure footprint, so timing error should attenuate estimates.

Variable alignment. All variables are aligned to fiscal year t: AIDI is constructed from year-t texts and entered with a one-year lag; operational outcomes are measured in t, and downside risk is based on July t–June t+1 returns assigned to t; GHGRP coverage in the cleaned linkage panel runs through 2023, so lagged-AIDI emissions models are limited to outcome years 2019–2023.

Validation and discriminant checks. Validation uses a stratified, double-coded sentence sample; intercoder agreement and precision, recall, and F1 are reported in Section 4 and the Supplementary Appendix. At the firm-year level, AIDI correlates modestly with external technology benchmarks (e.g., AI/IoT patent families) and investment proxies (R&D, IT, and capex intensity), but only weakly with generic disclosure volume and a placebo "digital buzz" measure; a broad digital-terms

placebo index does not reproduce the main patterns, supporting AIDI as a deployment-oriented disclosure measure rather than narrative intensity (The Lens, 2023; WRDS, 2026).

3.3 Measures

3.3.1 Environmental and Safety Operations

Operational outcomes are measured at the firm-year level in two domains: emissions intensity and workplace safety (injury rates). Facility-level emissions are aggregated to the consolidated firm-year level to align with AIDI. Lower values indicate better performance across outcomes.

Emissions intensity (GHGRP). Greenhouse gas emissions (metric tons CO_2e) are obtained from the U.S. EPA GHGRP, aggregated to corporate parents, and linked to SEC registrants using established issuer–facility mappings (see Supplementary Appendix S2; U.S. Environmental Protection Agency, 2025). Emissions intensity is GHGRP emissions scaled by firm-year net sales; the analysis uses ln(1 + emissions intensity).

Injury rate (TRIR). Workplace safety is measured using TRIR from a commercial occupational-safety and ESG dataset that harmonizes OSHA-style disclosures. TRIR is recordable incidents per 200,000 hours worked; where incident counts and hours are available, TRIR is reconstructed as 200,000 × (incidents/hours; OSHA, 2025). When components are unavailable, the standardized TRIR value is used subject to the denominator and unit verification rules (Supplementary Appendix S1; Table SA2).

These measures capture specific operational outcomes rather than a composite ESG construct and do not cover other social dimensions. To limit common-source bias, outcomes derived from the same texts used to construct AIDI are excluded.

3.3.2 Financial Performance and Downside Risk

Firm outcomes include an accounting-based profitability measure and market-based downside-risk metrics capturing lower-tail exposure.

Profitability. The primary measure is return on assets (ROA), defined as income before extraordinary items divided by average total assets in year t. Robustness checks use operating margin and return on sales (ROS). Supplementary analyses examine Tobin's Q, defined as (market value of equity + book value of debt) divided by book assets, as a forward-looking valuation outcome.

Downside risk. Market-based risk is computed from daily returns over a 12-month window aligned to outcome year t (July t to June t+1). Measures include (i) idiosyncratic volatility (the standard deviation of residual returns from a market or factor model), (ii) downside semideviation (the standard deviation of returns below the firm's mean), and (iii) a tail-loss proxy (e.g., the 5th or 10th percentile of returns; Ang et al., 2006; Engle & Manganelli, 2004; Sortino & van der Meer, 1991). These metrics are estimated in linear fixed-effects models and interpreted as proxies for realized lower-tail exposure. Because market-based measures may also reflect expectations and disclosure content, interpretation emphasizes operational outcomes and accounting profitability when assessing the capability → operations → outcomes mechanism.

3.3.3 Industry Carbon Salience

Industry carbon salience is constructed using sector taxonomies and emissions-intensity benchmarks. It captures cross-industry differences in the materiality and visibility of emissions and efficiency and moderates the AIDI–operations relationship, reflecting contexts with greater regulatory and stakeholder scrutiny and stronger incentives to manage emissions and efficiency (Financial Stability Board, 2023; NewClimate Institute, 2024). The baseline specification uses a binary indicator for high–carbon-salience industries (e.g., energy, utilities, materials, chemicals,

metals, cement, transportation, and selected industrial manufacturing), with robustness checks using continuous sector emissions-intensity measures and alternative classifications (IFRS Foundation, 2023; Intergovernmental Panel on Climate Change [IPCC], 2023; International Energy Agency [IEA], 2024).

3.3.4 Controls and Primary Outcomes

Baseline models include one-year lagged firm controls—size (log assets), age, leverage (debt-to-assets), capital-expenditure intensity, and R&D intensity—along with reporting controls (standalone sustainability reporting, third-party assurance, and report length). Where feasible, text-quality measures (e.g., boilerplate share and readability) are included to limit reporting noise (Huang et al., 2024). Core specifications include firm and year fixed effects; heterogeneity is assessed using prespecified interactions (including industry carbon salience) and robustness specifications. Accordingly, primary outcomes and the moderator are prespecified as: (i) operational outcomes—GHGRP-based emissions intensity and TRIR; (ii) profitability—ROA; (iii) downside risk—downside semideviation; and (iv) moderator—industry carbon salience. P-values for the two primary operational outcomes are adjusted using the Holm–Bonferroni procedure, and additional outcomes and robustness analyses are treated as exploratory and reported in the Supplementary Appendix.

3.4 Baseline Empirical Specification

Hypothesis 1 is tested using firm-year panel models with firm and year fixed effects:

$$Ops_{i,t} = \alpha_i + \lambda_t + \beta_1 AIDI_share_{i,t-1} + \gamma' X_{i,t-1} + \varepsilon_{i,t}.$$

$Ops_{i,t}$ is ln(1 + emissions intensity) or TRIR; $AIDI_share_{i,t-1}$ and controls ($X_{i,t-1}$) are lagged one year. Standard errors are clustered by firm. The coefficient β_1 captures within-firm associations and is interpreted as a panel association rather than a causal effect.

3.5 Mechanism Tests: Operations as a Channel (H2)

To evaluate whether environmental and safety operations are consistent with a channel linking AIDI to firm outcomes (H2), a sequence of firm-year fixed-effects panel models with one-year lags is estimated. Results are interpreted as mechanism-consistent associations.

Step 1: AIDI → operations

Estimate the baseline operational model (Section 3.4) relating AIDI_share_{i,t−1} to Ops_{i,t}.

Step 2: Operations → firm outcomes

$$Perf_{i,t} = \alpha_i + \lambda_t + \theta_1 Ops_{i,t-1} + \gamma' X_{i,t-1} + u_{i,t},$$

where Perf_{i,t} is a firm outcome (ROA or a downside-risk measure) and Ops_{i,t−1} is lagged emissions intensity or TRIR.

Step 3: Joint model (AIDI + operations)

$$Perf_{i,t} = \alpha_i + \lambda_t + \varphi_1 AIDI^{(v)}_{i,t-1} + \varphi_2 Ops_{i,t-1} + \gamma' X_{i,t-1} + \nu_{i,t},$$

where AIDI^(v)_{i,t−1} is the unscaled index variant AIDI_index_{i,t−1} (deploy_relaxed). Evidence consistent with an operational channel is indicated when AIDI_share_{i,t−1} is associated with Ops_{i,t} and Ops_{i,t−1} is associated with Perf_{i,t} in the expected direction; in the joint model, attention focuses on whether the AIDI^(v) coefficient changes (e.g., attenuates) when Ops_{i,t−1} is included.

