Poslovna izvrsnost – Business Excellence
Journal of Business Performance, Quality and Organizational Excellence

Vol. 20 (2026) No. 1, 43–64

p-ISSN: 1846-3355
e-ISSN: 1848-9060

Journal website: businessexcellence.net.efzg.hr

University of Zagreb Faculty of Economics & Business

ORIGINAL RESEARCH ARTICLE

AI Disclosure Dynamics in Large Global Corporations

Serban-Vladimir Galani *

Bucharest University of Economic Studies, Bucharest, Romania
ROR ID: https://ror.org/04yvncj21, galaniserban23@stud.ase.ro https://orcid.org/0009-0008-0506-216X

George-Cristinel Rotaru

Bucharest University of Economic Studies, Bucharest, Romania
rotarugeorge18@stud.ase.ro https://orcid.org/0009-0006-9151-6062

Alexandra-Mihaela Dumitru

Bucharest University of Economic Studies, Bucharest, Romania
dumitrualexandra18@stud.ase.ro https://orcid.org/0009-0004-1061-2526

Submission: 21 April 2026 | Revision: 3 June 2026 | Acceptance: 10 June 2026
Published: 15 June 2026

DOI: https://doi.org/10.22598/pi-be/2026.1.42656

Abstract

Purpose: *This study examines how the world's largest corporations disclose artificial intelligence (AI) in public reporting and whether disclosure intensity corresponds to short-term financial and employment outcomes.* ***Design/Methodology***: *The analysis covers public reports issued from 2020 to 2024 by the top 300 companies in the 2024 Fortune Global 500. A computer-aided text analysis framework combines dictionary-based extraction, semantic-similarity validation, ChatGPT-assisted screening, and human review.* ***Findings***: *AI disclosure follows a two-phase trajectory: stable, incremental growth from 2020 to 2022, followed by rapid acceleration from 2023, largely associated with generative AI and large language models. Disclosure varies substantially across sectors and regions, with Technology, Media & Telecommunications and Financial & Business Services acting as early and intensive disclosers, while Consumer & Commerce expands later. AI disclosure frequency is not materially associated with short-term revenue, profitability, or employment changes.* ***Practical Implications***: *Disclosure counts should be interpreted cautiously as signals of corporate communication rather than direct evidence of operational AI adoption or economic impact.* ***Originality/Value***: *The study provides large-scale empirical evidence on AI disclosure dynamics and proposes a validated computational approach for distinguishing substantive AI references from generic or symbolic communication.*

Keywords: *artificial intelligence disclosure; corporate reporting; computer-aided text analysis; Fortune Global 500; strategic signaling*

JEL codes: O33; M15; L25;

* Corresponding author: Serban-Vladimir GALANI

Available online 15 June 2026

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

Data availability statement (DAS): The supplementary files supporting the findings of this study are available with the article. Additional data and methodological materials are available from the corresponding author upon reasonable request, subject to restrictions related to the size, structure, and source-specific nature of the collected corporate disclosure documents.

Author contributions (CRediT statement): Conceptualization – S.-V.G., G.-C.R.; Methodology – S.-V.G., G.-C.R.; Software – S.-V.G.; Formal analysis – S.-V.G., G.-C.R.; Investigation – S.-V.G., G.-C.R., A.-M.D.; Data curation – S.-V.G., G.-C.R.; Writing – original draft – S.-V.G., G.-C.R., A.-M.D.; Writing – review & editing – S.-V.G., A.-M.D.; Visualization – S.-V.G., G.-C.R.; Supervision – S.-V.G.

Conflict of interest: The authors declare that they have no financial, institutional, or personal conflicts of interest that could have influenced the research, analysis, interpretation, or publication of this article.

Ethics statement: This study does not involve human participants or survey data requiring ethics approval.

Sex and Gender Reporting (SAGER statement): Sex and gender were not relevant variables in this study because the analysis focuses on corporate AI disclosures, financial indicators, employment data, sectors, regions, and technology categories. Therefore, sex- or gender-disaggregated analyses were not applicable.

AI tools declaration: ChatGPT 3.5. by OpenAI was used as an assistive tool for screening and classification of AI-related disclosure sentences, as described in the Methodology section and Appendix B. InstaText software was used to support language editing, grammar correction, and stylistic refinement. All AI-assisted analytical outputs were reviewed and validated by the authors, and final analytical decisions, interpretations, and conclusions were made solely by the authors.

Cite as:
Galani, S.-V., Rotaru, G.-C., & Dumitru, A.-M. (2026). AI disclosure dynamics in large global corporations. *Poslovna izvrsnost – Business Excellence*, *20*(1), 43–64.
https://doi.org/10.22598/pi-be/2026.1.42656

1. Introduction

The past three years since the emergence of generative artificial intelligence have been marked by intense buzz and excitement. McKinsey & Company (2025) reports that the proliferation of agentic AI is increasingly recognized across global corporations, with firms reporting not only high levels of use but also actual business deployments. Despite widespread adoption claims, business challenges persist, especially regarding workflow alignment, AI effectiveness, and operational integration. AI continues to evolve, and so do the organizational challenges associated with its implementation. AI has become increasingly important in the contemporary business environment, as well as in technology, science, and everyday life. For small and medium-sized enterprises, the recent AI wave represents substantial opportunities to expand, enhance productivity, create specific operational routines, reduce costs, and support continuous improvement. However, Schwaeke et al. (2025) suggest that adoption still depends on substantial investments, technological readiness, and organizational capabilities.

A recent McKinsey & Company (2025) report examines the current state of AI deployment in companies and large corporations across various industries, confirming AI's continued presence in the business landscape. Approximately 88% of organizations use AI in at least one business function, up from 78% the previous year. However, the same report notes that only about one-third of companies are scaling AI across their businesses, while two-thirds remain in earlier phases of testing, experimentation, or piloting. This distinction is important because pilot phases are often considered a recommended step toward adoption. Uren and Edwards (2023) and Lahlali et al. (2021) emphasize that experimentation and pilot implementation are common phases in technology adoption processes. Neumann et al. (2024) similarly describe adoption as an ongoing process rather than a single point in time.

This study draws a clear and deliberate distinction between two conceptual terms often conflated in the literature: AI adoption and AI disclosure. Cubric (2020) discusses AI adoption as a process involving implementation, organizational readiness, and technology integration, while Khalid et al. (2025) connect AI with broader governance, managerial, and organizational transformation. In this study, AI adoption refers to the actual implementation, integration, and use of AI technologies within organizational processes, infrastructures, and decision systems. It encompasses the technical, organizational, and strategic dimensions of embedding AI into operational routines. Its effects are typically assessed through investment data, labor market signals, patent activity, or productivity metrics. Disclosure, in contrast, captures how firms communicate their engagement with AI to external audiences, often reflecting strategic positioning, reputation management, and legitimacy considerations, for example, in annual reports, sustainability reports, or regulatory filings. While disclosure may signal innovative intent or technological orientation, it does not necessarily provide direct evidence of operational deployment or realized value. Aldabbous and Riyath (2024) and Yang and Zhou (2025) provide useful context for understanding why equating the volume of AI communication with substantive adoption risks conflating symbolic representation with material transformation.

These two constructs are related but fundamentally distinct: disclosure may accompany, precede, or even substitute for actual adoption. A firm may disclose AI activity to signal innovation intent or technological readiness to investors without yet having operationally deployed the technology at scale. Conversely, firms with deep AI integration may communicate it sparingly, especially in competitive or regulated environments. This asymmetry means that the volume of AI communication does not map directly onto the degree of operational transformation. Cao et al. (2024) show that corporate AI disclosure can contain forward-looking information, but it should not automatically be treated as direct evidence of operational implementation. Barrios et al. (2025) similarly suggest that AI disclosure becomes more informative when supported by complementary real-world investment, particularly AI-related labor capabilities. Accordingly, AI disclosure is best understood not as a direct measure of adoption, but as a communicative proxy that reflects how firms wish to be perceived by external stakeholders with respect to AI, rather than what they have necessarily achieved internally.

From a business excellence perspective, AI disclosure is part of organizational transparency, strategic performance communication, and digital transformation governance. Examining whether such disclosure corresponds to observable performance outcomes is relevant not only for investors but also for understanding how firms communicate excellence-oriented transformation. Babina et al. (2024) show that AI-related investments can be associated with firm-level growth outcomes when measured through labor market and investment signals, while McElheran et al. (2024) demonstrate that AI use remains unevenly distributed across firms. These findings suggest that disclosure remains analytically valuable when measured longitudinally and at scale across comparable firms. Although it is an imperfect proxy, it can capture communicated adoption trajectories, strategic prioritization signals, and legitimacy-seeking behavior, all important for understanding how AI diffuses through the corporate economy.

Barrios et al. (2025) further examine the determinants and informativeness of corporate AI disclosure by linking public AI-related communication with AI-related labor investment and subsequent firm outcomes. Their findings indicate that AI disclosures are more informative when accompanied by complementary investments in AI-capable employees, while high disclosure without corresponding AI employment may signal strategic overstatement or AI washing. This supports the argument that disclosure alone is an incomplete indicator of substantive AI adoption.

