AI Adoption and Stock Performance in S&P 500 Firm Disclosure

ABSTRACT

This study examines whether AI-related language in mandatory corporate disclosures functions as a credible strategic signal that investors incorporate into firm valuation. Drawing on signaling theory, we evaluate whether the intensity of AI-related terminology in the Management Discussion and Analysis (MD&A) section of 10-K filings predicts future stock performance for S&P 500 firms. Using a lexicon-based measure of AI discourse and multi-horizon return models, we link Item 7 textual content to one-, three-, and five-year forward stock returns, controlling for past performance. The results show no significant association between AI disclosure intensity and future returns across any horizon, suggesting that narrative references to AI do not operate as priced market signals in this context. This outcome may indicate that managers are less transparent or overly symbolic in communicating AI-related strategies, reducing the credibility of disclosure-based signals. The study contributes to research on technological signaling and narrative disclosure by identifying boundary conditions under which language-based indicators fail to influence market expectations. Methodologically, it offers a transparent and replicable text-analysis pipeline integrating natural language processing with financial modeling.

Keywords:

AI adoption, Firm Disclosure, MD&A, 10-K Filings, Textual Analysis, Stock Performance

1. INTRODUCTION

Artificial intelligence (AI) is increasingly positioned as a transformative force capable of reshaping organizational processes, innovation pathways, and competitive advantage across industries. Firms frequently reference AI in strategic communications, shareholder reports, and regulatory filings to convey technological preparedness and strategic vision. However, scholarly and practitioner communities alike continue to debate whether such discourse reflects substantive adoption efforts or merely symbolic posturing in response to market pressures (Nambisan et al., 2019). This ambiguity raises fundamental questions about the relationship between AI disclosure, organizational strategy, and capital market interpretation.

Regulatory filings, particularly the annual Form 10-K submitted to the U.S. Securities and Exchange Commission (SEC), are critical sites for understanding how firms communicate strategic priorities. The Management Discussion and Analysis (MD&A, Item 7) section plays a central role in reducing information asymmetry by summarizing managerial expectations, planned initiatives, and perceived risks. A growing literature demonstrates that textual elements of the MD&A convey incremental information that investors incorporate into valuation decisions (Azmi Shabestari & Romero, 2024; Karapandza, 2016; Loughran & McDonald, 2020). Yet, despite increasing references to AI in these filings, there is limited evidence on whether the language itself constitutes meaningful information from the perspective of capital markets.

This uncertainty can be framed through the lens of signaling theory. Signaling theory proposes that actors with private information attempt to communicate their quality, capabilities, or intentions through observable signals, and that receivers evaluate these signals based on credibility, costliness, and context (Spence, 1978). In the context of AI, MD&A language may serve as a strategic signal intended to demonstrate innovation orientation, technological readiness, or

alignment with future industry trajectories. If investors view such language as credible and costly to imitate, AI-related terminology should reduce information asymmetry and shape expectations of future performance. Conversely, if the signal is perceived as cheap, imprecise, or lacking evidence of deployment, investors may rationally discount it.

The contemporary digital transformation environment intensifies this tension. Firms face increasing pressure to demonstrate alignment with high-velocity technological change (Hinings et al., 2018) and to leverage general-purpose technologies such as AI as catalysts for growth (Vial, 2019). Disclosure thus becomes not only an information mechanism, but also a stage on which technological legitimacy and strategic identity are performed (Nambisan et al., 2019). This creates conditions where AI discourse may proliferate even when underlying adoption is uneven or exploratory. In such settings, textual claims risk becoming decoupled from operational reality, creating space for symbolic adoption dynamics consistent with the broader digital transformation literature (Sebastian et al., 2020). Understanding how markets interpret these signals is therefore vital for assessing the boundary conditions of strategic communication.

Despite rapid advances in computational text analysis, financial linguistics, and digital transformation research, we know relatively little about how topic-specific claims about AI in mandatory reports relate to investor interpretation. Research using keyword frequency analysis indicates that carefully selected lexical proxies can capture organizational attributes that are otherwise difficult to measure (Bellstam et al., 2021; Humphreys & Wang, 2018). Furthermore, studies show that narrative content can contain predictive information about financial outcomes (Bodnaruk et al., 2015; Li, 2010) and that forward-looking language may be linked to market reactions (Karapandza, 2016). However, the existing literature does not isolate whether AI-specific

discourse, as opposed to general forward-looking tone or sentiment, functions as a priced signal in capital markets.