3.6 Moderation: Industry Carbon Salience (H3)

Hypothesis 3 is tested by adding an interaction between lagged AIDI and industry carbon salience to the operational-outcome model:

$$Ops_{i,t} = \alpha_i + \lambda_t + \beta_1 AIDI_share_{i,t-1} + \beta_2 (AIDI_share_{i,t-1} \times CarbonSalience_j) + \gamma' X_{i,t-1} + \varepsilon_{i,t},$$

where CarbonSalience_j is an industry-level carbon-salience measure. The baseline specification uses a time-invariant high–carbon-salience indicator (sector taxonomy); robustness checks use continuous emissions-intensity benchmarks and alternative classifications (NewClimate Institute, 2024). For outcomes where lower values indicate better performance (emissions intensity and TRIR), β_2 captures whether the within-firm association between lagged AIDI and operational outcomes differs systematically in high–carbon-salience industries. Moderation is assessed via β_2 and marginal effects (simple slopes) for low versus high carbon salience. The next section reports the main results from these models.

4. Results

4.1 Descriptive statistics and correlations

The text corpus includes 8,605 U.S. firms and 36,514 firm-year observations from 2018–2024 (lagged specifications begin in 2019). To situate the analyses, Supplementary Appendix Table SA3 reports descriptive statistics and variable-level availability for AIDI_share, operational outcomes (emissions intensity and TRIR), firm outcomes (ROA and downside risk), and contextual variables (EDGAR filing intensity and carbon salience). Supplementary Appendix Table SA4 reports Pearson correlations computed in the GHGRP-anchored linked correlation dataset using pairwise-complete observations (maximum N = 489). Effective sample sizes vary by pair because variables are drawn from additional linked sources. In particular, correlations involving TRIR rely on limited overlap (N = 116) and are reported for context only and should not be interpreted as stable estimates; regressions therefore use outcome-specific locked panels.

4.2 Validation of the AI/IoT deployment proxy (AIDI)

AIDI is validated as capturing disclosed, implemented AI/IoT deployments in core operations rather than generic “digital transformation” rhetoric. Validation draws on (i) sentence-level

human-coded ground truth, (ii) convergent evidence from external innovation benchmarks, and (iii) a discriminant placebo test that separates deployment language from narrative intensity (Supplementary Appendix Table SA5).

Human coding and classification performance. Reliability and classification performance are assessed using a stratified sentence-level validation set (n = 400) combining high-recall screened sentences with randomly sampled background sentences. Two independent coders label AI/IoT deployment content, with high agreement (Cohen's $\kappa = 0.92$). Relative to adjudicated labels, the rule-based classifier achieves precision = 0.80, recall = 0.67, and F1 = 0.73. Errors concentrate in borderline analytics statements near the deployment threshold, consistent with the conservative inclusion rule requiring an implementation action and a concrete asset or process reference. Because deployments are rare and the evaluation set oversamples borderline candidates, precision, recall, and F1 are emphasized over overall accuracy (Opitz, 2024).

Convergent validity (innovation benchmarks). At the firm-year level, AIDI is positively associated with external AI/IoT innovation measures derived from Lens-based patent data linked to SEC registrants (e.g., Pearson $r = 0.11$ for $\log(1 + \text{AI/IoT patent families})$; Spearman $\rho = 0.09$; n = 296; The Lens, 2023; WRDS, 2026). Although patenting is sparse (families > 0 in 8.8% of firm-years), the positive association is consistent with AIDI capturing AI/IoT-related innovation activity (Dussaux et al., 2023).

Discriminant validity (placebo "digital buzz"). A placebo "digital buzz" index based on broad digital and innovation terms that do not meet the deployment inclusion rule behaves more like narrative intensity, showing a weak association with report length ($r = 0.05$) and a stronger association with readability ($r = -0.22$).

Taken together, these checks support interpreting AIDI as a deployment-oriented disclosure measure rather than a proxy for general disclosure volume or narrative emphasis.

4.3 Associations between AIDI and environmental and safety operations (H1)

Supplementary Appendix Table SA6 reports two-way fixed-effects (firm and year) estimates relating lagged AI/IoT deployment disclosure, AIDI_share_{i,t−1}, to GHGRP-based emissions intensity (ln(1 + emissions intensity)) and workplace injury rates (TRIR). Each outcome is estimated using (i) a minimal fixed-effects specification and (ii) a full specification that adds the feasible set of lagged controls in the outcome-specific locked sample, with firm-clustered standard errors.

Emissions intensity. In the GHGRP-linked locked panel (344 firm-years; 106 firms; 2019–2023), the AIDI_share_{i,t−1} coefficient is positive in both minimal and full specifications and statistically indistinguishable from zero (Supplementary Appendix Table SA6, Panel A). Because intensity requires non-missing sales, the locked sample is a narrow intersection (GHGRP × sales × feasible controls). Within-firm variation in AIDI_share_{i,t−1} is concentrated in seven firms (Supplementary Appendix Table SA6, Panel B), and leave-one-firm-out diagnostics estimated on the maximal feasible emissions-intensity sample (Supplementary Appendix Table SA6.1, Panels A–B) show substantial coefficient dispersion, including sign changes, while remaining statistically indistinguishable from zero. A supplementary levels specification using ln(1 + GHGRP emissions), which avoids the sales requirement and expands coverage, yields qualitatively similar results (Supplementary Appendix Table SA6.1, Panel C).

Workplace injuries (TRIR). The TRIR locked panel is larger (3,483 firm-years; 1,074 firms; 2019–2024). The AIDI_share_{i,t−1} estimate is negative and similar across specifications (Supplementary Appendix Table SA6, Panel A) but is not statistically distinguishable from zero

after Holm–Bonferroni adjustment (adjusted p = 1.00; nominal p-values > 0.4). The evidence is therefore consistent with, but does not establish, a modest within-firm association between higher deployment-oriented AI/IoT disclosure and subsequent safety performance.

Overall, the H1 results do not provide robust statistical support for improvements in emissions intensity or TRIR in the locked-sample TWFE design. Estimates are imprecise, emissions results are sensitive to narrow sample intersections and mover composition, and TRIR associations remain non-rejectable after multiplicity adjustment.

4.4 Mechanism-consistent patterns: operations, profitability, and downside risk (H2a–H2c)

This subsection evaluates whether environmental and safety operations are associated with firm outcomes and whether joint models are consistent with a deployment-oriented disclosure → operations → outcomes pathway. Supplementary Appendix Table SA7 reports two-way fixed-effects (firm and year) regressions linking lagged operational outcomes to subsequent firm outcomes (H2a–H2b), and Supplementary Appendix Table SA8 reports joint TWFE models that include both lagged AIDI and lagged operational measures (H2c). Inference relies on firm-clustered standard errors with Holm–Bonferroni adjustments within each table and panel.