The literature on technology adoption notes that adoption typically occurs in a sequence: exploration, experimentation, implementation, routinization, and eventually reconfiguration of processes and capabilities. AI is no exception. Berente et al. (2021) argue that AI differs from earlier waves of technology because of its autonomy, learning capacity, and implications for organizational decision-making. This sequence depends on technological opportunity, organizational readiness, and environmental pressures. In large companies, these pressures rarely remain abstract; they come simultaneously from competitors, regulators, customers, and capital markets. AI adoption is therefore considered not only an IT decision but also a strategic management objective. As a result, there is a growing need to study and observe AI adoption and disclosure to determine whether their expected outcomes differ from those associated with earlier technologies.

One practical challenge in the adoption literature concerns measurement. Much of what is known about AI adoption comes from surveys or interviews, which are valuable but limited by response bias, social desirability effects, and small sample sizes. Corporate narratives, by contrast, offer a useful lens for understanding how large companies communicate AI-related priorities. Annual reports, regulatory filings, investor presentations, and sustainability reports record what businesses choose to convey about strategy, risk, innovation, and transformation. Although corporate disclosure is not a perfect proxy for actual capability, it can serve as a consistent indicator of communicated adoption, strategic prioritization, and perceived legitimacy, especially when measured longitudinally and at scale.

Empirical work further shows that AI investment decisions are increasingly disclosed in corporate reports as part of broader digital transformation narratives. McElheran et al. (2024) show that measured AI use is concentrated among larger firms, suggesting that AI-related communication may function both as an informational signal and as part of broader strategic positioning. For large publicly traded firms, such disclosures may signal future growth orientation even when immediate operational impact remains limited. Consequently, strategic AI investment contains both economic and communicative dimensions of adoption.

This paper contributes to the existing literature through its systematic empirical approach, analytical framework, and use of disclosure documents from large, globally visible corporations. The study employs a computational mixed-method approach that combines computer-aided text analysis, semantic modeling, and human validation to identify, verify, and classify how large corporations disclose artificial intelligence initiatives in public reporting between 2020 and 2024.

The main aim of this paper is to clarify corporate AI disclosure by providing new insights into the industrial, technological, and organizational contexts in which AI is communicated. To structure the empirical inquiry, the study is guided by the following research questions:

- RQ1: How does AI disclosure evolve over time among the largest corporations, and what phases characterize this trajectory?
- RQ2: How does AI disclosure vary across industry sectors, and which sectors emerge as early or lower-intensity disclosers?
- RQ3: How do AI disclosure patterns differ across technology categories, including generative AI, machine learning, and robotics?
- RQ4: Is AI disclosure frequency associated with short-term changes in corporate financial performance, specifically revenue, profitability, and employment?

Based on AI references identified in corporate disclosure documents; the empirical findings show that corporate communication about artificial intelligence follows a two-stage trajectory: an initial period of modest growth followed by rapid expansion. This pronounced increase becomes evident beginning in 2023, when both the share of corporations mentioning AI and the volume of AI references rise sharply. This inflection point coincides with the broader emergence and diffusion of generative AI and large language models.

Overall, these findings suggest that AI disclosures primarily serve as instruments of strategic communication and legitimacy signaling rather than as direct indicators of realized economic transformation. Therefore, disclosure-based measures should be used cautiously as proxies for substantive AI adoption, especially in studies examining near-term performance effects.

2. Literature Review

Technological progress is rapidly transforming corporate governance, with artificial intelligence playing a central role in this shift. As Khalid et al. (2025) argue, AI is increasingly embedded in organizational functions, from routine automation to advanced analytical support for strategic decision-making. Beyond improving efficiency, AI is also influencing the ethical foundations that guide corporate behavior. Its impact extends beyond performance gains, prompting organizations to reconsider long-standing assumptions about transparency, accountability, and control. The growing reliance on intelligent systems is reshaping how decisions are formulated, assessed, and justified, thereby altering managerial responsibilities and governance structures.

The integration of artificial intelligence into governance and managerial processes has not progressed in a linear or uninterrupted manner. Uren and Edwards (2023) describe the history of AI development as a sequence of alternating "AI springs" and "AI winters," marked by periods of strong enthusiasm followed by disappointment or stagnation. Earlier waves of optimism, particularly those associated with expert systems, created high expectations about the transformative potential of intelligent technologies within organizations. Yet many of these initiatives struggled to advance beyond pilot stages or symbolic implementation. This pattern of inflated expectations followed by reassessment has shaped how firms, regulators, and stakeholders approach contemporary AI solutions, fostering both ambition and caution in the pursuit of digital transformation.

Another research stream examines AI disclosure in corporate annual reports by combining traditional textual analysis with large language model capabilities. Cao et al. (2024) show that the number of companies disclosing some form of AI engagement has increased since 2010 across different targeted categories. Their approach provides deeper insight into how firms communicate AI-related activities, while also warning that AI-related success is not guaranteed, as integration may involve multiple risks, unforeseen effects, and uncertainties. Cao et al. (2024) further find that AI disclosures may provide forward-looking information related to firm growth, operational efficiency, and risk. This supports treating disclosure as a signal with forward-looking informational content, rather than as a direct measure of operational AI adoption.

From a risk perspective, limited clarity about how organizations develop and use artificial intelligence creates significant challenges for market participants. Investors depend on accurate and comparable information to evaluate exposure to technological, legal, and reputational vulnerabilities. As Cao et al. (2024) emphasize, the current reporting environment offers little formal guidance on what should be communicated, making AI-related disclosures discretionary, uneven in depth, and difficult to verify. These conditions intensify information asymmetries and may encourage overly optimistic narratives about technological sophistication. Regulatory authorities have already signaled concerns about this tendency, particularly when companies exaggerate or misrepresent their AI capabilities to project an image of innovation.

A significant source of risk arises from the evolving regulatory landscape surrounding artificial intelligence, as policymakers continue to address ethical dilemmas, data privacy challenges, and the broader societal consequences of widespread automation. The OECD, BCG, and INSEAD (2025) highlight the need for appropriate AI policies and governance mechanisms to guide responsible development and deployment. Similarly, Cao et al. (2024) point to the importance of transparent reporting practices in reducing uncertainty around AI-related corporate risks. Policymakers need to remain alert to the risks of placing excessive confidence in the intensity or apparent sophistication of AI technologies, especially when such optimism may outpace proven capabilities or oversight mechanisms. Khalid et al. (2025) further suggest that governments are expected to promote updated policy frameworks that can keep pace with technological change.

Practitioner-oriented contributions also stress the importance of context-specific AI implementation. LeHong et al. (2025), for example, emphasize that selecting AI applications appropriate to the industry and organizational context may be as important as the implementation process itself. Executive leaders should develop a clear understanding of what AI can support, how it can be implemented, and where it can genuinely contribute to business value. The Gartner AI Opportunity Radar is therefore relevant as a practitioner-oriented tool to help organizations describe their AI ambitions in relation to concrete use cases.

Previous empirical evidence linking AI adoption to measurable corporate outcomes reveals two well-defined thematic paths. The first concerns the measurement of AI adoption, and the second addresses its downstream impact on firm performance. Babina et al. (2024), using labor market signals to measure firm-level AI investments, show that firms investing in AI experience faster growth in sales, employment, and market valuation. McElheran et al. (2024), using representative establishment-level surveys, find that measured AI use is relatively concentrated in large corporations. These findings indicate that AI adoption is uneven and that aggregate diffusion can mask substantial heterogeneity in realized outcomes.

Another body of research emphasizes that even when AI adoption generates value, it may follow the logic of the productivity J-curve. According to McElheran et al. (2025), early AI implementation may involve experimentation, integration costs, data re-engineering, and workflow redesign, which can temporarily flatten measured productivity before gains appear. Their US census-linked evidence identifies patterns consistent with initial adjustment frictions followed by later performance improvements. These findings imply that the absence of short-run effects is theoretically plausible even under genuine adoption.

These findings are directly relevant to the present empirical study because the main independent variable is AI disclosure intensity, not verified operational adoption. Prior evidence indicates that disclosure-based measures can diverge from real activity: firms may communicate AI to signal innovativeness, legitimacy, or readiness to investors, even when deployments remain limited or benefits are not yet captured by accounting outcomes. In this context, McElheran et al. (2025) suggest that short-term performance effects may be weak during periods of adjustment, while OECD, BCG, and INSEAD (2025) emphasize that AI-related organizational value often depends on complementary capabilities, governance, and implementation conditions. Therefore, contemporaneous correlations between AI-related communication and revenue, profit, or headcount changes may be weak or unstable.

Rech et al. (2026) offer a methodologically grounded perspective on the disclosure-adoption gap by constructing firm-level AI adoption proxies from two complementary sources: textual disclosures derived from corporate annual reports using a 72-term AI lexicon, and patent data. Their analysis, based on a panel of Chinese A-share non-financial firms from 2007 to 2023, demonstrates that disclosure-based AI indicators can improve out-of-sample financial distress forecasting beyond standard accounting fundamentals. This supports the informational value of AI proxies as early-warning signals. Importantly, however, Rech et al. (2026) acknowledge that such proxies are indirect measures. They capture how firms represent their AI engagement in public documents, rather than whether AI has been operationally deployed or has generated realized value. AI proxies may reflect short-run transition costs and organizational disruption associated with early-stage implementation, but they may also reflect disclosure incentives rather than demonstrated capability. This distinction is particularly consequential in periods of rapid AI diffusion, when the density of AI references in corporate reports tends to increase even among firms that remain in early experimental phases. The framework proposed by Rech et al. (2026) thus reinforces the conceptual position adopted in the present study: disclosure-based AI measures are analytically useful and informationally non-trivial, but they function as communicative proxies for adoption rather than as its direct equivalent.