This paper addresses this gap by examining whether AI-related language in the MD&A section of 10-K filings serves as a credible strategic signal from the perspective of financial markets. Using a lexicon-based measure of AI discourse in Item 7 filings for S&P 500 firms, we quantify the intensity of AI-specific terminology and link it to subsequent market performance, while accounting for past return dynamics. Rather than assuming that AI mentions reflect realized capability, this design treats disclosure language itself as a theoretically meaningful construct: a potential signal that may succeed or fail depending on how markets evaluate its informativeness.

This approach contributes to strategic management research by situating AI discourse within established theories of information asymmetry and signaling. It also extends disclosure research by focusing on topic-specific language rather than broad sentiment or readability measures, and by contextualizing digital transformation narratives as strategic communication devices rather than pure reflections of operational reality. Finally, it offers a transparent, replicable methodological pipeline that integrates domain-specific lexicons, natural-language-processing preprocessing, and publicly accessible market data.

In doing so, this study advances an emerging research agenda at the intersection of technological signaling, digital transformation, and capital market interpretation. If AI language functions as a credible signal, investors should adjust expectations when such discourse appears in filings. If it does not, this would indicate that markets either discount such language or require more granular or costly evidence before updating beliefs. In either case, the findings have implications for managers, analysts, and policymakers concerned with how narrative strategy shapes economic outcomes.

The remainder of the paper proceeds as follows. Section 2 reviews the literature on strategic disclosure, digital transformation narratives, and keyword-based text analysis in financial settings. Section 3 describes the dataset construction and the lexicon-based measurement of AI discourse. Section 4 outlines the empirical framework and variable construction. Section 5 presents the results and interprets their theoretical implications. Section 6 concludes by identifying limitations and directions for future research.

2. LITERATURE REVIEW

The use of textual disclosures as strategic communication has become central to understanding how firms convey information under conditions of uncertainty. When firms operate in dynamic environments, narrative elements of disclosure can provide meaningful insight into strategic direction, technological positioning, and perceived capability (Hinings et al., 2018; Vial, 2019). Prior research demonstrates that the content and structure of MD&A language can influence investor beliefs and valuation outcomes because it captures managerial perspectives unavailable in financial statements alone (Karapandza, 2016; Loughran & McDonald, 2020).

Text analysis of corporate disclosures has become an important methodology in the fields of accounting, finance, and strategic management. Among various sources, the U.S. Securities and Exchange Commission (SEC) Form 10-K serves as a rich and standardized document containing companies' annual performance, risk factors, strategies, and forward-looking statements (Bodnaruk et al., 2015; Cooper et al., 2022; Dyer et al., 2017; Harris, 2024; Karapandza, 2016; Liu et al., 2020; Mushtaq et al., 2022). Researchers have used computational linguistics techniques to analyze these disclosures to understand companies' behavior and results (Alduais, 2024; Azmi Shabestari & Romero, 2024; Bochkay et al., 2023; Du et al., 2024; Loughran & McDonald, 2020; Mao et al., 2024; Ravula, 2021).

While there are various approaches, including latent Dirichlet allocation (LDA) (Blei et al., 2003; Bochkay et al., 2023; Du et al., 2024; Dyer et al., 2017) and sentiment analysis (Chan & Chong, 2017; Du et al., 2024; Kouloumpis et al., 2011; Mao et al., 2024; Mushtaq et al., 2022), this review focuses on the keyword frequency analysis approach, which is appropriate when researchers use a dictionary or list of predefined terms. For this study, technology-related keywords taken from previous literature (Bellstam et al., 2021; Hinings et al., 2018; Kraus et al., 2022; Nambisan et al., 2019) will serve as a proxy to measure a company's level of technology adoption.