H2a (operations → profitability). Results provide limited support for an operations-to-profitability channel in the locked-sample TWFE design. For both emissions intensity and TRIR, ROA coefficients are close to zero and remain non-rejectable after Holm–Bonferroni adjustment.

H2b (operations → downside risk). Downside-risk results are sharper. Lagged workplace safety shows a strong and stable within-firm association with downside risk: higher TRIR at t−1 is associated with higher downside semi-deviation at t in both minimal and full specifications (Holm-adjusted $p < 0.001$), with little change across specifications. By contrast, emissions intensity exhibits weaker and specification-sensitive associations with downside risk; a nominal full-

specification effect ($p = 0.035$) does not survive Holm–Bonferroni adjustment (Holm-adjusted $p = 0.106$).

H2c (joint-model patterns). In joint TWFE models, GHGRP-linked emissions results remain close to zero and non-rejectable for both ROA and downside risk, with limited within-firm identifying variation in lagged AIDI. In TRIR joint models, lagged TRIR retains a strong, specification-invariant association with downside risk (Holm-adjusted $p < 0.001$), while lagged AIDI remains non-rejectable; ROA coefficients remain close to zero.

4.5 Moderation by industry carbon salience (H3)

Supplementary Appendix Table SA9 tests moderation by industry carbon salience using two-way fixed-effects models that interact standardized lagged AIDI (AIDI_share_{i,t−1}^z) with a high–carbon-salience indicator (SAL), estimated on outcome- and specification-specific feasible samples for GHGRP emissions intensity (2019–2023) and workplace injury rates (TRIR; 2019–2024). Across outcomes and specifications, the interaction terms are small relative to their standard errors and statistically non-rejectable (all $p > 0.30$). Point estimates are positive but imprecise, so any implied differences in marginal slopes should be treated as descriptive. Overall, the TWFE results provide no reliable evidence of moderation by industry carbon salience, and H3 is not supported.

4.6 Economic significance and scale interpretation

Supplementary Appendix Table SA10 uses the deploy_relaxed count/index variant, AIDI_index_{i,t−1}, to avoid absorption and collinearity in fully saturated models and to calibrate scale within the linked samples, rather than to provide a direct magnitude comparison to the share-based results in Supplementary Appendix Table SA6. Because disclosure is sparse in the linked samples (3.94–7.29% of firm-years have AIDI_index_{i,t−1} > 0; Panel B), magnitudes are

interpreted using adopter-relevant shifts within each outcome's estimation sample (0→p95 and the interquartile range among positive observations, IQR_pos). Coefficients from the benchmarking specifications (Panel A) are translated into implied changes under these shifts (Panel C).

Across outcomes, implied magnitudes are modest. For emissions intensity, a 0→p95 shift implies approximately +0.017% (and +0.012% for IQR_pos). For TRIR, implied changes are on the order of -5.5×10^{-5} to -7.9×10^{-5} TRIR units. For ROA, implied changes are approximately 0.002–0.003 basis points, and for downside semi-deviation approximately 8×10^{-6} to 1.2×10^{-5} units. These benchmarks likely reflect limited within-firm identifying variation and measurement constraints in the linked panels and may understate associations in settings with richer operational measurement or stronger exogenous variation.

4.7 Secondary and exploratory analyses

Supplementary Appendix analyses probe the robustness and interpretability of the baseline TWFE results along three dimensions: alternative AIDI constructions, additional controls, and dynamic and event-time diagnostics. Inference is qualitatively consistent with the main specifications, while reinforcing that AIDI–operations associations are estimated with limited precision in the within-firm locked-sample design.

Alternative AIDI constructions. Supplementary Appendix Tables SA15–SA16 report estimates using alternative text measures, including technology-mention indices and count/index variants without length normalization. Across constructions, associations with emissions intensity and TRIR are broadly consistent with the baseline deployment-oriented measure in sign and statistical (non)rejection, suggesting the main patterns are not driven by a particular operationalization or normalization choice.

Generic IT proxy. Supplementary Appendix Table SA12 adds an R&D-based proxy in a fixed-row design that holds the estimation sample constant. AIDI coefficients are essentially unchanged across outcomes, mitigating (though not eliminating) the concern that AIDI primarily captures broad IT intensity rather than deployment-oriented disclosures.

Dynamic diagnostics. Event-time and distributed-lag checks are treated as exploratory because they require substantially tighter data and shorter usable windows. A joint lead and lag specification is not identified for emissions in the restricted sample; emphasis is therefore placed on a feasible one-period lead test estimated separately (Supplementary Appendix Table SA13, Panel B). Lead coefficients are not statistically distinguishable from zero across outcomes, providing no evidence of systematic pre-trends in the feasible test window. Distributed-lag specifications (Supplementary Appendix Table SA14) do not reveal a stable lag structure, consistent with short event-time coverage in the linked panels; results are suggestive rather than definitive.

Additional checks. For downside risk, alternative AIDI constructions and standard sample exclusions yield qualitatively similar estimates in both minimal and full specifications (Supplementary Appendix Table SA11). Heterogeneity checks using alternative carbon-salience definitions and within-industry subsamples likewise do not produce stable slope differences, consistent with the null interaction results in the main moderation models.

Taken together, the Supplementary Appendix supports interpretation of AIDI as a deployment-oriented text measure and indicates that core conclusions are not an artifact of a single measurement choice or robustness decision. At the same time, operational-outcome associations remain imprecisely estimated in the locked-sample TWFE framework, warranting cautious

interpretation of effect sizes and causal direction. Section 5 discusses the theoretical, methodological, managerial, and policy implications of these findings.

5. Discussion and Implications

5.1 Interpreting the Findings: Deployment-Oriented Digital Disclosure, Operations, and Downside Risk

This study tests whether deployment-oriented AI/IoT disclosures (AIDI) are associated with subsequent changes in environmental and safety operations and firm outcomes (profitability and downside risk). Guided by a capability-and-appropriability perspective, digital technologies are expected to generate performance-relevant benefits when they translate into operational routines that reduce environmental and safety externalities and when industry conditions support appropriation (Ologeanu-Taddei et al., 2025; Teece, 1986; WIPO, 2024). In linked firm-year outcome panels, evidence for an AIDI → operations pathway is limited, whereas the safety–risk linkage is strong.

Outcome-specific locked-sample two-way fixed-effects models with multiplicity-adjusted inference provide no robust evidence that lagged AIDI is associated with GHGRP-based emissions intensity or workplace injury rates (TRIR); TRIR estimates are directionally negative and stable across specifications but imprecise and do not survive multiplicity adjustment.

The clearest result is a strong operations–risk linkage: lagged TRIR is consistently associated with downside risk across specifications and adjustments, consistent with firm-year exposure to disruptions and liabilities that increase downside risk (Yang et al., 2024). By contrast, associations between emissions intensity and downside risk are weaker and do not survive multiple-testing adjustment, and neither emissions intensity nor TRIR robustly predicts profitability (ROA) in the linked samples.

Moderation tests provide no reliable evidence that the AIDI–operations association varies by industry carbon salience. Although the framework predicts stronger effects in high-salience settings, interaction estimates are small and imprecise. The null results are consistent with similar within-firm slopes across industries or attenuation from limited power due to sparse disclosures and thin outcome linkages.