Building on this literature, the present empirical analysis is framed as a conservative external validity check. It tests whether the communicated acceleration of AI, especially GenAI and LLM references, is accompanied by near-term changes in performance metrics. This framing is consistent with Babina et al. (2024), who find positive firm-level outcomes when AI is measured through investment and labor-market signals, and with McElheran et al. (2025), who show that productivity effects may appear only after adjustment frictions. Since the most credible positive estimates in prior work often rely on adoption proxies tied to skills, innovation, and production data rather than disclosure counts, a disclosure-based proxy should be expected to be noisier and more sensitive to signaling. This is precisely the distinction highlighted in the present study by separating adoption from disclosure.

3. Methodology

Our study employs a computational mixed-method research design to examine how large corporations present information related to artificial intelligence in their disclosure documents. Given the scale and heterogeneity of corporate disclosures – including multiple document types and reporting years – human reading is infeasible and could result in inconsistency among different reviewers. To ensure replicable and scalable analysis, we use computer-aided text analysis (CATA), a method widely used in management and organizational research for structured text extraction and classification (Short et al., 2010). In this study, AI disclosure refers to corporate communication about AI in public reports; AI references are the validated textual units used to measure disclosure; disclosure breadth refers to the share of firms with at least one AI reference; and disclosure intensity refers to the number of validated AI references per firm-year.

CATA provides a framework for identifying sentences associated with AI terminology; however, dictionary-based approaches alone may overlook semantically relevant content or misclassify ambiguous references. To address these limitations, we integrate semantic similarity models that capture contextual meaning beyond exact keyword matches, following recent advances in natural language processing (Reimers & Gurevych, 2019, 2020). This hybrid design enables us to detect AI-related statements even when firms use diverse terminology, while maintaining precision through multiple validation layers (Manning et al., 2008). Overall, the methodological approach is designed to balance analytical rigor and interpretability.

The dataset consists of corporate disclosure documents from the top 300 companies listed in the 2024 edition of the Fortune Global 500 (Fortune, 2024). We collected publicly available reports covering 2020–2024, including annual reports, sustainability and environmental, social, and governance (ESG) reports, integrated reports, proxy statements, and other relevant regulatory filings. Documents were primarily sourced from the official websites of the respective companies and, when necessary, from the websites of the stock exchanges on which they are listed.

To ensure comparability and data quality, several inclusion and exclusion criteria were applied. For a limited number of Chinese state-owned or state-controlled conglomerates, no consolidated English-language disclosures were available. In these cases, we included reports from relevant listed subsidiaries within the conglomerate, as English reports were publicly accessible. When only non-English documents were available, we excluded them because the current version of the analytical software operates exclusively on English-language text. Future extensions of the research may incorporate multilingual processing or pre-translation pipelines to address this limitation.

For companies whose fiscal year does not align with the calendar year, documents were assigned to the year with the most activity. For example, a report covering a fiscal year ending in March 2025 was attributed to the 2024 reporting year. In total, the processed dataset comprises 2,305 documents, amounting to approximately 240 million words and 9.8 million sentences. A complete inventory of all documents included in the analysis is available in Supplementary File S1 (Document Coverage Matrix).

The analysis used a multi-stage text-processing algorithm to extract, validate, and structure AI-related sentences. The application combines standard document processing tools with semantic modeling techniques to improve precision. (1) Documentation was collected in PDF format and processed with two complementary extraction libraries: pdfplumber (Singer-Vine, n.d.) for structured text retrieval and PyMuPDF (fitz) for high-fidelity parsing of complex layouts (Artifex Software, Inc., 2023). To preserve context, the system extracts each sentence along with the two preceding and two following sentences. When PDF documents contained corrupted text layers or embedded images, pages were converted to high-resolution images at 300 DPI and processed using Tesseract optical character recognition (Tesseract OCR, n.d.), an open-source engine. (2) Extracted text is scanned using a predetermined set of terms associated with AI concepts, including keywords and multi-word patterns. This step filters for potentially relevant segments, consistent with established CATA practices (Short et al., 2010). (3) To capture semantically relevant content beyond keyword matches, candidate sentences are evaluated using a sentence transformer model (all-MiniLM-L6-v2), which generates high-dimensional embeddings for both extracted segments and reference descriptions of core AI concepts (Reimers & Gurevych, 2019, 2020). Cosine similarity scores are computed using scikit-learn, and a conservative threshold ensures high precision. (4) After semantic validation, sentences are checked to confirm that AI terminology is used in an organizational context. Generic digitalization language, aspirational statements, or regulatory boilerplate lacking AI specificity are removed. (5) Corporate reports often repeat the same initiatives in different sections or documents. To avoid artificially inflating the frequency of AI disclosures, semantically similar sentences are clustered using cosine similarity-based grouping. All keywords, regex patterns, false-positive filters, and context validators used to identify potentially relevant AI-related sentences before semantic validation, AI-assisted screening, and human review are documented in Supplementary File S2 (AI Disclosure Keyword Dictionary).

After the initial run of the software, we observed a substantial number of false-positive sentences. To mitigate this, we introduced an additional validation layer using ChatGPT (OpenAI). The exact prompt used for ChatGPT-assisted screening and classification is provided in Appendix B. ChatGPT was used as a secondary screening and classification layer to assess whether each candidate sentence represented a genuine AI disclosure, complementing semantic similarity by capturing contextual cues that automated rules may overlook. All ChatGPT-assisted classifications were reviewed by the authors for final inclusion. In addition, a systematic random sample of 952 observations, representing 10.0% of the 9,528 ChatGPT-flagged false positives, was independently double-coded by two authors to validate the model's classification performance. Details are provided in Supplementary File S3 (False-Positive Validation). Inter-rater agreement reached 98.5%, with Cohen's Kappa of 0.74, indicating substantial agreement (Landis & Koch, 1977). From the reviewed observations, 96.4% were confirmed as false positives. While the initial taxonomy focused on business applications and strategic dimensions of AI disclosure, it provided limited insight into the technologies used. To address this gap, ChatGPT was used to assign each validated sentence to one of eight predefined technology categories: (1) general/unspecified AI, (2) machine learning/deep learning, (3) generative AI, (4) large language models (LLMs), (5) computer vision, (6) NLP/speech, (7) edge AI/responsible AI governance, and (8) robotics/autonomous systems. The application extracts the full sentence and its surrounding ±2 sentences to ensure accurate classification. In both stages, ChatGPT served strictly as an assistive tool. Final classifications were made by the authors, who manually reviewed the outputs to ensure alignment with the study's conceptual framework. This hybrid human-AI approach was necessary for efficiently analyzing the extracted data.

To analyze cross-industry differences in AI disclosure, the research consolidated the original Fortune industry classifications into six master sectors. This aggregation serves both methodological and conceptual objectives, as a limited number of sector indicators reduces model dimensionality. Collapsing industries into six master sectors yields more stable coefficient estimates and improves inference reliability. For the clusters, we used the Global Industry Classification Standard (GICS®) as a benchmark. Ultimately, we estimate bivariate and multivariate specifications relating to AI reference counts to annual changes in revenue, net profitability, and employment to test whether disclosure frequency is associated with financial performance. Lagged models are included to test whether disclosure in year t predicts outcomes in year t + 1, and nonlinear threshold specifications are used to assess whether effects concentrate among heavy disclosers.

4. Results

4.1. AI Disclosure

By identifying AI-related references, including robotics and autonomous systems when contextually linked to AI, this study aims to determine the disclosure rate of artificial intelligence among top global companies, identify which sectors were early or late disclosers, assess which technology categories were most represented, examine which country clusters focused more on communicating their AI engagement, and evaluate whether disclosure frequency is associated with observable changes in employment or financial performance. As noted in the conceptual framing of this paper, the purpose is to test whether disclosure serves as a communicative proxy for adoption, not as its equivalent. The disclosure rate reflects how firms represent their AI engagement externally; it does not directly confirm the scale, maturity, or operational impact of their AI deployments.

Table 1
AI Disclosure Evolution

Year	Companies with at least one AI reference	Companies with available reports	Companies with at least one AI reference (%)	Total AI References	AI references per company with at least one AI reference
2020	181	284	63.73%	1,371	7.57
2021	180	284	63.38%	1,517	8.43
2022	183	284	64.44%	1,502	8.21
2023	207	284	72.89%	2,292	11.07
2024	224	284	78.87%	2,992	13.36

Note. Authors' own research.

The evolution of AI-related disclosures between 2020 and 2024 is shown in Table 1, measured by the percentage of companies with at least one AI reference in their public reports and the total volume of AI references per year. Among the 284 firms with available reports, AI disclosure remains relatively stable from 2020 to 2022, with approximately 63–64% of firms referencing AI at least once. During this period, the total number of AI references increases only modestly, indicating incremental progress rather than transformative change.