Keyword frequency from the 10-Ks can be used to construct a financial constraint index. Textual proxies can outperform traditional accounting measures when measuring subtle constructs (Bodnaruk et al., 2015). Frequency analysis can also examine the pricing of financial constraints and strengthen the validity of text-based indicators in empirical financial models (Buehlmaier & Whited, 2018). Text-based innovation measures derived from patents and 10-Ks show strong predictive validity for R&D output and firm performance. This approach shows how often innovation-related terms appear in annual reports (Bellstam et al., 2021). Together, these findings establish precedent for using dictionary-based measures as proxies for underlying strategic constructs.

The 10-K report is divided into several sections, notably Item 1 ("Business"), Item 1A ("Risk Factors"), and Item 7 ("Management Discussion and Analysis"). These sections have been shown to be particularly rich in forward-looking disclosures (Azmi Shabestari & Romero, 2024; Bochkay et al., 2023; Cooper et al., 2022; Dyer et al., 2017; Liu et al., 2020). There is a significant relationship between textual features and company performance, suggesting that textual signals in these sections are meaningful for strategic analysis (Azmi Shabestari & Romero, 2024). Analyzing

the frequency of technological and strategic terms in annual reports can be used to predict company performance (Alduais, 2024; Azmi Shabestari & Romero, 2024; Mushtaq et al., 2022).

If AI-related language operates as a credible strategic signal, it should convey information about firms' technological orientation that investors incorporate into expectations of future value creation (Loughran & McDonald, 2011, 2020). Under signaling theory, managers with genuine capability should reference AI in disclosure because doing so reflects costly internal investments, organizational restructuring, and complementary resource alignment (Hinings et al., 2018; Nambisan et al., 2019). Investors, in turn, should price these signals to the extent they indicate differentiation and competitive advantage.

Based on this reasoning, we propose the following hypothesis:

H1: Firms that exhibit greater intensity of AI-related language in the MD&A section of 10-K filings will experience higher future stock returns, controlling for past performance.

Because the valuation impact of strategic signals may not be uniform across investment horizons, and because markets may process information at different speeds, we test this hypothesis across three forward-looking windows:

H1a: AI-related language intensity is positively associated with one-year forward returns.

H1b: AI-related language intensity is positively associated with three-year forward returns.

H1c: AI-related language intensity is positively associated with five-year forward returns.

Together, these hypotheses reflect the expectation that AI language carries informational value and thus may function as a priced signal in capital markets.

However, no existing study isolates AI-specific language in mandatory filings to determine whether it functions as a priced strategic signal or a symbolic narrative device. This paper addresses this gap by examining whether AI-related terminology in Item 7 of the Form 10-K serves

as a credible informative signal in capital markets, using lexicon-based measurement and multi-horizon performance modeling.

3. DATA & METHODS

The sample consists of firms listed in the S&P 500 index for which complete 10-K filings and corresponding market data are available. Firm identifiers (ticker and CIK) are obtained from the SEC EDGAR index. We focus on the 2021 filing for each firm to align textual information with forward-looking market horizons. Textual data are collected from the Form 10-K annual report, specifically Item 7 (Management Discussion and Analysis), which contains forward-looking strategic content (Azmi Shabestari & Romero, 2024; Cooper et al., 2022).

To construct return-based variables, we use January 1, 2021, as the reference date. If the reference date is not a trading day, the nearest trading day within ±7 days is used. We retrieve prices for three pre-periods (1, 3, 5 years prior) and three post-periods (1, 3, 5 years after), which enables the construction of forward (dependent) and past (control) return measures.

Item 7 text is extracted, converted to plain text, and tokenized using a standard natural-language-processing pipeline: lowercase transformation, removal of non-alphabetical characters, stopword removal, lemmatization, and token filtering for length (≤2 characters removed). This ensures that the frequency distribution of tokens reflects semantically meaningful language.

To measure AI-related disclosure, we construct a domain-specific lexicon from prior research on digital transformation and AI terminology (Bianchini et al., 2025; Damioli et al., 2025; Zeng et al., 2025). The lexicon includes AI-associated terms that appear in managerial discourse (e.g., “algorithm”, “neural”, “automation”, “predictive”, “machine learning”, “AI”, “analytics”, “data-driven”). The final AI lexicon consisted of 255 terms, which, when applied to the corpus, produced over 4.6 million tokens across 500 firm-level documents.