Taken together, the findings warrant a cautious interpretation and imply two points: (i) stronger designs and richer operational measures are needed to detect when digital deployments yield verifiable emissions and safety improvements; and (ii) workplace safety remains a consequential operational lever for managing downside exposure, even when deployment-oriented AI/IoT disclosure shows no detectable effects in this setting.

5.2 Theoretical and Conceptual Contributions

This study's theoretical contribution is intentionally bounded. It applies a capability-and-appropriability logic to clarify when deployment-oriented digital capabilities could become rent-yielding through environmental and safety operations and what standard firm-year designs can and cannot infer about that mechanism (Ologeanu-Taddei et al., 2025; WIPO, 2024).

First, the analysis advances capability-based theorizing by separating deployment-oriented digital capability signals from broad digital rhetoric and providing a reusable measurement infrastructure for mechanism-focused research (Elsayed, 2025; Roszkowska-Menkes et al., 2024). AIDI operationalizes an observable footprint—disclosed AI/IoT deployments—but the results highlight the challenges of detecting operational effects in standard firm-year panels.

Second, the findings refine appropriability-based expectations. Although the framework predicts stronger translation of deployment into operational gains where environmental stakes and rewards for verified performance are higher, no robust evidence of moderation by carbon salience is

detected. This suggests that appropriability is difficult to test with coarse firm-year moderators in disclosure-linked panels and that key mechanisms may operate at finer levels (facility, asset, process) than firm-year data capture. Accordingly, claims that digital capabilities generate rents through environmental operations should be treated as contingent, shaped by feasibility, appropriability conditions, and complementary investments (process redesign, governance, human capital) that aggregate panels capture imperfectly (Teece, 1986).

Third, the strongest empirical regularity is the safety–risk linkage: within-firm variation in workplace safety robustly predicts downside risk, whereas AIDI–operations estimates are imprecise (Yang et al., 2024). This pattern positions operational performance as central to downside protection, not only average returns. Workplace safety therefore emerges as a strategically material operational domain even when deployment-oriented digital signals do not translate into measurable operational improvements. Research on environmental performance may benefit from incorporating operational risk channels—especially safety—into theories of value protection and resilience (Tran et al., 2024; Yang et al., 2024).

The results caution against treating digitalization as a resource that systematically improves sustainability outcomes (Benuzzi et al., 2025; Ologeanu-Taddei et al., 2025). The evidence is more consistent with translation and appropriation as necessary conditions: deployments are more likely to be value-creating when embedded in measurable, verifiable routines and supported by complementary assets and governance, yet disclosure-based signals may still map imperfectly to operational change in publicly linked firm-year data (Ologeanu-Taddei et al., 2025; WIPO, 2024). Theoretical progress is therefore likely to come from sharper boundary conditions and micro-to-macro mechanisms specifying where AI/IoT is deployed, how it alters monitoring and control, and

how verification and incentives enable appropriation (Khadam et al., 2025; Papagiannidis et al., 2025; Pizzi et al., 2024).

In sum, the study contributes to strategic management theory by (i) strengthening measurement of deployment-oriented AI/IoT capability signals, (ii) clarifying the identification limits of firm-year panels for testing capability-to-operations mechanisms, and (iii) documenting a robust safety–risk channel linking workplace safety performance to downside risk (Yang et al., 2024).

5.3 Measurement and Methodological Contributions

The study's core methodological contribution is the development and validation of an open, deployment-oriented text-as-data measure of disclosed AI/IoT use in core operations. AIDI is derived from sentence-level disclosures in U.S. firm reports and targets concrete deployments—implemented systems tied to identifiable assets, processes, or operational routines—rather than generic "digital transformation" rhetoric (Elsayed, 2025; Roszkowska-Menkes et al., 2024). A stratified, hand-coded validation set indicates high labeling reliability and acceptable classifier performance, while convergent comparisons to external innovation benchmarks (e.g., AI/IoT patenting) and discriminant tests using a placebo "digital buzz" measure support interpreting AIDI as a deployment-oriented reporting footprint rather than narrative intensity or report length.

A further contribution is the separation of measurement validity from downstream estimability: even a validated deployment-disclosure measure can yield weak within-firm associations when linked outcomes are sparse, noisy, or aggregated above the operational level where deployments occur (e.g., facilities and assets). Finally, releasing the dictionary, code, and firm-year scores provides reusable infrastructure for cumulative research on digital capability–operations mechanisms. Unless otherwise noted, main inferences use the share-based measure (AIDI_share); count/index variants (AIDI_index) are used where necessary for estimability in fully saturated

specifications and for scale benchmarking, without altering the qualitative pattern of (non)rejection.

5.4 Managerial Implications

Managerial implications emphasize caution in interpreting AI/IoT deployment disclosures and highlight workplace safety as a practical lever for managing downside risk.

Do not equate AI/IoT deployment disclosure with operational improvement. In this design, deployment-oriented disclosure is not a reliable indicator of firm-year changes in emissions intensity or injury rates, even when disclosures describe concrete deployments. Operational impacts may vary by site and process, take time to surface in firm-level metrics, be confounded by shifts in production mix, asset utilization, or reporting practices, or be obscured by noise in aggregated firm-year measures. Treat "AI/IoT deployed" as an input signal—not a performance claim—and avoid using narrative deployment indicators as substitutes for verified environmental or safety gains.

Anchor AI/IoT programs in measurable, process-level routines. When the goal is environmental or safety improvement, tie deployments to ESG-critical assets and workflows by ensuring sensor calibration and data quality, clear ownership for interventions, integration into maintenance and operating procedures, and closed-loop learning (alerts → actions → verification → process redesign). Evaluate impact with granular internal metrics—facility- or line-level emissions, near-miss and incident logs, reliability measures, and audit trails—rather than disclosure volume. If improvements do not materialize, invest in complements such as governance, incentives, training, and process redesign rather than adding more "digital."

Use workplace safety as a concrete risk-management lever. Lagged TRIR robustly predicts downside risk, so strengthen safety management systems—hazard identification and controls,

maintenance discipline, training, incident investigation and learning, and leadership accountability—to reduce exposure to adverse events that matter to investors and stakeholders. Use digital tools (e.g., predictive maintenance, anomaly detection, safety analytics) to support these systems, but prioritize operational discipline; safety performance shows the most reliable link to financial downside risk in this setting.

Treat industry context as a design input, not a presumed moderator. The analysis finds no reliable carbon-salience heterogeneity in the AIDI–operations relationships, so avoid assuming AI/IoT deployments will automatically yield measurable environmental or safety improvements in either high- or low-salience sectors. Design initiatives around verifiable operational mechanisms: specify the process changes, the assets and sites affected, the decision rights and control loops, and the measurable targets used to confirm impact.

5.5 Policy and Regulatory Implications

Treat deployment-oriented digital disclosure as a screening signal, not evidence of operational improvement. In linked firm-year panels, lagged AIDI does not robustly predict subsequent emissions intensity or injury rates. Frameworks that reward narrative "digital" claims can incentivize visibility over verified operational change (Elsayed, 2025; Roszkowska-Menkes et al., 2024).