A notable shift occurs from 2023 onward, driven by the rise of GenAI and LLMs. The proportion of firms mentioning AI rises to 72.89% in 2023 and further to 78.87% in 2024. This increase is accompanied by growth in total AI references, which nearly doubles between 2022 and 2024. The growth in reference volume outpaces the increase in the number of disclosing firms, indicating that the post-2022 period is marked not only by broader disclosure but also by substantially higher disclosure intensity among disclosers. This pattern suggests that, from 2023 onward, AI became a strategic theme, reflected in both broader firm-level disclosure and more frequent, detailed mentions in company reports. This timing aligns with the widespread commercialization and public visibility of generative AI technologies, which appear to have accelerated companies' incentives to communicate AI-related activities to external stakeholders.

4.2. AI Disclosure by Sector

To examine cross-industry differences in AI disclosure, Fortune's industry categories were consolidated into six clusters, as shown in Appendix A. These clusters are: Industrial & Mobility (55 companies), Energy & Materials (63 companies), Financial & Business Services (72 companies), Technology, Media & Telecommunications (TMT) (35 companies), Consumer & Commerce (37 companies), and Healthcare & Life Sciences (22 companies). This consolidation reduces sectoral noise while maintaining economically meaningful distinctions, enabling robust comparisons across industries.

Sectoral differences in AI disclosure intensity and the proportion of companies disclosing AI between 2020 and 2024 are shown in Figure 1 and Figure 2. By 2020, AI disclosure was present across all sectors, though with significant variation. TMT and Financial & Business Services had high initial disclosure, indicating that AI was already incorporated into corporate reporting and strategic communication.

Figure 1

AI Disclosure Breadth by Sector (2020–2024)

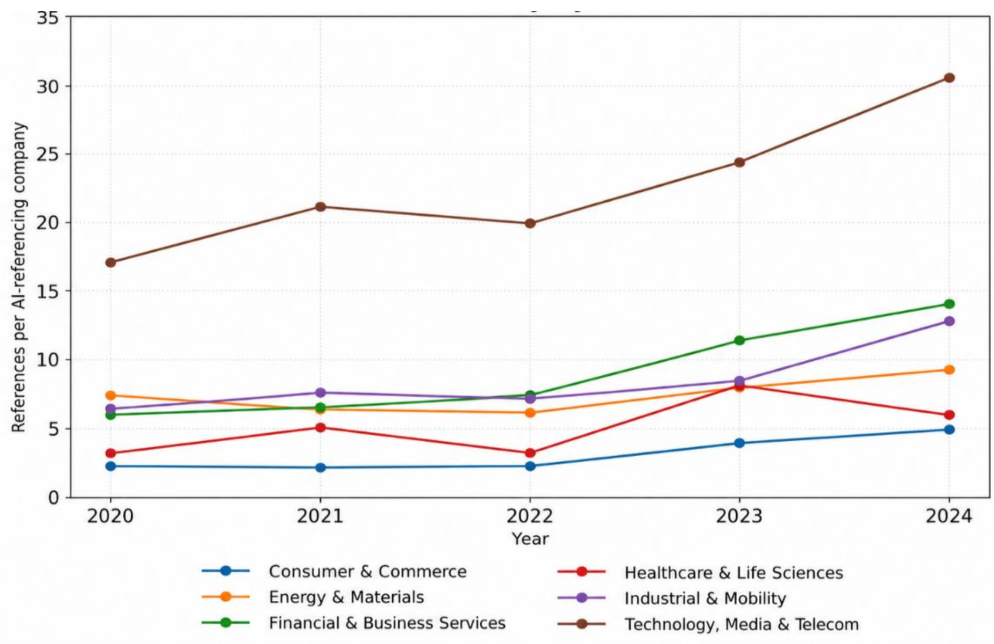

Note. Authors' own research.

Figure 2
AI Disclosure by Sector (2020–2024)

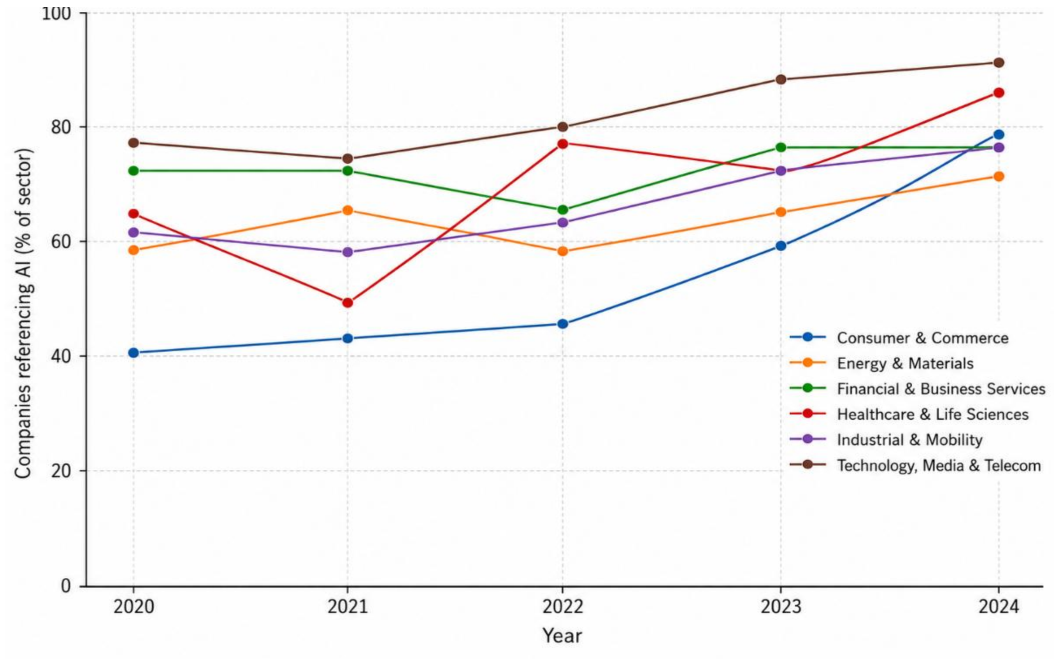

Note. Authors' own research.

However, disclosure intensity was relatively modest, suggesting that AI was treated as a background development rather than a strategic narrative. In contrast, Energy & Materials and Industrial & Mobility showed moderate initial disclosure rates, reflecting their higher capital intensity and longer implementation cycles. Consumer & Commerce and Healthcare & Life Sciences lagged, with fewer firms referencing AI, consistent with experimental use cases.

The data do not indicate an immediate disclosure response to the commercial launch of ChatGPT 3.5 on November 30, 2022. Disclosure rates increased only marginally across most sectors and, in some cases, even decreased. This pattern suggests that, before 2023, AI disclosure was mainly driven by gradual AI-related reporting and communication rather than strategic repositioning or external signaling. Notably, TMT continues to show higher disclosure intensity than other sectors, highlighting its role as an early discloser of AI.

A shift occurs in 2023 across nearly all industries. Reference rates increased in Financial & Business Services, Industrial & Mobility, and Energy & Materials. This acceleration coincides with an increase in total AI references, indicating that not only are more companies disclosing AI, but they are also discussing it more extensively in their financial reports. The rise in disclosure volume outpaces the growth in the number of disclosing firms, suggesting that 2023 marks a transition from AI exploration toward strategic communication.

By 2024, AI disclosure was high across all sectors. TMT leads, exhibiting exceptionally high references per firm. One outlier is the Consumer & Commerce sector, the "late mass adopter," with the largest disclosure jump from 2020 to 2024 (+37.9 percentage points), but references per adopter remained relatively low. Another is Financial & Business Services, a "high-intensity signaler," with disclosure rates almost flat but references per adopter increasing significantly (+7.9). Industrial & Mobility shows delayed reference behavior due to higher capital expenditure and integration costs.

4.3. AI Disclosure by Technology Category

Based on the text identified in the reports and the surrounding context, each reference was labeled in one or more of eight categories: general/unspecified AI, computer vision, edge AI/responsible AI governance, generative AI (GenAI), large language models (LLMs), machine learning/deep learning, NLP/speech, and robotics/autonomous systems. Some references were labeled in multiple categories, as LLMs are part of GenAI, and GenAI is part of artificial intelligence. Each category is kept distinct to analyze AI as a whole while allowing for a closer examination of its most frequently mentioned subcategories.

The evolution of AI-related disclosures by technology type is presented in Figure 3 and Figure 4, with the former showing the share of companies disclosing each AI category and the latter showing references by technology among companies with references.

The results reveal a clear shift toward generative AI, LLMs, and AI in general. The share of companies referencing AI in general terms increases steadily from 56% in 2020 to 77% in 2024, while references per firm nearly double over the same period, from 6.5 to 12.0. The sharp acceleration after 2022 suggests that firms are expanding not only AI-related disclosure but also their business emphasis on AI, positioning it as a core element of corporate strategy. Categories such as Machine Learning/Deep Learning and Robotics/Autonomous Systems exhibit relatively stable disclosure rates over the entire sample period. Machine Learning disclosure fluctuates within a narrow range (32–39%), and disclosure intensity remains remarkably constant (2.6–2.9 references per firm). Robotics follows a similar pattern, with disclosure stabilizing around one-third of firms.