For each firm:

$$AI_rate_i = \frac{AI\ hits_i}{total\ tokens_i}$$

We standardize this measure to obtain a comparable signal score:

$$AI_rate_z_i = \frac{AI\ hits_i - \mu_AI_rate)}{\sigma(AI_rate)}$$

This variable represents the primary independent variable, capturing the intensity of AI-related language. For each firm i, we compute forward-looking stock returns at 1-year, 3-year, and 5-year horizons:

$$ret_{k_fwd,i} = \frac{(P_{i,t+k} - P_{i,t})}{P_{i,t}} \text{ for } k \in \{1, 3, 5\} \text{ years}$$

Similarly, control variables are constructed to capture past performance momentum:

$$ret_{k_past,i} = \frac{(P_{i,t} - P_{i,t-k})}{P_{i,t-k}} \text{ for } k \in \{1, 3, 5\} \text{ years}$$

To test the hypotheses, we estimate a cross-sectional linear regression for each time horizon:

$$ret_{k_\text{fwd},i} = \beta_0 + \beta_1 AI_rate_z_i \ + \beta_2 ret_{1y_past,i} + \beta_3 ret_{3y_past\ ,i} + \beta_4 ret_{5y_past,i} + \epsilon_i$$

The dependent variables are firms' forward-looking stock returns, denoted as $ret_{k_fwd,i}$. The primary independent variable is $AI_rate_z_i$, which captures the standardized intensity of AI-related disclosure in firm i's annual filing. To condition on prior performance, the models include historical stock returns, $ret_{k_past,i}$, as control variables. The error term, ϵ_i, is assumed to have a mean of zero.

A positive and significant β_2 would support H1, indicating that AI language functions as a priced strategic signal.

4. RESULTS

This section presents the empirical results of the regression models testing whether AI-related language in Item 7 of 10-K filings predicts future stock performance. Table 1 reports the estimates for the effect of standardized AI disclosure intensity on one-year, three-year, and five-year forward returns, controlling for prior performance over corresponding historical periods.

The findings provide empirical evidence that AI-related disclosure does not function as an effective strategic signal in this setting. Although signaling theory suggests that AI claims may convey information about technological capability (Loughran & McDonald, 2011, 2020; Spence, 1978), the absence of valuation response implies that investors do not interpret disclosure language alone as credible evidence of implementation or advantage. Put differently, markets may view AI discourse in 10-K filings as non-costly or easily imitated, limiting its informativeness as a signal.

Insert Table 1 about here

Prior studies consistently show that firms emphasizing AI adoption experience enhanced efficiency and profitability (Mishra et al., 2022), though such benefits depend on the technological readiness of their industry (Zeng et al., 2025). Effective AI integration requires complementary human, institutional, and social capabilities (Bianchini et al., 2025). More broadly, AI is emerging as a general-purpose technology driving a new wave of innovation and productivity growth (Damioli et al., 2025).

Across all three models, the coefficient for AI_rate_z is negative, small in magnitude, and statistically insignificant. This pattern holds for all three forward horizons. These findings provide

no empirical support for H1 or sub-hypotheses H1a–H1c, which predicted a positive association between AI disclosure intensity and future returns.

In contrast, several control variables show meaningful associations. Prior performance over longer horizons, particularly the five-year past return, is positively and significantly related to future performance in the five-year forward model. This indicates persistence or momentum effects in long-run valuation patterns, consistent with prior financial research.

The absence of significant effects suggests that capital markets do not treat AI-related language in Item 7 disclosures as a priced strategic signal. This implies that investors may discount narrative claims of AI orientation unless accompanied by additional, verifying information. These findings contribute to signaling theory by indicating that AI language in 10-K filings does not currently meet the criteria of a credible signal as theorized by Spence (1973).

