 http://dx.doi.org/10.17287/kmr.2026.55.4.1589

Aspiring to Lead: Performance Feedback and AI Disclosure after the Generative AI Shock

전인재(주저자)
Injae Jeon(First Author)

고려대학교 기업경영연구원 연구위원 Institute for Business Research and Education, Korea University(*jeonsom@korea.ac.kr*)

The emergence of generative artificial intelligence (AI) has introduced substantial uncertainty about the basis of competitive advantage. Extending the behavioral theory of the firm beyond internal search and adaptation, this study examines how performance feedback shapes firms' external strategic communication during a major technological discontinuity. I argue that firms performing above social aspiration levels are more likely to increase AI-related disclosure following the generative AI shock. Favorable peer-relative performance provides a credible basis for connecting current success to claims of technological readiness, managerial foresight, and future competitiveness. I further predict that this response is stronger in industries where AI was already salient before the shock. Among firms that discuss AI, I also expect above-aspiration firms to place greater relative emphasis on AI's competitive implications. Using a panel of 10-K filings by U.S. public firms from 2017 to 2024 and an LLM-assisted text analysis approach, I find broad support for these predictions. This study extends the behavioral theory of the firm by showing that performance feedback shapes not only internal adaptation but also public communication under uncertainty. It also contributes to research on voluntary disclosure and technology framing by showing how firms use public discourse to interpret a major technological shock and position themselves in relation to it for external audiences.

Keyword: Generative AI, Corporate Disclosure, Performance Feedback, Social Aspirations, Competitive Framing, Technological Discontinuity

Ⅰ. Introduction

Generative artificial intelligence (AI) has quickly become one of the most consequential technological discontinuities facing modern firms (Brynjolfsson et al., 2021; Tushman and Anderson, 1986). In a short period, generative AI has moved from a specialized technical capability to a widely recognized general-purpose technology with broad implications for productivity, innovation, and competitive positioning (Yoon et al., 2025). Unlike earlier waves of computerization that

Submission Date: 03. 28. 2026 Revised Date: (1st: 05. 27. 2026) Accepted Date: 06. 11. 2026

primarily affected routine tasks, generative AI reaches into higher-order cognitive activities such as legal writing, financial analysis, and software development that were long viewed as the domain of human expertise (Acemoglu and Restrepo, 2022; Autor et al., 2003; Frey and Osborne, 2017). The public release of models such as ChatGPT therefore represents not only a technological advance but also an interpretive shock, disrupting established assumptions about value creation and raising uncertainty about which organizational capabilities will matter in the next competitive regime (Eisfeldt et al., 2023; Felten et al., 2021).

Such shocks affect not only what firms do, but also what they say. When a major discontinuity triggers an era of ferment, the meaning of the new technology and the basis of future competition are not yet settled (Kaplan, 2008; Kaplan and Tripsas, 2008). Under these conditions, corporate disclosure becomes an important vehicle through which firms interpret environmental change for external audiences. Through disclosure, firms can reduce information asymmetries, shape stakeholder beliefs, and present themselves as prepared for technological transition (Granqvist et al., 2013; Healy and Palepu, 2001; Leuz and Verrecchia, 2000; Lounsbury and Glynn, 2001). This is especially true in the case of generative AI, where many capabilities that may determine success, such as proprietary data, specialized human capital, and managerial foresight, are difficult for outsiders to observe directly (Brynjolfsson et al., 2021). In this setting, disclosure can serve as a visible form of external communication about a firm's perceived preparedness for technological change.

Firms, however, differ markedly in how they respond to this shock in their disclosures. Some firms sharply expand AI-related discussion and connect AI to future opportunities, risks, and competitive implications, whereas others remain cautious or largely silent. What explains this heterogeneity? I argue that an important part of the answer lies in performance feedback. Drawing on the behavioral theory of the firm, I suggest that firms interpret and communicate the generative AI shock through the lens of their performance relative to social aspiration levels.