Prioritize operational measurement, comparability, and assurance. To stimulate AI/IoT applications that materially improve environmental and safety outcomes, emphasize standardized, auditable, decision-relevant metrics—especially at the facility or asset level where digital systems operate. Strengthen measurement standards for emissions and safety reporting, raise assurance expectations, and align enforcement so the payoff to verified improvement exceeds symbolic

adoption (IFRS Foundation, 2023; Pizzi et al., 2024). In sum, pair AI/IoT deployment with reporting architectures that make operational performance observable and comparable.

Enable linkable, high-resolution data infrastructures. Public firm-year linkages are often too thin to detect meaningful effects, so improve monitoring value by promoting interoperability across reporting systems (e.g., consistent facility identifiers, stable reporting boundaries over time, and machine-readable formats; OSHA, 2025; U.S. Environmental Protection Agency, 2025; U.S. Securities and Exchange Commission, 2025). This enables reliable linkage of emissions, production denominators, and safety outcomes to the operating units where AI/IoT is deployed, strengthening oversight and cumulative research while distinguishing operational impacts from reporting noise.

Do not treat digitalization as a substitute for compliance and oversight. Treat digital tools as complements to credible measurement and governance, not stand-alone solutions. Avoid policies that assume digital transformation will automatically reduce emissions or injuries; require strong monitoring and accountability to deliver measurable gains (Ologeanu-Taddei et al., 2025).

Recognize workplace safety as a financially material operational domain. Within-firm evidence links poorer safety performance to higher downside risk, reinforcing the policy value of strong safety enforcement and high-quality reporting (Yang et al., 2024). Improve reporting quality, inspections, and compliance incentives to reduce exposure to adverse events with clear economic consequences. Collectively, shift policy incentives toward verifiable, standardized, and linkable operational performance rather than narrative technology claims (IFRS Foundation, 2023; Pizzi et al., 2024).

5.6 Limitations and Future Research

Disclosure-based capability proxy (AIDI). AIDI captures reported AI/IoT deployments rather than latent capability or verified implementation. Strategic nondisclosure, investor-relations incentives, and cross-industry differences in terminology and reporting style may introduce systematic, time-varying measurement error that attenuates within-firm estimates and complicates comparisons across industries and over time, especially as the meanings of "AI," "machine learning," and "IoT" evolve.

Operational data coverage and sample intersections. Linking the text corpus to external operational datasets produces much thinner firm-year panels. For emissions intensity, GHGRP matching and the sales denominator requirement sharply reduce usable observations, and within-firm identifying variation in lagged AIDI is concentrated in a small set of mover firms, making estimates sensitive to panel composition. Injury-rate coverage is larger, but intervals remain wide and multiplicity-adjusted inference does not support a robust AIDI → TRIR association.

Identification and timing. The analyses are observational and rely on within-firm variation in lagged two-way fixed-effects models with feasible controls, placebo indices, and limited lead and event-time diagnostics. These design choices mitigate—but do not eliminate—selection, omitted-variable bias, and reverse causality (e.g., operational performance, regulatory scrutiny, or managerial quality may jointly influence both deployment disclosure and outcomes). Effects may also be delayed, localized, or nonlinear, and firm-year aggregation can mask facility- or process-level improvements.

Aggregation and organizational perimeter. Mechanisms likely operate at the facility, asset, or line level, but AIDI and outcomes are measured at the consolidated firm-year level. Deployments concentrated in a subset of sites—or occurring in joint ventures or supply-chain partners—may fail to produce detectable firm-level changes in emissions or injury rates.

Outcome scope and unobserved trade-offs. The analysis focuses on GHGRP-based emissions intensity, TRIR, ROA, and a downside-risk proxy, but does not capture implementation costs, project failures, cybersecurity and privacy incidents, or workforce and distributional consequences that may accompany AI/IoT deployments. Estimates should therefore be read as net associations in observed outcomes, not a complete benefit–harm or welfare assessment.

Future research directions

Pair deployment-oriented text measures with higher-resolution operational data. Link AIDI (or similar measures) to facility-level emissions, process metrics, maintenance logs, safety incident records, and—where feasible—IoT telemetry to better match measures to the locus of action and reduce attenuation from firm-year aggregation.

Exploit quasi-experimental variation in adoption constraints and feasibility. Use regulatory shocks and measurement mandates, shifts in enforcement intensity, staggered site-level rollouts, exogenous changes in sensor availability or connectivity, and vendor-driven implementation timelines to identify when digital deployments produce verifiable operational change rather than primarily reflecting disclosure strategy.

Clarify heterogeneity in appropriability and implementation regimes. Distinguish contexts where gains are more likely to be appropriable (e.g., tight regulation, credible verification, operational complexity) and separate in-house capability building from vendor adoption, which may differ in distinctiveness and persistence.

Extend external validity. Replicate the analysis in private firms, SMEs, and non-U.S. settings—especially across regimes with different disclosure and enforcement architectures—to assess whether identification constraints and null AIDI–operations results are context-specific. Use cross-

jurisdiction comparisons to test whether stronger verification and standardized reporting increase the detectability and economic relevance of digital deployments.

6. Conclusion

Amid climate regulation and the low-carbon transition, this study examines how firms disclose AI/IoT deployments in core operations and what those disclosures imply for environmental and safety operations and firm outcomes. Using U.S. firm reports from 2018–2024, the analysis develops and validates an open, deployment-oriented text-as-data measure—the AI/IoT Deployment Index (AIDI)—that distinguishes concrete operational deployments from generic digital rhetoric and supports mechanism-focused research at scale.

Empirically, lagged two-way fixed-effects models with multiplicity-adjusted inference provide no robust within-firm evidence that lagged AIDI is associated with subsequent GHGRP-based emissions intensity or workplace injury rates (TRIR). The clearest empirical regularity instead appears in the operations–risk relationship: across specifications and adjustments, lagged TRIR is a strong, stable predictor of downside risk, consistent with safety capturing a strategically material operational dimension tied to disruptions, incidents, and liabilities. In light of these findings, the study provides open measurement infrastructure—a validated, deployment-oriented AI/IoT disclosure index with the accompanying dictionary, code, and firm-year scores—that enables cumulative research distinguishing operational deployment from generic "digital transformation" rhetoric.

These results underscore the need for strategy research to specify boundary conditions and align measurement with the operational locus of change, while encouraging managers and policymakers to interpret AI/IoT deployment narratives cautiously and prioritize verifiable, granular operational measurement and governance—especially for workplace safety, where links to downside risk are

strongest. Future work should map deployments to facility- and process-level outcomes using designs with stronger identification.

Statements and Declarations

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Declaration of competing interests

The author declares no competing interests.

Ethics approval / consent to participate

Not applicable. The study uses publicly available firm reports and externally sourced firm-level operational datasets and does not involve human participants, personal data, or interventions.

Data availability

A blinded replication package (code and constructed indices, including AIDI) is available via an OSF view-only link:

https://osf.io/2wauy/overview?view_only=de7809f034a942a0ab0dcea03ad9ef70

Availability of any third-party datasets is subject to the original providers' licensing terms.