Figure 3

AI Disclosure by Technology Category (2020–2024)

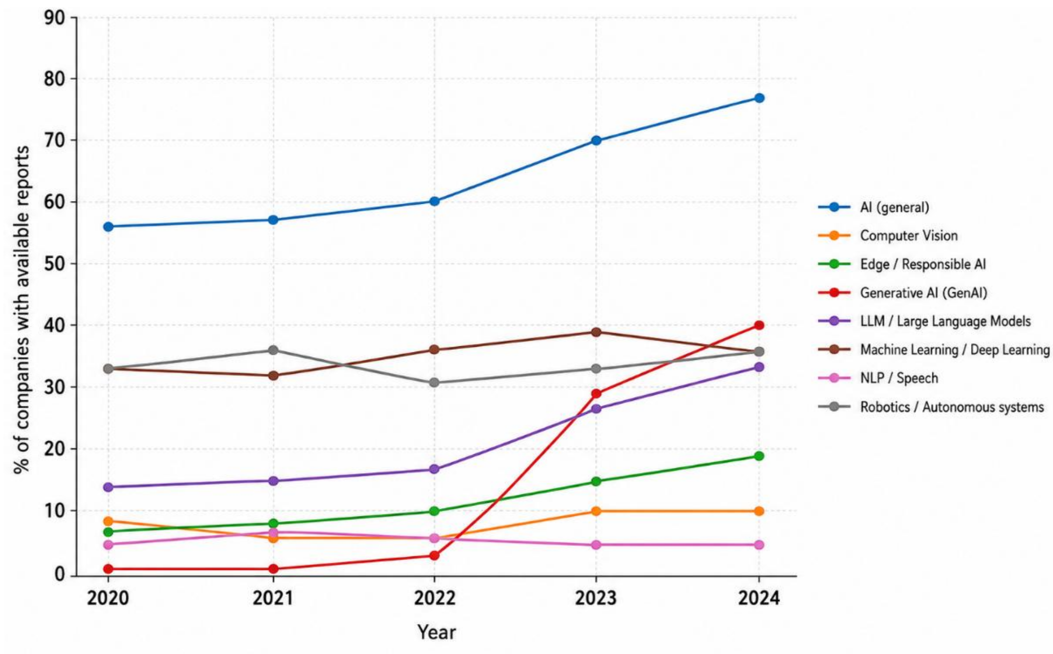

Note. Authors' own research.

Figure 4

AI References by Technology Category Among Companies with AI References (2020–2024)

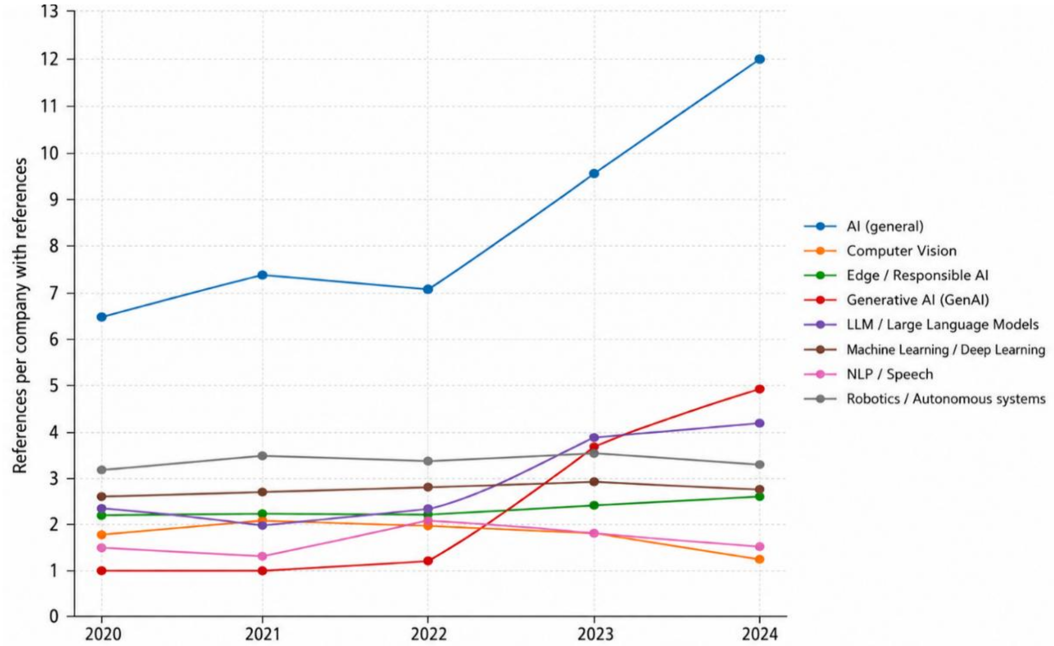

Note. Authors' own research.

The results reveal a clear shift toward generative AI, LLMs, and AI in general. The share of companies referencing AI in general terms increases steadily from 56% in 2020 to 77% in 2024, while references per firm nearly double over the same period, from 6.5 to 12.0. The sharp acceleration after 2022 suggests that firms are expanding not only AI-related disclosure but also their business emphasis on AI, positioning it as a core element of corporate strategy. Categories such as Machine Learning/Deep Learning and Robotics/Autonomous Systems exhibit relatively stable disclosure rates over the entire sample period. Machine Learning disclosure fluctuates within a narrow range (32–39%), and disclosure intensity remains remarkably constant (2.6–2.9 references per firm). Robotics follows a similar pattern, with disclosure stabilizing around one-third of firms.

In contrast, Generative AI and LLMs have seen significant increases in reference rates since 2023. Before 2022, these categories were referenced by fewer than 5% of firms. In 2023, references jumped sharply: GenAI accounted for 30% of firms and LLMs for 28%. Disclosure frequency also increased rapidly, reaching 4.9 references per company for Generative AI and 4.2 for LLMs in 2024. This simultaneous expansion may indicate a transition from experimentation to the development of commercially viable products and strategic communication. Additionally, LLMs appeared earlier than GenAI in corporate disclosures, suggesting that LLMs entered as a capability while GenAI later emerged as a strategic, board-level framing.

Edge and Responsible AI references increase steadily over the period, rising from 7% of firms in 2020 to 19% in 2024. Unlike Generative AI, this growth is incremental rather than abrupt, consistent with an institutionalization process driven by regulatory scrutiny, governance concerns, and risk management requirements.

The results reveal a two-speed evolution of AI disclosures. Established AI technologies (Machine Learning, Robotics) show stable patterns with little change in disclosure frequency. In contrast, novel technologies (Generative AI and LLMs) have experienced rapid diffusion since 2022, signaling a qualitative shift in how companies conceptualize and communicate AI. Notably, only 11 companies mentioned GenAI in their corporate reports before 2023: Deutsche Telekom, Nissan Motor, Enel, Tencent Holdings, Microsoft, Banco Bradesco, Meta

Platforms, Huawei Investment & Holding, ArcelorMittal, Assicurazioni Generali, and Mizuho Financial Group.

4.4. AI Disclosure by Country

When looking at the disclosure by country, five clusters were created: China, Europe (excluding Russia), Japan, South Korea & Taiwan, the United States of America (USA), and the rest of the world.

Regional differences in AI disclosure intensity and in the share of companies disclosing AI across country clusters from 2020 to 2024 are illustrated in Figure 5 and Figure 6. The results reveal substantial regional heterogeneity in both the timing and nature of AI disclosure practices. Chinese companies exhibit steady referencing rates, increasing from 58% in 2020 to 67% in 2024. While disclosure breadth grows incrementally, disclosure intensity increases substantially, from 9.1 to 20.3 references per company. The post-2022 acceleration in references suggests a shift toward more intensive AI-related corporate communication.

Europe shows high disclosure rates throughout the study period, with almost three-quarters of firms referencing AI as early as 2020. Disclosure intensity increases gradually, from 7.2 to 12.8 references per company, suggesting early normalization of AI within European firms, followed by incremental intensification. The Korea–Japan–Taiwan cluster is characterized by consistently high disclosure relative to the number of companies that referenced AI, with references per company among the highest across all regions, rising from 10.7 in 2020 to 18.0 in 2024. This pattern indicates deep AI-related disclosure, consistent with advanced manufacturing, electronics, and robotics ecosystems where AI-related capabilities were strategically salient.

Figure 5

AI Disclosure Intensity by Country 2020–2024

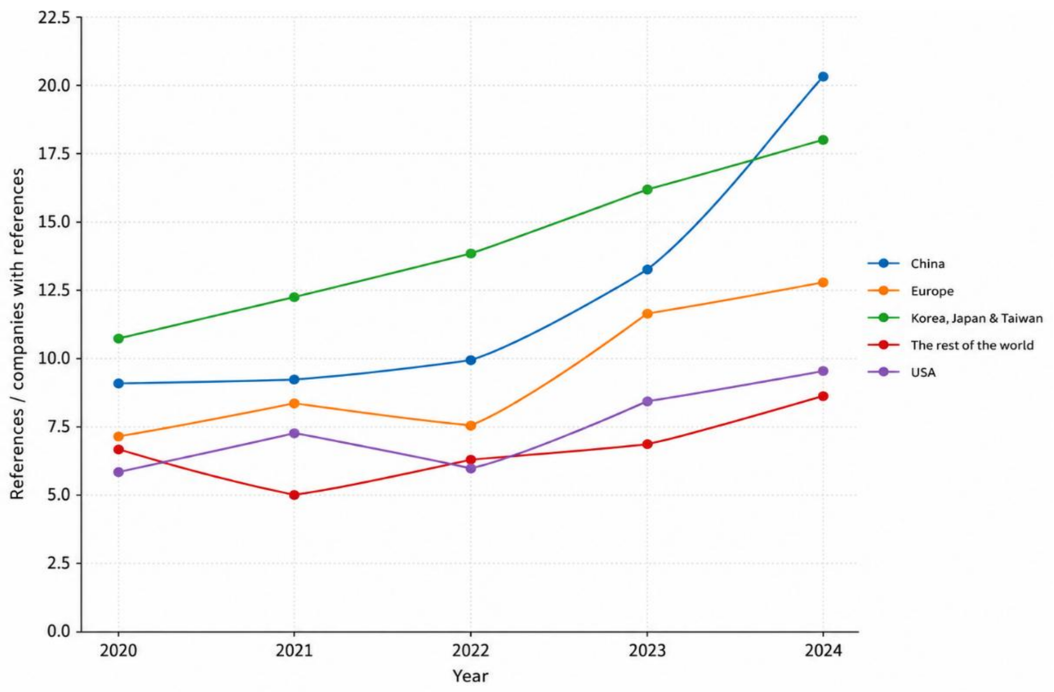

Note. Authors' own research.