Prior research on performance feedback has focused primarily on the consequences of performance below aspirations, emphasizing problemistic search, risk taking, and corrective change (Greve, 2008). In contrast, I focus on the strategic implications of performance above aspirations (Baum et al., 2005; Iyer and Miller, 2008). I argue that, following the generative AI shock, firms exceeding social aspiration levels are especially likely to increase AI-related disclosure. Favorable peer-relative performance gives managers a credible basis for presenting current success as evidence of managerial capability and future readiness (Audia and Greve, 2006; Baum et al., 2005; Kim et al., 2015). Rather than treating disclosure as a direct measure of AI adoption, I examine it as a form of public communication through which firms connect present performance to claims about preparedness for technological change.

I further argue that this logic extends beyond the amount of disclosure to its framing. During periods of technological upheaval, firms do not simply report on new technologies; they also participate in shaping how those technologies are interpreted by external audiences (Kaplan and Tripsas, 2008). Firms performing above social aspirations are more likely to frame AI in competitive terms by linking it to rivalry, market positioning, industry shifts, substitution threats, platform competition, or changes in the basis of advantage. Such framing allows managers to connect current performance to a forward-looking account of continued competitiveness. I also propose that this effect depends on technological salience. In industries where AI is more central to core tasks and value creation, the competitive meaning of AI is likely to be more visible and consequential.

I examine these arguments using a panel of U.S. public firms' 10-K filings from 2017 to 2024. I use an LLM-assisted text analysis approach to identify substantive AI-related discussion in 10-K filings and to classify the content of that discussion across multiple categories. I then relate variation in AI disclosure to firms' performance relative to social aspiration levels. The results show that firms exceeding social aspiration levels significantly increase AI-related disclosure following the generative AI shock. Among firm-years with AI disclosure, these firms also place greater relative emphasis on competitive AI framing. In addition, the positive disclosure response is stronger in industries where AI was more salient before the shock. Historical aspirations yield weaker and less consistent patterns. These findings suggest that, in the context of a field-wide technological shock, managers orient strongly toward social comparison when communicating their strategic position to external audiences.

This study contributes to research on the behavioral theory of the firm and voluntary disclosure by showing how performance feedback shapes external strategic communication under conditions of technological uncertainty. Prior work has focused primarily on how aspiration shortfalls trigger internal search and adaptation (Iyer and Miller, 2008; Kuusela et al., 2017). By contrast, I show that favorable peer-relative performance is associated with more extensive AI-related disclosure following a major technological discontinuity (Healy and Palepu, 2001; Verrecchia, 2001).

The study also extends this argument from disclosure intensity to strategic framing, showing that, conditional on discussing AI, firms performing above social aspirations place greater emphasis on the competitive implications of AI (Benner and Tripsas, 2012; Kaplan, 2008). More broadly, the findings contribute to research on the organizational consequences of generative AI by suggesting that, in the early stages of a general- purpose technology shift, competition unfolds not only through operational investments but also through firms' efforts to publicly interpret the technology and position themselves in relation to it (Eisfeldt et al., 2023; Felten et al., 2021).

Ⅱ. Literature Review and Hypothesis Development

2.1 Generative AI as a Technological Discontinuity and an Interpretive Shock

Generative AI has emerged as a consequential technological discontinuity that unsettles prevailing assumptions about how firms create value, organize production, and compete. Research on technological discontinuities argues that major breakthroughs destabilize established competitive positions by making the value of existing competencies, routines, and complementarities more uncertain (Tripsas, 1997; Tushman and Anderson, 1986). Such periods are often marked by an era of ferment, in which the trajectory of the new technology remains unsettled, and market participants differ in their interpretations of what the technology means for industry evolution (Kaplan and Tripsas, 2008). Generative AI exhibits these characteristics. It is widely regarded as a general-purpose technology whose effects are likely to diffuse across a broad range of tasks, products, and organizational functions rather than remain confined to a narrow technical domain (Brynjolfsson et al., 2021).