CRediT authorship contribution statement

Conceptualization; Methodology; Data curation; Formal analysis; Writing—original draft; Writing—review & editing; Visualization; Project administration.

References

Ang, A., Hodrick, R. J., Xing, Y., & Zhang, X. (2006). The cross-section of volatility and expected returns. The Journal of Finance, 61(1), 259–299. https://doi.org/10.1111/j.1540-6261.2006.00836.x

Benuzzi, M., Bax, K., Paterlini, S., & Taufer, E. (2025). Chasing ESG performance: How methodologies shape outcomes. International Review of Financial Analysis, 104, 104239. https://doi.org/10.1016/j.irfa.2025.104239

Bernardo, B. M. V., São Mamede, H., Barroso, J. M. P., & dos Santos, V. M. P. D. (2024). Data governance & quality management—Innovation and breakthroughs across different fields. Journal of Innovation & Knowledge, 9(4), 100598. https://doi.org/10.1016/j.jik.2024.100598

Center for Research in Security Prices. (2025, November 20). CRSP/Compustat merged database guide [PDF]. https://www.crsp.org/wp-content/uploads/guides/CRSP_Compustat_Merged_Database_Guide.pdf

Chen, X., Martínez-Ruiz, M. P., Bulmer, E., & Yáñez-Araque, B. (2025). Unpacking the black box: How AI capability enhances human resource functions in China's healthcare sector. Information, 16(8), 705. https://doi.org/10.3390/info16080705

Corporate Register. (2026). Corporate Register: Global CR reporting directory [Database]. Retrieved February 8, 2026, from https://corporateregister.com/

Dussaux, D., Agnelli, A., & Es-Sadki, N. (2023, July 19). Exploring new metrics to measure environmental innovation (OECD Environment Working Paper No. 221). Organisation for Economic Co-operation and Development. https://one.oecd.org/document/ENV/WKP(2023)13/en/pdf

El-Helaly, M. (2024). Artificial intelligence and occupational health and safety, benefits and drawbacks. La Medicina del Lavoro, 115(2), e2024014. https://doi.org/10.23749/mdl.v115i2.15835

Elsayed, N. (2025). Strategic AI disclosures and legitimacy: Impression management in UK FTSE100 annual reports. International Journal of Accounting Information Systems, 56, 100757. https://doi.org/10.1016/j.accinf.2025.100757

Engle, R. F., & Manganelli, S. (2004). CAViaR: Conditional autoregressive value at risk by regression quantiles. Journal of Business & Economic Statistics, 22(4), 367–381. https://doi.org/10.1198/073500104000000370

European Securities and Markets Authority. (2025, October 14). Materiality matters (!): Results of a fact-finding exercise on 2024 corporate reporting practices under ESRS Set 1 (ESMA32-846262651-5288). https://www.esma.europa.eu/sites/default/files/2025-10/ESMA32-846262651-5288_Fact_finding_on_materiality_disclosures_in_sustainability_statements.pdf

Financial Stability Board. (2023, October 12). Progress report on climate-related disclosures. https://www.fsb.org/uploads/P121023-1.pdf

Gartner. (2024, May 21). Gartner predicts 50% of organizations will manage hybrid cloud energy consumption with sustainability monitoring by 2026. https://www.gartner.com/en/newsroom/press-releases/2024-05-21-gartner-predicts-half-of-organizations-will-manage-hybrid-cloud-energy-consumption-with-sustainability-monitoring-by-2026

Huang, J., Wang, D. D., & Wang, Y. (2024). Textual attributes of corporate sustainability reports and ESG ratings. Sustainability, 16(21), 9270. https://doi.org/10.3390/su16219270

Huber, J., Anzengruber-Tanase, B., Schobesberger, M., Haslgrübler, M., Fischer-Schwarz, R., & Ferscha, A. (2025). Evaluating user safety aspects of AI-based systems in industrial occupational safety: A critical review of research literature. International Journal of Environmental Research and Public Health, 22(5), 705. https://doi.org/10.3390/ijerph22050705

IFRS Foundation. (2023, June). Effects analysis: IFRS Sustainability Disclosure Standards—IFRS S1 General Requirements for Disclosure of Sustainability-related Financial Information and IFRS S2 Climate-related Disclosures. https://www.ifrs.org/content/dam/ifrs/project/general-sustainability-related-disclosures/effects-analysis.pdf

IFRS Foundation. (2024, November). Progress on corporate climate-related disclosures—2024 report. https://www.ifrs.org/content/dam/ifrs/supporting-implementation/issb-standards/progress-climate-related-disclosures-2024.pdf

Intergovernmental Panel on Climate Change. (2023). Climate change 2023: Synthesis report. https://www.ipcc.ch/report/ar6/syr/

International Energy Agency. (2024, March 1). CO2 emissions in 2023. https://www.iea.org/reports/co2-emissions-in-2023

Kathan, M. C., Utz, S., Dorfleitner, G., Eckberg, J., & Chmel, L. (2025). What you see is not what you get: ESG scores and greenwashing risk. Finance Research Letters, 74, 106710. https://doi.org/10.1016/j.frl.2024.106710

Khadam, U., Davidsson, P., & Spalazzese, R. (2025). A systematic literature review on AI in IoT systems: Tasks, applications, and deployment. Internet of Things, 34, 101779. https://doi.org/10.1016/j.iot.2025.101779

Khan, M. I., Yasmeen, T., Khan, M., Ul Hadi, N., Asif, M., Farooq, M., & Al-Ghamdi, S. G. (2025). Integrating Industry 4.0 for enhanced sustainability: Pathways and prospects. Sustainable Production and Consumption, 54, 149–189. https://doi.org/10.1016/j.spc.2024.12.012

KPMG. (2025). AI's dual promise: Enabling positive climate outcomes and powering the energy transition. https://assets.kpmg.com/content/dam/kpmg/nl/pdf/2025/ai-dual-promise.pdf

Lublóy, Á., Keresztúri, J. L., & Berlinger, E. (2025). Quantifying firm-level greenwashing: A systematic literature review. Journal of Environmental Management, 373, 123399. https://doi.org/10.1016/j.jenvman.2024.123399

NewClimate Institute. (2024, April). A repository of sector-specific decarbonisation benchmarks informing 1.5°C-aligned corporate climate action (Version 1.0). https://newclimate.org/sites/default/files/2024-04/NewClimate_CorporateBenchmarksRepository_April2024.pdf

Occupational Safety and Health Administration. (2025, December). Injury tracking application (ITA) user guide. https://www.osha.gov/sites/default/files/ita_user_guide.pdf

Ologeanu-Taddei, R., Hönigsberg, S., Weritz, P., Wache, H., Mittermeier, F., Tana, S., Dang, D., Hautala-Kankaanpää, T., & Pekkola, S. (2025). The relationship of digital transformation and corporate sustainability: Synergies and tensions. Technological Forecasting and Social Change, 210, 123809. https://doi.org/10.1016/j.techfore.2024.123809

Opitz, J. (2024). A closer look at classification evaluation metrics and a critical reflection of common evaluation practice. Transactions of the Association for Computational Linguistics, 12, 820–836. https://doi.org/10.1162/tacl_a_00675