Figure 6

AI Disclosure of Companies by Country 2020–2024

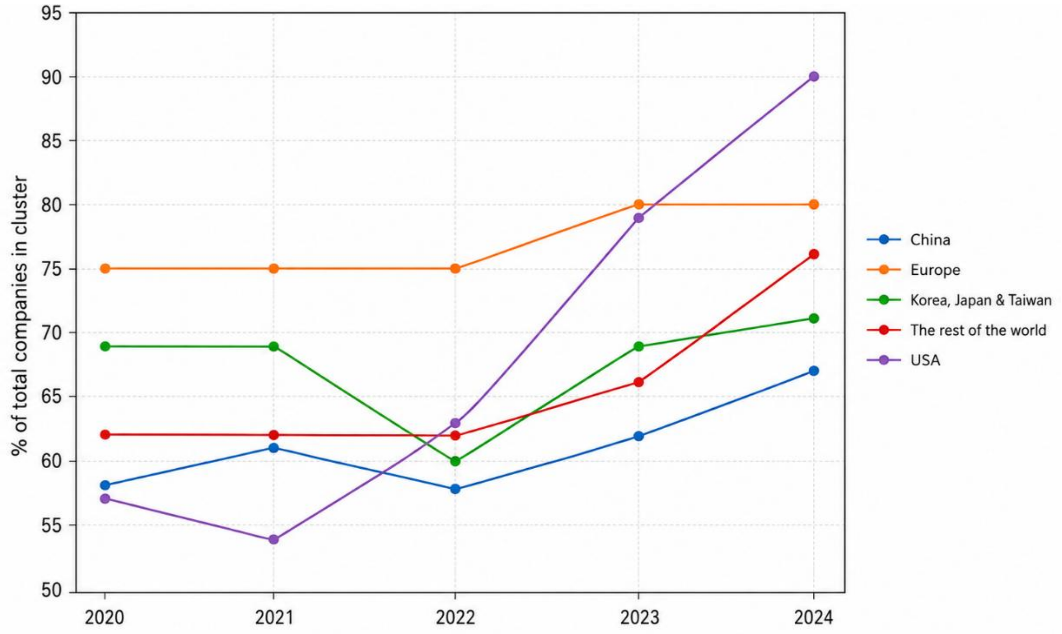

Note. Authors' own research.

The United States exhibits a distinct pattern. Disclosure breadth is relatively low and even declines slightly between 2020 and 2021, before accelerating sharply after 2022. By 2024, 90% of US companies reference AI, the highest rate of all. Firms in the Rest of the World cluster exhibit moderate disclosure rates and shallower disclosure depth, consistent with uneven access to AI capabilities.

Overall, the results highlight a clear distinction between breadth and disclosure intensity across regions. The United States leads in rapid late-stage disclosure, Europe exhibits early signs of potential adoption with steady intensification, and China and East Asia display comparatively lower or fluctuating breadth rates but markedly higher disclosure frequency among adopters. US companies show comparatively early GenAI signaling, pivoting from ML to LLM to GenAI language earlier than others. Chinese and Japanese firms show high AI presence, particularly in machine learning (ML) and robotics, but generative AI (GenAI) appears later and is framed more as an applied capability. European firms, due to earlier and stricter regulation, place stronger emphasis on responsible AI, edge AI, and governance and trust language. Cross-national differences reflect not only technological capability but distinct disclosure cultures: market-driven economies accelerate GenAI narratives, industrial economies delay naming until operational readiness, and regulated environments embed GenAI within governance frameworks.

4.5. Correlation Between Disclosure of AI and Business Performance

4.5.1. Correlation between AI referencing and revenue change

To assess whether AI-related disclosure frequency can be associated with financial performance, we examined the bivariate relationship between annual revenue percent change and the number of AI references per firm-year. The Pearson correlation between revenue growth and AI references is $r = 0.003$, and the rank-based Spearman correlation is $\rho = -0.026$, indicating no clear monotonic relationship. Consistent with these correlations, Figure 7 shows that the fitted bivariate regression line is effectively flat, with an estimated slope of approximately 0.005 percentage points of revenue growth per additional AI reference.

Figure 7

Bivariate Association: AI References vs Revenue Change

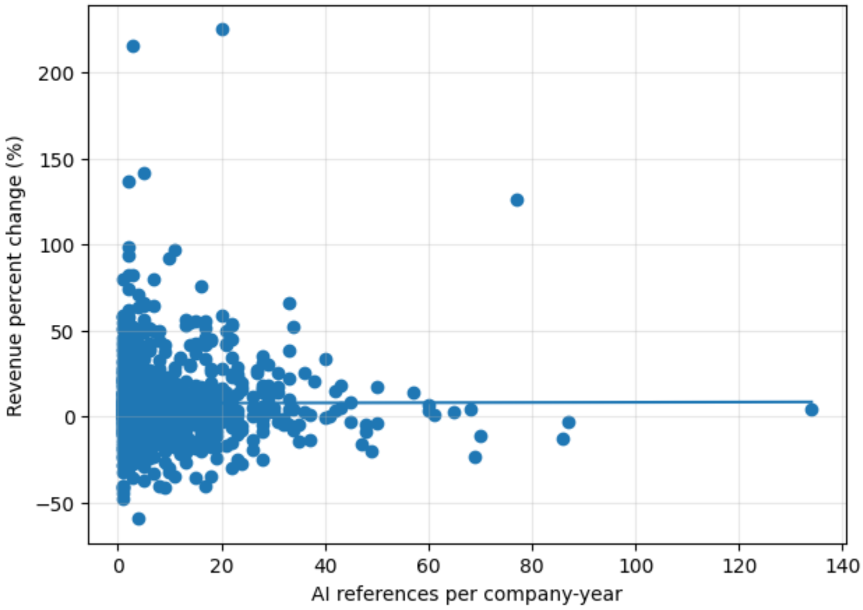

Note. Authors' own research.

Revenue growth varies widely across the entire range of AI reference data, with high- and low-growth observations present at both low and high levels of AI disclosure. As such, these results provide no evidence that companies that highlight AI more often in their disclosures experience higher revenue growth.

To better understand whether there is any correlation, we conducted a lagged model to examine the percent change in revenue in year t+1 from AI disclosure frequency in year t. We found no evidence of a positive association. Estimates are small and weakly negative but not statistically significant with company-clustered standard errors (β = −0.00038, *SE* = 0.00152, *p* = .804 in the lag specification with year/sector/country fixed effects). These results suggest that AI disclosure frequency is not a reliable predictor of short-term future revenue growth. A nonlinear threshold effects analysis further found no robust evidence of threshold effects, as being in the top decile or top 5% of AI disclosure does not consistently predict higher revenue growth.

In a technology-specific analysis, GenAI references are the only category that correlates with future revenue growth by a narrow magnitude. In lagged specifications with year, industry, and country fixed effects and company-clustered standard errors, GenAI disclosure frequency in year t is weakly positively associated with revenue percent change in year t+1 (β = 0.071, *p* = .007 for log(1+GenAI references in t)). Because the predictor is logged, this coefficient implies negligible economic impact: moving from 1 to 10 GenAI references is associated with merely 0.12 percentage points higher revenue growth in the following year. As it is a single specification across multiple hypothesis tests, it may reflect statistical noise, and we interpret this association with caution. Therefore, more research is required once more annual data on GenAI becomes available.

4.5.2. Correlation between AI disclosure and profitability change

To evaluate whether AI disclosure frequency is associated with changes in profitability, we examine the bivariate relationship between the percent change in net profits published by Fortune and the number of AI references per company per year. The association is rather weak, as the Pearson correlation is r = 0.066, and the Spearman correlation is ρ = 0.026, indicating little evidence of a systematic relationship. The data shows substantial dispersion in profitability results across the full range of AI reference counts, with several outliers. This weak and dispersed relationship is illustrated in Figure 8.