The significance of generative AI extends beyond automation. Unlike earlier digital technologies that primarily improved information processing or transactional efficiency, generative AI directly affects knowledge work, content production, product design, and the recombination of organizational knowledge. As a result, it has implications for both labor substitution and new task creation (Acemoglu and Restrepo, 2019, 2022), as well as for product innovation and firm growth (Babina et al., 2024; Eisfeldt et al., 2023). These features make generative AI more than just another technological input; they make it a technology that may alter the basis of strategic action. When such shifts occur, firms face not only technical uncertainty but also interpretive uncertainty, as managers, investors, and analysts must determine what the technology means and which firms are best positioned to benefit from it (Kaplan, 2008).

Under such interpretive uncertainty, corporate disclosure becomes especially important. Disclosure does more than transmit information; it also helps shape how external audiences understand ambiguous environmental change (Granqvist et al., 2013). This is particularly true when firms face pressure to explain how they understand and intend to respond to a major technological shift. In the case of generative AI, many of the capabilities that may matter most—such as proprietary data, specialized human capital, and organizational experimentation capacity—are difficult for outsiders to observe directly (Brynjolfsson et al., 2021). AI-related disclosure can therefore serve as both an informational and a symbolic mechanism through which firms position themselves in response to a highly salient technological shock (Granqvist et al., 2013; Lounsbury and Glynn, 2001).

2.2 Performance Feedback, Above-Aspiration Outcomes, and Strategic Action

To explain cross-firm heterogeneity in AI-related disclosure, I draw on the behavioral theory of the firm. A central insight of this tradition is that firms evaluate performance relative to aspiration levels, which serve as reference points for interpreting whether current outcomes are satisfactory (Cyert and March, 1963; Greve, 1998). Prior work has identified multiple sources of aspiration levels, including firms' own past performance and the performance of comparable peers. Importantly, the relative salience of these aspiration sources should depend on the decision context. Historical aspirations are likely to be especially relevant when firms evaluate current outcomes against their own prior trajectory and update expectations through accumulated experience. Social aspirations, by contrast, should become more salient when performance is interpreted in explicitly peer-relative terms and when organizational actions are visible to external audiences who also evaluate firms comparatively (Audia and Greve, 2006; Greve, 1998). Because the present study examines public disclosure following a field- wide technological shock, I focus on social aspirations as the primary reference point. In this setting, managers are not simply assessing whether the firm is doing better than before; they are also communicating whether the firm is keeping pace with, lagging behind, or leading comparable rivals.

Research on performance feedback has primarily emphasized the consequences of performance below aspirations, showing that shortfalls trigger problemistic search, risk taking, and corrective organizational change (Greve, 2003, 2008; Han and Park, 2014; Iyer and Miller, 2008; Kuusela et al., 2017). By contrast, I focus on the strategic implications of performance above aspirations. Firms that exceed aspiration levels are not simply inactive. Favorable attainment discrepancies can reinforce managerial confidence in the firm's strategic judgment and capabilities and provide a basis for interpreting current performance as evidence of effective strategic management (Baum et al., 2005; Kim et al., 2015). Above-aspiration performance may also reduce immediate performance pressure and ease resource constraints, but in the present setting, its most important implication is communicative: it gives firms a credible peer-relative position from which to interpret and present their response to technological change.

These implications become especially important during periods of technological upheaval. When a discontinuity such as generative AI unsettles the basis of competition, external audiences face uncertainty about which firms are prepared for the new technological regime. Firms performing above social aspirations can more plausibly connect current success to claims of managerial foresight, technological readiness, and future competitiveness. In this context, above-aspiration performance creates not only confidence inside the firm but also credibility outside the firm. It allows managers to frame AI-related disclosure as part of a broader narrative of continued competence rather than as speculative rhetoric or remedial impression management (Healy and Palepu, 2001; Lounsbury and Glynn, 2001; Singh, 1986).