Papagiannidis, E., Mikalef, P., & Conboy, K. (2025). Responsible artificial intelligence governance: A review and research framework. The Journal of Strategic Information Systems, 34(2), 101885. https://doi.org/10.1016/j.jsis.2024.101885

Phan, T. T., Nguyen, P. V., Corvello, V., & Pham, T. T. P. (2026). AI capacity in the public sector: Pathways from the environmental context to an organizational impact. Journal of Innovation & Knowledge, 13, 100931. https://doi.org/10.1016/j.jik.2025.100931

Pizzi, S., Venturelli, A., & Caputo, F. (2024). Restoring trust in sustainability reporting: The enabling role of the external assurance. Current Opinion in Environmental Sustainability, 68, 101437. https://doi.org/10.1016/j.cosust.2024.101437

Popescu, S. M., Mansoor, S., Wani, O. A., Kumar, S. S., Sharma, V., Sharma, A., Arya, V. M., Kirkham, M. B., Hou, D., Bolan, N., & Chung, Y. S. (2024). Artificial intelligence and IoT driven technologies for environmental pollution monitoring and management. Frontiers in Environmental Science, 12, 1336088. https://doi.org/10.3389/fenvs.2024.1336088

Ramadan, M. N. A., Ali, M. A. H., Khoo, S. Y., Alkhedher, M., & Alherbawi, M. (2024). Real-time IoT-powered AI system for monitoring and forecasting of air pollution in industrial environment. Ecotoxicology and Environmental Safety, 283, 116856. https://doi.org/10.1016/j.ecoenv.2024.116856

Roszkowska-Menkes, M., Aluchna, M., & Kamiński, B. (2024). True transparency or mere decoupling? The study of selective disclosure in sustainability reporting. Critical Perspectives on Accounting, 98, 102700. https://doi.org/10.1016/j.cpa.2023.102700

Seera, M., Subramaniam, R. K., & Samuel, S. D. (2025). Collaboration or compliance? Unpacking ESG performance and carbon penalties. Sustainable Futures, 9, 100758. https://doi.org/10.1016/j.sftr.2025.100758

Shui, X., Zhang, M., Wang, Y., & Smart, P. (2024). Do climate change regulatory pressures increase corporate environmental sustainability performance? The moderating roles of foreign market exposure and industry carbon intensity. British Journal of Management, 36(1), 223–239. https://doi.org/10.1111/1467-8551.12841

Sneideriene, A., & Legenzova, R. (2025). Greenwashing prevention in environmental, social, and governance (ESG) disclosures: A bibliometric analysis. Research in International Business and Finance, 74, 102720. https://doi.org/10.1016/j.ribaf.2024.102720

Sortino, F. A., & van der Meer, R. (1991). Downside risk. The Journal of Portfolio Management, 17(4), 27–31. https://doi.org/10.3905/jpm.1991.409343

Teece, D. J. (1986). Profiting from technological innovation: Implications for integration, collaboration, licensing and public policy. Research Policy, 15(6), 285–305. https://doi.org/10.1016/0048-7333(86)90027-2

The Lens. (2023, September 3). Patent family info tab. Lens Support. https://support.lens.org/knowledge-base/family-info-tab/

Tran, V. T., Phan, D. H. B., Tee, C.-M., & Nguyen, D. T. (2024). Unmasking the carbon conundrum: How emissions impact stock price crash risk. Finance Research Letters, 64, 105443. https://doi.org/10.1016/j.frl.2024.105443

U.S. Environmental Protection Agency. (2025, December 1). Learn about the Greenhouse Gas Reporting Program (GHGRP). https://www.epa.gov/ghgreporting/learn-about-greenhouse-gas-reporting-program-ghgrp

U.S. Securities and Exchange Commission. (2025, November 28). About EDGAR. https://www.sec.gov/submit-filings/about-edgar

von Hippel, P. T., & Schuetze, B. A. (2025). How not to fool ourselves about heterogeneity of treatment effects (EdWorkingPaper No. 25-1116). Annenberg Institute at Brown University. https://doi.org/10.26300/4qtt-0w17

Wharton Research Data Services. (2026). WRDS SEC linking tables. Retrieved February 8, 2026, from https://wrds-www.wharton.upenn.edu/pages/grid-items/wrds-sec-linking-tables/

World Intellectual Property Organization. (2024). WIPO guide to trade secrets and innovation. World Intellectual Property Organization. https://doi.org/10.34667/tind.49735

Yang, M., Wang, Y., Hammer, R., & Maresova, P. (2024). Occupational health and safety management system and stock price crash risk. International Review of Financial Analysis, 93, 103141. https://doi.org/10.1016/j.irfa.2024.103141

Zhang, L., Innab, N., Shuhidan, S. M. M., Pan, Y., Zhang, Y., Som, H. M., & Alasbali, N. (2025a). Artificial intelligence-driven internet of things-based green supply chain for carbon reduction in sustainable manufacturing. Journal of Environmental Management, 389, 126170. https://doi.org/10.1016/j.jenvman.2025.126170

Zhang, Q., Hu, Z., Zhang, Z., & Zhao, R. (2025b). Carbon emission and idiosyncratic risk: Role of environmental regulation and disclosure. International Review of Financial Analysis, 105, 104419. https://doi.org/10.1016/j.irfa.2025.104419

Zhu, Y., & Li, C. (2025). Research on the impact of digital and green transformation on corporate sustainability performance from a synergistic perspective. Systems, 13(9), 820. https://doi.org/10.3390/systems13090820