Figure 8

Bivariate Association: AI References vs Profitability Change

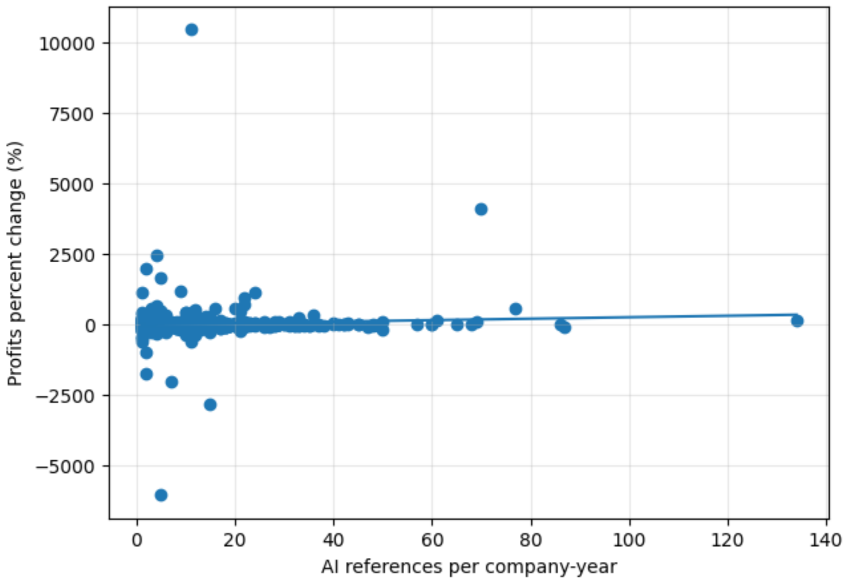

Note. Authors' own research.

Lagged, nonlinear, and category-specific tests provide limited evidence that AI disclosure frequency is associated with changes in profitability. Neither overall AI reference frequency nor threshold indicators predict next-year profitability change once year, sector, and country fixed effects are included.

To address the heavy tails in net profitability changes, we assess robustness to outlier treatment by winsorizing profitability change at the 1st/99th and 5th/95th percentiles. Findings remain unchanged; the association between AI reference intensity and profitability change is weak (Spearman $\rho \approx 0.03$, $p > .35$ across specifications). Even when considering technology type, neither GenAI nor LLM disclosures show statistically significant associations with profitability change.

4.5.3. Correlation between AI disclosure and employee growth

To examine whether AI disclosure is correlated with workforce expansion or contraction, the research relates AI references per firm-year to year-on-year employee growth. The relationship is negligible: Pearson's r = −0.017 (p = .635) and Spearman's ρ = .059 (p = .108), with substantial dispersion in employment growth at all levels of AI disclosure. The bivariate fit (R^2 ≈ 0.0003) is essentially flat. This negligible relationship is illustrated in Figure 9. Robustness checks by winsorizing employee growth at the 1st/99th and 5th/95th percentiles confirm that the null association is not driven by extreme observations.

In lagged models predicting employee growth in year t+1 from disclosure intensity in year t, log(1+AI references) is statistically insignificant (β = −0.012, p = .989). Category-specific models similarly show no observable association for specific technologies such as GenAI or LLM disclosures. One relevant finding is that firms in the top 5% of AI disclosure intensity in year t exhibit lower employee growth in t+1 (≈ −7.6 percentage points, p = .040), consistent with restructuring dynamics rather than broad-based headcount effects.

Figure 9

Bivariate Association: AI References vs Employee Growth

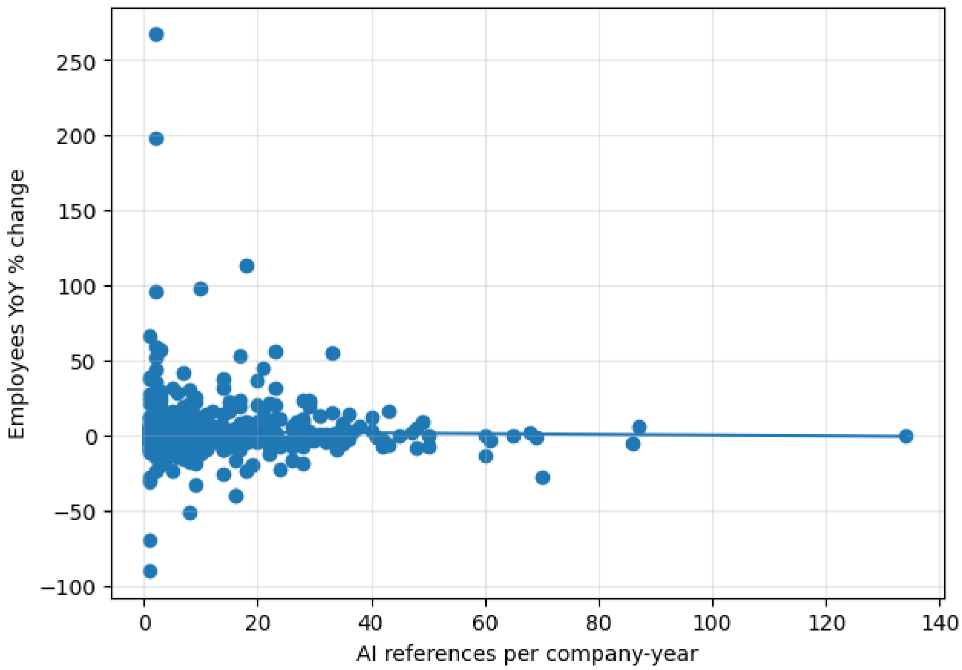

Note. Authors' own research.

5. Discussion

The findings contribute to the growing literature on AI disclosure, corporate communication, and technology signaling. The two-phase trajectory identified – gradual growth from 2020 to 2022, followed by sharp acceleration from 2023—directly reflects the broader AI landscape. The 2023 inflection point coincides with the widespread commercialization and public visibility of generative AI, suggesting that external technological shocks can quickly reshape corporate narrative strategies, even when operational implementation lags.

The lack of robust and economically meaningful associations between AI disclosure and short-term financial performance – across revenue, profitability, and employment – is consistent with several theoretical perspectives. First, it aligns with the J-curve hypothesis of technology adoption (McElheran et al., 2025), which suggests that early-stage AI investments may incur adjustment costs that temporarily suppress measurable returns. Second, it highlights the signaling function of disclosure: firms may communicate about AI not necessarily to describe operational realities, but to manage stakeholder expectations and legitimize their strategic positioning. Third, it supports the view that disclosure-based measures are noisier proxies than investment-based or skill-based measures of actual AI adoption (Babina et al., 2024).

The heterogeneity observed across sectors and regions further enriches the discussion. TMT's dominance in AI disclosure intensity is expected, given its proximity to the underlying technologies. However, the pattern in Financial & Business Services – stable disclosure breadth but sharp intensification – suggests a strategic reorientation following the emergence of GenAI, particularly in customer-facing applications and compliance governance. Consumer

& Commerce's late but rapid expansion indicates a democratization effect, as AI narratives become accessible to a broader set of firms with wider technology availability.

Regional differences in AI disclosure reflect institutional and cultural factors beyond technological capability. The US pattern of late but rapid disclosure breadth expansion, combined with early GenAI narrative adoption, suggests a market-driven signaling dynamic responsive to investor pressures. The European pattern of early breadth but gradual intensification, with strong emphasis on Responsible AI and governance language, reflects the influence of regulatory frameworks such as the EU AI Act. Chinese firms' high disclosure intensity among a relatively stable group of disclosing firms may indicate concentrated, system-level AI communication rather than broad communicative diffusion.

These findings suggest that AI disclosure, as currently practiced by large corporations, serves primarily as a strategic communication tool rather than a transparent window into operational AI deployment. This has important implications for investors, regulators, and researchers who rely on disclosure-based measures to assess AI adoption and its effects.

6. Conclusion

The findings show that AI disclosure occurs in two distinct phases. Between 2020 and 2022, references to AI increased gradually. From 2023 onward, disclosure accelerated in both the number of companies and the frequency of references, coinciding with the rise of generative AI and large language models. Notably, the growth in disclosure intensity outpaced the increase in the number of disclosing firms, suggesting that post-2022 dynamics may reflect a narrative-focused approach rather than purely new operational AI adoption.

Despite this rapid shift in AI disclosure, we found no evidence that disclosure frequency is statistically associated with short-term changes in revenue, net profitability, or employment across contemporaneous, lagged, nonlinear, or category-specific specifications. AI reference frequency does not clearly predict firm-level economic outcomes. The only exception is generative AI disclosure, which shows a statistically significant but very small association with next-year revenue growth; however, this association may reflect statistical noise. Therefore, more research is needed to determine whether a robust association exists once more annual data on generative AI become available.

Employment effects are similarly not evident, with only extremely high disclosure levels coinciding with subsequent workforce contraction, consistent with restructuring rather than generalized labor displacement. This may suggest that AI is communicated as adding organizational complexity rather than simply substituting labor, as companies increasingly present AI as a multi-layered capability portfolio.

These findings suggest that AI disclosure primarily reflects corporate strategic communication and signaling rather than realized economic impact. Disclosure-based measures should therefore be interpreted cautiously, as they do not provide direct evidence of operational AI adoption, particularly in short-term performance analyses. Regarding RQ1, AI disclosure follows a two-phase trajectory, with gradual growth from 2020 to 2022 and sharp acceleration from 2023 onward. Regarding RQ2, disclosure varies across sectors, with TMT and Financial & Business Services emerging as early and intensive disclosers, while Consumer & Commerce expands later. Regarding RQ3, generative AI and large language models drive the post-2022 acceleration, whereas machine learning and robotics remain comparatively stable. Regarding RQ4, AI disclosure frequency is not materially associated with short-term changes in revenue, profitability, or employment.