2.3 AI Disclosure as a Signal of Managerial Capability Under Uncertainty

Voluntary disclosure research suggests that managers disclose information strategically when doing so can reduce information asymmetries, direct stakeholder attention, shape stakeholder beliefs, and influence market valuation (Ahn et al., 2007; Healy and Palepu, 2001; Li, 2010; Loughran and McDonald, 2011; Stocken, 2000). Disclosure becomes particularly valuable when outside stakeholders have limited ability to independently assess a firm's preparedness for technological change (Kaplan and Tripsas, 2008; Verrecchia, 2001). This problem is especially acute in the case of generative AI because many relevant capabilities—such as proprietary data assets, complementary human capital, and internal experimentation capacity—are intangible and difficult for outsiders to observe directly (Brynjolfsson et al., 2021).

In this setting, AI-related disclosure can serve as an especially salient signal of otherwise unobservable managerial capabilities because generative AI has become a visible marker of strategic modernity and adaptive capability (Felten et al., 2021; Granqvist et al., 2013). Such communication is particularly useful when stakeholders are uncertain about managerial quality or about the firm's private assessment of an emerging opportunity (Healy and Palepu, 2001; Stocken, 2000). Firms performing above social aspirations are especially likely to perceive such an opportunity. Their recent success provides a credible basis for claiming superior judgment, and AI-related disclosure allows them to link that success to a forward-looking narrative about preparedness for technological change. Rather than merely reporting isolated initiatives, these firms can present themselves as recognizing the significance of generative AI early and responding in ways that support future competitiveness (Lounsbury and Glynn, 2001; Suddaby and Greenwood, 2005).

Importantly, firms performing above social aspirations should be more likely to use such disclosure proactively than firms performing below aspirations. Firms below aspirations may face stronger pressure to address immediate performance problems, and forward-looking claims about AI readiness may be discounted by external audiences when current peer-relative performance is weak (Audia and Greve, 2006; Iyer and Miller, 2008). One alternative possibility is that below-aspiration firms might attempt to compensate by inflating AI-related claims. However, disclosure in SEC Form 10-K filings is not costless. Exaggerated or weakly substantiated claims in formal regulatory filings can impose reputational and legal costs, which should constrain aggressive signaling by poorly performing firms (Loughran and McDonald, 2011). By contrast, firms performing above social aspirations have a stronger credibility basis for linking AI-related disclosure to managerial capability and future competitiveness. Accordingly, firms performing above social aspirations should be better positioned to increase AI-related disclosure following the generative AI shock.

Hypothesis 1 (H1): Firms exceeding social aspiration levels will increase AI-related disclosure following the generative AI shock.

2.4 From Disclosure to Framing: Why Firms Above Social Aspirations Cast AI in Competitive Terms

Although the amount of AI-related disclosure is important, the way firms frame AI is equally consequential. Communication during periods of technological change is shaped not only by whether firms mention a new technology, but also by how they define its strategic meaning. Research on framing and managerial cognition shows that firms interpret emerging technologies through schemas that shape both strategic action and stakeholder understanding (Kaplan and Tripsas, 2008). In periods of rapid technological change, when the meaning and implications of a new technology remain unsettled, firms also compete to influence how that technology is understood and which responses are viewed as legitimate or strategically appropriate (Benner and Tripsas, 2012; Kaplan, 2008).

I argue that firms performing above social aspirations will be especially likely to frame AI in competitive terms. Competitive framing refers to presenting AI as relevant to rivalry, market positioning, or the basis of competitive advantage. This framing should be more likely among firms performing above social aspirations for three reasons. First, competitive framing allows these firms to present their response to technological uncertainty as strategic rather than reactive. Rather than describing AI as a generic technological development, they can portray it as a means of protecting or extending their competitive position. Doing so helps managers present the firm as actively responding to a changing environment rather than passively adapting to it (Kaplan, 2008; Kaplan and Tripsas, 2008).

Second, competitive framing allows firms above social aspirations to connect current success to future competitiveness. Prior performance gives these firms reputational capital, but a major technological shift can unsettle the criteria by which firms are evaluated. By embedding AI within a competitive narrative, managers can suggest that the capabilities underlying current success will remain valuable under the new technological regime (Lounsbury and Glynn, 2001; Suddaby and Greenwood, 2005). In this way, competitive framing helps preserve the value of current performance when external audiences are uncertain about how future competition will unfold.