Figures

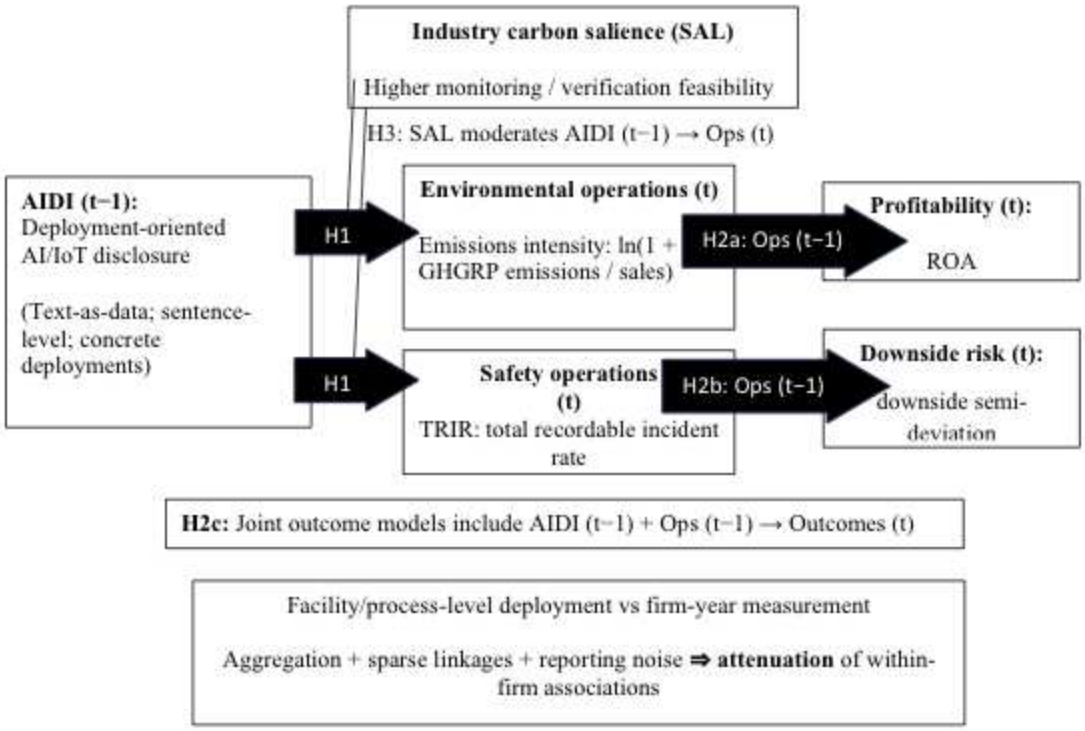


Figure 1

Conceptual framework and hypotheses.

Note. Deployment-oriented AI/IoT disclosure (AIDI_{t−1}) is linked to subsequent operational outcomes (H1: Ops_t, measured as emissions intensity and TRIR). Firm outcomes are modeled as functions of lagged operational outcomes (H2a: Ops_{t−1} → ROA_t; H2b: Ops_{t−1} → downside risk_t), as indicated on the H2 arrows. Mechanism-consistent patterns are assessed using joint models including both AIDI_{t−1} and Ops_{t−1} (H2c). Industry carbon salience (SAL) moderates the AIDI_{t−1} → Ops_t association (H3). The measurement boundary highlights attenuation arising from facility/process-level deployments measured against firm-year outcomes and sparse public linkages.

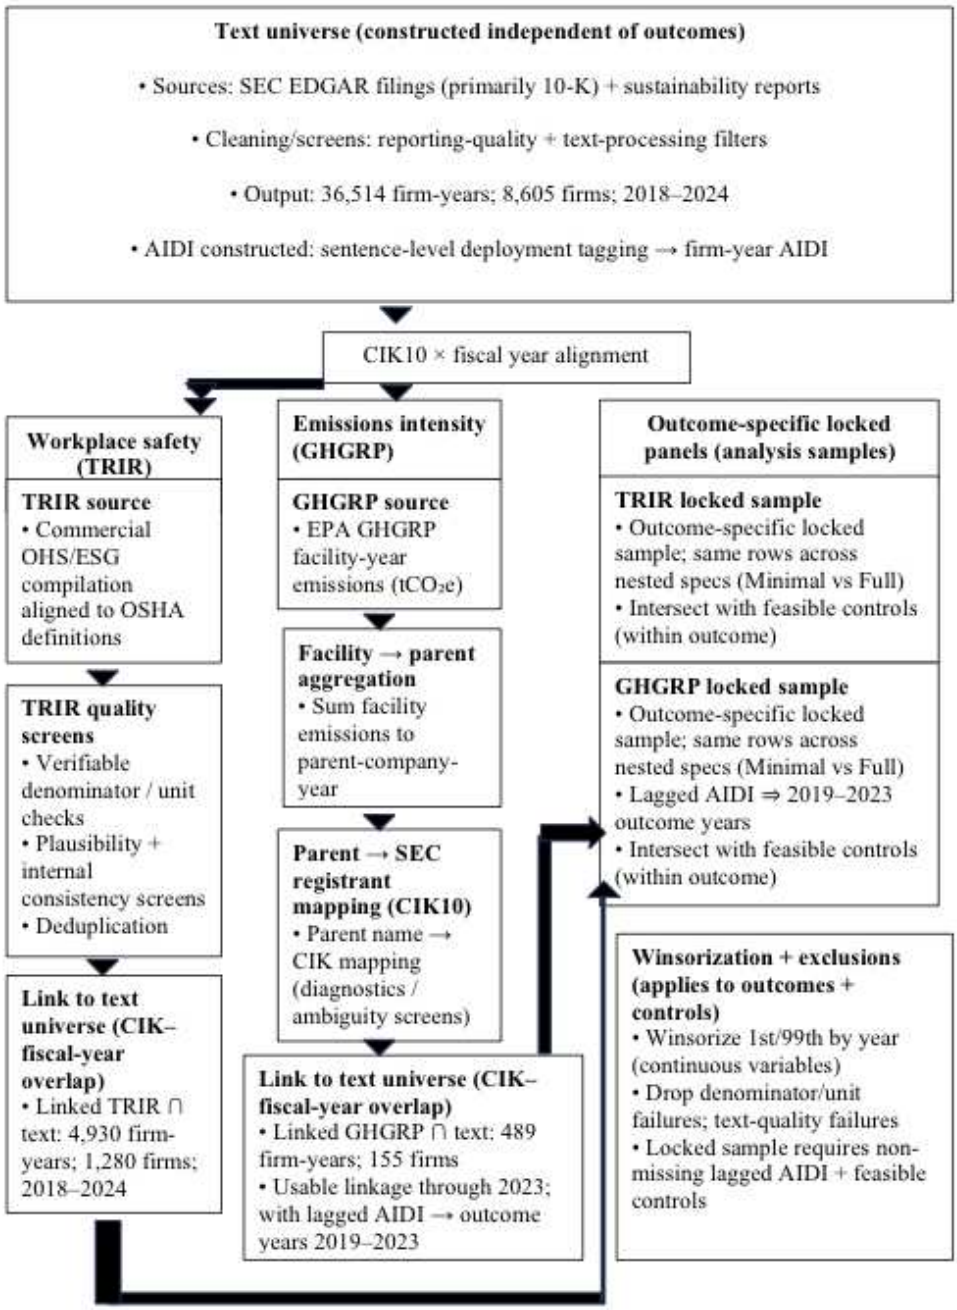


Figure 2

Data pipeline and linkage architecture: text universe → outcome-specific locked panels.

Note. The text universe is constructed from SEC EDGAR filings (primarily 10-K) and sustainability reports (2018–2024) after reporting-quality and text-processing screens (36,514 firm-years; 8,605 firms). Outcomes are linked after corpus construction using CIK10 × fiscal-year alignment to create outcome-

specific panels: (i) workplace safety (TRIR) from a commercial OHS/ESG compilation aligned to OSHA definitions, linked by CIK–fiscal-year overlap (4,930 firm-years; 1,280 firms; 2018–2024); and (ii) GHGRP emissions intensity linked via facility→parent aggregation and parent→SEC registrant (CIK10) mapping, then linked to the text universe (489 firm-years; 155 firms; usable linkage through 2023, with lagged AIDI implying outcome years 2019–2023). Analyses use outcome-specific locked samples so nested specifications are estimated on identical rows within each outcome (Minimal vs Full), intersected with feasible controls; locked samples require non-missing lagged AIDI. Continuous variables are winsorized at the 1st/99th percentiles by year, and observations failing denominator/unit checks or text-quality screens are excluded.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- SupplementaryAppendix.docx