5. DISCUSSION AND CONCLUSION

This study examined whether AI-related language in the Management Discussion and Analysis (MD&A) section of 10-K filings functions as a credible strategic signal in capital markets. Drawing on signaling theory (Spence, 1973), we theorized that AI discourse could convey information about firms' technological orientation that investors incorporate into future valuation. However, the empirical findings do not support this expectation: across one-, three-, and five-year horizons, the standardized AI disclosure variable does not predict forward stock returns, even after controlling for past performance. The coefficients on AI language intensity are small, negative, and statistically insignificant in all models. For managers, the findings highlight the limits of rhetorical positioning: merely increasing AI-related terminology in regulatory filings is unlikely to influence valuation. Firms seeking to communicate strategic transformation may need to pair linguistic signals with verifiable indicators to achieve credibility. Methodologically, this study

demonstrates the utility and limits of lexicon-based measurement in studying technology disclosure. However, the null findings also suggest that lexical proxies may under capture the qualitative context or semantic nuance required to differentiate symbolic from substantive technological communication. Future research may benefit from integrating lexicon-based approaches with semantic embeddings, sentence-level classification, or model-based validation using operational indicators.

Several limitations suggest avenues for future inquiry. First, the analysis focuses on a single reporting year and does not capture how temporal shifts in disclosure affect market interpretation; longitudinal designs could determine whether persistent AI discourse alters credibility. Second, our approach treats AI disclosure as homogeneous; industry-level variation in adoption pathways may moderate its signaling value. Third, integrating alternative data sources, patents, insider trading, earnings calls, or capital expenditures, may help distinguish between firms that are experimenting with AI and those that are commercializing AI.

In conclusion, the evidence suggests that AI-related language in 10-K filings does not serve as a priced strategic signal in capital markets. Although AI remains a central theme in strategic communication and corporate innovation narratives, investors appear to discount disclosure language without corroborating evidence of implementation.

ACKNOWLEDGMENT AND DISCLOSURE

The authors received no specific financial support for this research and report no financial relationships or other potential conflicts of interest. All disclosures are made in accordance with the National Bureau of Economic Research dissemination disclosure policy.

REFERENCES

Alduais, F. (2024). Textual analysis of the annual report and corporate performance: evidence from China. *Journal of Financial Reporting and Accounting*, *22*(5), 1221–1252.

Azmi Shabestari, M., & Romero, J. A. (2024). Textual analysis and future performance: Evidence from item 1a and item 7. *Journal of Accounting, Auditing & Finance*, *39*(3), 931–948.

Bellstam, G., Bhagat, S., & Cookson, J. A. (2021). A text-based analysis of corporate innovation. *Management Science*, *67*(7), 4004–4031.

Bianchini, S., Müller, M., & Pelletier, P. (2025). Drivers and barriers of AI adoption and use in scientific research. *Technological Forecasting and Social Change*, *220*. https://doi.org/10.1016/j.techfore.2025.124303

Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet Allocation. *Journal of Machine Learning Research*, *3*.

Bochkay, K., Brown, S. V, Leone, A. J., & Tucker, J. W. (2023). Textual analysis in accounting: What's next? *Contemporary Accounting Research*, *40*(2), 765–805.

Bodnaruk, A., Loughran, T., & McDonald, B. (2015). Using 10-K text to gauge financial constraints. *Journal of Financial and Quantitative Analysis*, *50*(4), 623–646.

Buehlmaier, M. M. M., & Whited, T. M. (2018). Are financial constraints priced? Evidence from textual analysis. *The Review of Financial Studies*, *31*(7), 2693–2728.

Chan, S. W. K., & Chong, M. W. C. (2017). Sentiment analysis in financial texts. *Decision Support Systems*, *94*, 53–64.

Cooper, H. B., Ewing, M. T., & Mishra, S. (2022). Text-mining 10-K (annual) reports: a guide for B2B marketing research. *Industrial Marketing Management*, *107*, 204–211.

Damioli, G., Van Roy, V., Vertesy, D., & Vivarelli, M. (2025). Is artificial intelligence leading to a new technological paradigm? *Structural Change and Economic Dynamics*, *72*, 347–359.

Du, K., Xing, F., Mao, R., & Cambria, E. (2024). Financial sentiment analysis: Techniques and applications. *ACM Computing Surveys*, *56*(9), 1–42.