Third, competitive framing is likely to be more credible when it comes from firms already performing well relative to peers. Because stakeholders can observe favorable current outcomes, they are more likely to view claims about AI-enabled competitiveness as plausible rather than rhetorical overstatement (Lounsbury and Glynn, 2001; Stocken, 2000). By contrast, similar claims from firms performing below aspirations may be discounted as impression management or overcompensation for weak performance. Accordingly, firms performing above social aspirations should have both stronger incentives and a higher expected return from framing AI in competitive terms.

This argument complements research on corporate rhetoric and symbolic management.

Firms use language not only to describe reality but also to shape how new categories, technologies, and strategic possibilities are interpreted by external audiences (Fiol, 1989; Granqvist et al., 2013). In the case of generative AI, competitive framing is particularly salient because AI has quickly become associated with innovation, adaptability, and strategic relevance. Firms performing above social aspirations should therefore be more likely not only to disclose AI but also to frame it in ways that emphasize its implications for rivalry, market positioning, industry change, and competitive advantage.

Hypothesis 2 (H2): Among firms that disclose AI-related information, firms exceeding social aspiration levels will devote a larger share of AI disclosure to competitive framing following the generative AI shock.

2.5 Technological Salience as a Boundary Condition

The benefits of AI disclosure and AI framing should not be uniform across industries. I argue that the relationship between performance above social aspirations and AI-related disclosure will be stronger when AI is more technologically salient within the focal industry. I define technological salience as the degree to which AI is perceived as consequential for value creation, product innovation, or competitive differentiation in an industry.

Research on general-purpose technologies suggests that broad technological advances diffuse unevenly because their practical relevance depends on complementary assets, task structures, and the extent to which existing activities can be reorganized around the new technology (Brynjolfsson et al., 2021). Realizing the value of such technologies often requires intangible investments, new routines, and managerial experimentation, all of which vary across industries. Related research on automation and task exposure likewise shows substantial cross-industry variation in the extent to which AI is likely to affect core activities rather than peripheral ones (Autor et al., 2003; Eisfeldt et al., 2023; Felten et al., 2021).

This heterogeneity matters because AI-related disclosure should have greater signaling value when stakeholders view AI as strategically central to the industry. In industries where AI is highly salient, investors, analysts, and other external audiences are more likely to treat AI readiness as informative about managerial capability and long-run competitiveness (Eisfeldt et al., 2023). Under these conditions, firms performing above social aspirations have more to gain from discussing AI because such disclosure speaks directly to an issue that external audiences already view as important (Granqvist et al., 2013).

Technological salience should also increase the scrutiny attached to AI-related claims. When AI is widely seen as central to industry competition, external audiences are likely to evaluate such claims more carefully. This heightened scrutiny should make AI-related disclosure more valuable for firms that can credibly link current performance to future preparedness, while limiting the attractiveness of aggressive signaling by firms whose claims may be viewed with greater skepticism. Accordingly, the positive relationship between performance above social aspirations and AI-related disclosure should be stronger in industries where technological salience of AI is high.

Hypothesis 3 (H3): The positive relationship between performance above social aspiration levels and AI-related disclosure will be stronger in industries with greater technological salience of AI.

Ⅲ. Methods

3.1 Data and Sample

I examine how firms respond to a major technological shock through their corporate disclosures. The empirical setting is the emergence of generative AI following the public release of ChatGPT on November 30, 2022, which sharply increased managerial and investor attention to AI-related opportunities and risks (Eisfeldt et al., 2023). My primary data source is annual Form 10-K filings of U.S. publicly traded firms. The filings are well-suited to this study because they provide standardized narrative disclosures about firms' operations, risks, and strategic positioning for external audiences (Dyer et al., 2017; Jegadeesh and Wu, 2013; Li, 2010; Loughran and McDonald, 2011).