This study has several limitations that suggest areas for further research. First, corporate disclosure is an imperfect proxy for AI-related engagement, as it reflects stated priorities rather than the actual scope or effectiveness of deployment. To better distinguish symbolic communication from substantive adoption, future research could combine disclosure-based indicators with additional data, such as AI-related capital expenditures, patent activity, workforce skill composition, or internal usage metrics. Second, the study examines short-term outcomes, even though performance benefits may appear over longer time horizons. Longitudinal studies could assess whether AI disclosure predicts medium- or long-term productivity improvements. Another limitation is that the empirical design is associational rather than causal. Although fixed effects for industry, country, and year are included, endogeneity issues may still be present. Future research could employ quasi-experimental methods, exogenous shocks, or regulatory changes to strengthen causal inference.

The article is relevant to UN Sustainable Development Goals: 9 INDUSTRY, INNOVATION AND INFRASTRUCTURE; 16 PEACE, JUSTICE AND STRONG INSTITUTIONS

References

Aldabbous, N., & Riyath, M. I. M. (2024). Role of AI and digital transformation in corporate reporting: A bibliometric review. *International Business Research, 17*(5), 52–73. https://doi.org/10.5539/ibr.v17n5p52

Artifex Software, Inc. (2023). *PyMuPDF* (Version 1.23.0) [Computer software]. PyPI. https://pypi.org/project/PyMuPDF/1.23.0/

Babina, T., Fedyk, A., He, A., & Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics, 151*, Article 103745. https://doi.org/10.1016/j.jfineco.2023.103745

Barrios, J. M., Campbell, J. L., Johnson, R. G., & Liu, C. (2025). *Signals or smoke? The determinants and informativeness of corporate artificial intelligence (AI) disclosures* (SSRN Scholarly Paper No. 5133107). SSRN. https://doi.org/10.2139/ssrn.5133107

Berente, N., Gu, B., Recker, J., & Santhanam, R. (2021). Managing artificial intelligence. *MIS Quarterly, 45*(3), 1433–1450. https://doi.org/10.25300/MISQ/2021/16274

Cao, Y., Liu, M., Qiu, J., & Zhao, R. (2024). *Information in disclosing emerging technologies: Evidence from AI disclosure* (SSRN Scholarly Paper No. 4987085). SSRN. https://doi.org/10.2139/ssrn.4987085

Cubric, M. (2020). Drivers, barriers and social considerations for AI adoption in business and management: A tertiary study. *Technology in Society, 62*, Article 101257. https://doi.org/10.1016/j.techsoc.2020.101257

Fortune. (2024). *Fortune Global 500*. https://fortune.com/ranking/global500/

Khalid, F., Srivastava, M., & Toumi, F. (2025). AI technology adoption and corporate financial misconduct. *Journal of Financial Reporting and Accounting*. Advance online publication. https://doi.org/10.1108/JFRA-06-2025-0474

Lahlali, M., Berbiche, N., & El Alami, J. (2021). How enterprise must be prepared to be "AI first"? *International Journal of Advanced Computer Science and Applications, 12*(5), 346–351. https://doi.org/10.14569/IJACSA.2021.0120542

Landis, J. R., & Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics, 33*(1), 159–174. https://doi.org/10.2307/2529310

LeHong, H., Selassie, B., Cribbs, J., Mesaglio, M., & Scheibenreif, D. (2025, February 6). *Gartner AI Opportunity Radar: Set your enterprise's AI ambition* (G00801057). Gartner.

Manning, C. D., Raghavan, P., & Schütze, H. (2008). *Introduction to information retrieval*. Cambridge University Press. https://doi.org/10.1017/CBO9780511809071

McElheran, K., Li, J. F., Brynjolfsson, E., Kroff, Z., Dinlersoz, E., Foster, L., & Zolas, N. (2024). AI adoption in America: Who, what, and where. *Journal of Economics & Management Strategy, 33*(2), 375–415. https://doi.org/10.1111/jems.12576

McElheran, K., Yang, M.-J., Kroff, Z., & Brynjolfsson, E. (2025). *The rise of industrial AI in America: Microfoundations of the productivity J-curve(s)* (CES Working Paper No. CES-WP-25-27). Center for Economic Studies, U.S. Census Bureau. https://doi.org/10.2139/ssrn.5036270

McKinsey & Company. (2025, November 5). *The state of AI: Global survey 2025*. QuantumBlack AI by McKinsey. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai

Neumann, O., Guirguis, K., & Steiner, R. (2024). Exploring artificial intelligence adoption in public organizations: A comparative case study. *Public Management Review, 26*(1), 114–141. https://doi.org/10.1080/14719037.2022.2048685

OECD, BCG, & INSEAD. (2025). *The adoption of artificial intelligence in firms: New evidence for policymaking*. OECD Publishing. https://doi.org/10.1787/f9ef33c3-en

Rech, F., Musa, H., Šebeňa, M., & Tuo, S. J. (2026). *Forecasting financial distress in dynamic environments: AI adoption signals and temporally pruned training windows*. arXiv. https://doi.org/10.48550/arXiv.2512.02510

Reimers, N., & Gurevych, I. (2019). Sentence-BERT: Sentence embeddings using Siamese BERT-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing* (pp. 3982–3992). Association for Computational Linguistics. https://doi.org/10.18653/v1/D19-1410

Reimers, N., & Gurevych, I. (2020). *SentenceTransformers: Multilingual sentence, paragraph, and image embeddings using BERT & Co.* https://www.sbert.net

Schwaeke, J., Peters, A., Kanbach, D. K., Kraus, S., & Jones, P. (2025). The new normal: The status quo of AI adoption in SMEs. *Journal of Small Business Management, 63*(3), 1297–1331. https://doi.org/10.1080/00472778.2024.2379999

Short, J. C., Broberg, J. C., Cogliser, C. C., & Brigham, K. H. (2010). Construct validation using computer-aided text analysis (CATA): An illustration using entrepreneurial orientation. *Organizational Research Methods, 13*(2), 320–347. https://doi.org/10.1177/1094428109335949

Singer-Vine, J. (n.d.). *pdfplumber* [Computer software]. GitHub. https://github.com/jsvine/pdfplumber

Tesseract OCR. (n.d.). *Tesseract documentation*. https://tesseract-ocr.github.io/

Uren, V., & Edwards, J. S. (2023). Technology readiness and the organizational journey towards AI adoption: An empirical study. *International Journal of Information Management, 68*, Article 102588. https://doi.org/10.1016/j.ijinfomgt.2022.102588

Yang, T., & Zhou, N. (2025). Artificial intelligence and the quality of corporate accounting information disclosure. *Finance Research Letters, 85*, Article 108136. https://doi.org/10.1016/j.frl.2025.108136

Appendix A

Economic Sectors Mapping

Original Sector	6-Sector Master
Retail	Consumer & Commerce
Specialty Retailers	Consumer & Commerce
Wholesale / Wholesalers	Consumer & Commerce
Trading	Consumer & Commerce
Consumer Goods	Consumer & Commerce
Household Products	Consumer & Commerce
Apparel	Consumer & Commerce
Food, Beverages & Tobacco	Consumer & Commerce
Technology	Technology, Media & Telecom
Electronics	Technology, Media & Telecom

	Continued
Telecommunications	Technology, Media & Telecom
Media	Technology, Media & Telecom
Healthcare / Health Care	Healthcare & Life Sciences
Industrials	Industrial & Mobility
Automotive	Industrial & Mobility
Motor Vehicles & Parts	Industrial & Mobility
Transportation	Industrial & Mobility
Engineering & Construction	Industrial & Mobility
Construction	Industrial & Mobility
Energy	Energy & Materials
Materials	Energy & Materials
Chemicals	Energy & Materials
Financials	Financial & Business Services
Business Services	Financial & Business Services
Conglomerate	Financial & Business Services

Note. Original Fortune industry categories were consolidated by the authors into six master sectors to reduce sectoral fragmentation and support cross-sector analysis. Original industry classifications are based on the Fortune Global 500 industry categories (Fortune, 2024).

Appendix B

ChatGPT Prompt

"You are an expert analyst of corporate AI disclosures and financial data. Analyze the uploaded document and identify every occurrence related to AI keywords (AI, Artificial Intelligence, Machine Learning, ML, Generative AI, GenAI, LLM, NLP, Computer Vision, Robotics). For each occurrence: (1) Classify it as TRUE_POSITIVE — real reference to artificial intelligence — or FALSE_POSITIVE — where the string appears inside unrelated words, is mentioned as part of another acronym not related to AI, is generic marketing language without actual technological meaning, or refers to a different concept. (2) If TRUE_POSITIVE, label the technology type mentioned using at least one of these categories: GENERATIVE_AI, MACHINE_LEARNING, LLM, NLP/SPEECH, COMPUTER_VISION, ROBOTICS/AUTOMATION, AI_GENERAL. Rules: Only label TRUE_POSITIVE if the sentence clearly refers to artificial intelligence technologies or applications. Be conservative and rely on context."