Dyer, T., Lang, M., & Stice-Lawrence, L. (2017). The evolution of 10-K textual disclosure: Evidence from Latent Dirichlet Allocation. *Journal of Accounting and Economics*, *64*(2–3), 221–245.

Harris, T. (2024). Managers' perception of product market competition and earnings management: a textual analysis of firms' 10-K reports. *Journal of Accounting Literature*.

Hinings, B., Gegenhuber, T., & Greenwood, R. (2018). Digital innovation and transformation: An institutional perspective. *Information and Organization*, *28*(1), 52–61.

Humphreys, A., & Wang, R. J.-H. (2018). Automated text analysis for consumer research. *Journal of Consumer Research*, *44*(6), 1274–1306.

Karapandza, R. (2016). Stock returns and future tense language in 10-K reports. *Journal of Banking & Finance*, *71*, 50–61.

Kouloumpis, E., Wilson, T., & Moore, J. (2011). Twitter sentiment analysis: The good the bad and the omg! *Proceedings of the International AAAI Conference on Web and Social Media*, *5*(1), 538–541.

Kraus, S., Durst, S., Ferreira, J. J., Veiga, P., Kailer, N., & Weinmann, A. (2022). Digital transformation in business and management research: An overview of the current status quo. *International Journal of Information Management*, *63*, 102466.

Li, F. (2010). The information content of forward-looking statements in corporate filings—A naïve Bayesian machine learning approach. *Journal of Accounting Research*, *48*(5), 1049–1102.

Liu, M., Wu, K., Yang, R., & Yu, Y. (2020). Textual Analysis for Risk Profiles from 10-K Filings. *CPA Journal*, *90*(6).

Loughran, T., & McDonald, B. (2011). When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *The Journal of Finance*, *66*(1), 35–65.

Loughran, T., & McDonald, B. (2020). Textual analysis in finance. *Annual Review of Financial Economics*, *12*(1), 357–375.

Mao, Y., Liu, Q., & Zhang, Y. (2024). Sentiment analysis methods, applications, and challenges: A systematic literature review. *Journal of King Saud University-Computer and Information Sciences*, 102048.

Mishra, S., Ewing, M. T., & Cooper, H. B. (2022). Artificial intelligence focus and firm performance. *Journal of the Academy of Marketing Science*, *50*(6), 1176–1197.

Mushtaq, R., Gull, A. A., Shahab, Y., & Derouiche, I. (2022). Do financial performance indicators predict 10-K text sentiments? An application of artificial intelligence. *Research in International Business and Finance*, *61*, 101679.

Nambisan, S., Wright, M., & Feldman, M. (2019). The digital transformation of innovation and entrepreneurship: Progress, challenges and key themes. *Research Policy*, *48*(8), 103773.

Ravula, S. (2021). Text analysis in financial disclosures. *ArXiv Preprint ArXiv:2101.04480*.

Sebastian, I. M., Ross, J. W., Beath, C., Mocker, M., Moloney, K. G., & Fonstad, N. O. (2020). How big old companies navigate digital transformation. In *Strategic information management* (pp. 133–150). Routledge.

Spence, M. (1978). Job market signaling. In *Uncertainty in economics* (pp. 281–306). Elsevier.

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

Zeng, S., Wang, X., & Sun, T. (2025). Artificial Intelligence, Domain AI Readiness, and Firm Productivity. *ArXiv Preprint ArXiv:2508.09634*.

TABLE 1

Table 1. Regression Results: AI Disclosure and Future Stock Returns

	1 Year Forward Return	**3-Year Forward Return**	**5 Year Forward Return**
AI_rate_z	-0.008	-0.016	-0.011
	(0.014)	(0.026)	(0.058)
1 Year Past Return	-0.010	-0.191**	-0.269+
	(0.036)	(0.069)	(0.154)
3 Year Past Return	-0.010	-0.008	-0.117***
	(0.008)	(0.015)	(0.033)
5 Year Past Return	0.009	0.014	0.135***
	(0.009)	(0.016)	(0.036)
Num. Obs.	449	449	448
R^2	0.005	0.030	0.038
R^2 (Adj.)	-0.004	0.021	0.029

+ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$