I collect 10-K filings from the SEC EDGAR database and merge them with financial data from Compustat. The panel covers fiscal years 2017 through 2024, allowing me to observe both the pre-shock baseline and firms' disclosure responses after the emergence of generative AI. I define the post-shock period as fiscal years 2023 and 2024. Because ChatGPT was released near the end of calendar year 2022, fiscal year 2022 filings may vary in the extent to which they reflect the shock, depending on firms' fiscal-year ends and filing dates. I therefore treat fiscal year 2022 as a transition year and do not classify it as part of the post-shock period. This conservative definition reduces ambiguity associated with filings submitted shortly after the release of ChatGPT and allows the post-shock period to capture disclosures after firms had more time to interpret the generative AI shock. The final sample consists of 13,527 firm-year observations from 2,898 firms across a broad range of industries.

3.2 Variables

3.2.1 AI Disclosure

To construct the dependent variables, I use a multi-stage text analysis procedure. Consistent with recent work encouraging the use of statistical learning methods in financial text analysis, I employ a large language model (LLM)-assisted procedure to identify and classify AI-related disclosure (Harrison et al., 2023; Li et al., 2021; Miric et al., 2023). I first partition each filing into text windows of approximately 800 words. This approach allows me to identify localized AI-related discussion while preserving enough context for reliable interpretation, which is important because 10-K filings are lengthy, heterogeneous, and often contain substantial boilerplate (Dyer et al., 2017).

I then apply a keyword gate to identify candidate windows using terms such as “artificial intelligence,” “machine learning,” and “generative AI.” Candidate windows are subsequently evaluated using a two-stage LLM procedure. In the first stage, a screening model determines whether the passage contains substantive AI-related discussion. In the second stage, a more capable model assigns continuous scores ranging from 0 to 1 to five categories of AI discourse: AI use, AI opportunity, competitive framing, operational risk, and regulation. This design is consistent with emerging work showing that LLMs can be used effectively for annotation and categorization tasks when combined with validation and careful prompt design (Carlson and Burbano, 2025; Eisfeldt et al., 2023; Harrison et al., 2023). In addition, this multidimensional approach is preferable to simple word counts because prior research shows that textual measures are sensitive to contextual ambiguity, weighting choices, and topic heterogeneity, and that richer text representations often better capture economically meaningful variation (Jegadeesh and Wu, 2013; Li et al., 2021).

My primary dependent variable is AI disclosure intensity, measured as the log-transformed sum of category scores across all windows in a filing. To test Hypothesis 2, I also construct a measure of competitive AI framing share, defined as the competitive AI framing score divided by the total AI category score in a filing. It captures, within AI-related disclosure, the relative emphasis on competitive framing that links AI to the competitive environment, including rivals, industry shifts, substitution threats, platform competition, market positioning, or competitive advantage. To test Hypothesis 3, I construct a disclosure-based measure of industry AI salience at the three-digit SIC industry level. I define industry AI salience as the extent to which AI was already salient in an industry's public disclosure environment prior to the generative AI shock. Specifically, for each three- digit SIC industry, I calculate the average pre-shock AI disclosure score among firms in that industry using fiscal years 2017 through 2022. I then take the log- transformed value and standardize the measure to have a mean of zero and standard deviation of one in the construction sample. Higher values indicate industries in which AI was more salient in firms' disclosures before the generative AI shock. This measure is conceptually motivated by prior work showing that AI exposure varies substantially across occupations and industries (Felten et al., 2021). Because the final regression sample differs slightly from the construction sample, the standardized measure has a mean close to, but not exactly equal to, zero in the estimation sample reported in Table 1.

3.2.2 Performance Feedback

The main independent variables are derived from the behavioral theory of the firm and capture firm performance relative to aspiration levels. In line with the theoretical focus of this study, I emphasize social aspirations. I define social aspiration as the leave-one-out average ROA of peer firms in the same three-digit SIC industry-year, excluding the

〈Table 1〉 Descriptive statistics and correlations

Variable	Mean	S.D.	1	2	3	4	5	6	7	8	9	10	11	12	13
1 AI Disclosure Intensity	0.28	0.63	1.00												
2 Above Social Aspiration	0.29	0.40	0.06	1.00											
3 Below Social Aspiration	0.26	1.26	-0.03	-0.15	1.00										
4 Above Historical Aspiration	0.18	0.93	-0.03	-0.01	0.16	1.00									
5 Below Historical Aspiration	0.17	0.79	-0.01	-0.12	0.77	-0.04	1.00								
6 Firm Size	6.15	2.85	0.09	0.11	-0.43	-0.27	-0.33	1.00							
7 Profitability	-0.37	1.37	0.02	0.13	-0.61	-0.26	-0.45	0.55	1.00						
8 Leverage	0.84	1.97	-0.03	-0.10	0.61	0.23	0.43	-0.39	-0.73	1.00					
9 Cash Holdings	0.28	0.29	0.03	0.03	0.12	0.12	0.11	-0.42	-0.19	0.01	1.00				
10 Intangible Asset Intensity	0.19	0.23	0.14	0.07	-0.09	-0.05	-0.08	0.30	0.13	-0.07	-0.42	1.00			
11 Tobin's Q	0.97	0.78	0.07	0.02	0.43	0.23	0.28	-0.40	-0.47	0.44	0.30	-0.12	1.00		
12 Filing Length	152.59	57.74	0.10	0.06	-0.13	-0.08	-0.08	0.35	0.09	-0.13	0.09	0.04	-0.21	1.00	
13 Industry AI Salience	0.04	1.00	0.42	0.06	0.01	0.02	0.02	-0.07	-0.01	0.01	0.09	0.26	0.14	-0.05	1.00

focal firm from the calculation (Audia and Greve, 2006). I then compute attainment discrepancy as the difference between focal-firm ROA and this social aspiration level.

Consistent with prior performance-feedback research, I use a spline specification to decompose attainment discrepancy into above-aspiration and below-aspiration components. This piecewise approach allows for asymmetric responses to favorable and unfavorable performance feedback (Iyer and Miller, 2008; Kuusela et al., 2017). The variable of primary theoretical interest is above social aspiration, which captures the extent to which a firm outperforms its peer-based reference point. This measure maps directly onto my argument that favorable peer-relative performance provides a credible basis for AI-related disclosure and strengthens firms' incentives to communicate preparedness for technological change. I also construct historical aspiration using an exponentially weighted moving average of the focal firm's past ROA, following prior research on adaptive aspiration updating (Kim et al., 2015).

3.2.3 Control Variables

I include a standard set of lagged control variables that may influence disclosure behavior and the strategic tone of firms' communications. These controls include firm size (log total assets), profitability (ROA), leverage, cash holdings, Tobin's Q, and intangible asset intensity. Tobin's Q accounts for differences in growth opportunities and investment opportunities that may shape forward-looking disclosure (Li, 2010). Intangible asset intensity is included because AI-related capabilities often depend on data, knowledge, and organizational capital.

I also control for filing length, measured as the total number of text windows in the filing, because longer filings mechanically create more opportunities to mention emerging technologies (Dyer et al., 2017; Loughran and McDonald, 2011). This helps account for document-level heterogeneity in the textual measures.

3.3 Empirical Specification

To test the hypotheses, I examine whether firms performing above social aspirations increase AI-related disclosure after the emergence of generative AI following the public release of ChatGPT in late 2022. I estimate firm-year panel regressions of the following form:

$$\begin{aligned} AI\ Disclosure_{it} = \\ &\beta_1(Above\ Social\ Aspiration_{it-1} \times Post_t) \\ &+ \beta_2(Below\ Social\ Aspiration_{it-1} \times Post_t) \\ &+ \gamma Controls_{it-1} + \alpha_i + \delta_t + \epsilon_{it} \end{aligned}$$

where α_i denotes firm fixed effects to control unobserved, time-invariant firm and industry characteristics, and δ_t denotes year fixed effects to absorb macroeconomic trends and baseline temporal shifts in disclosure practices. The post-shock indicator equals one for fiscal years 2023 and 2024 and zero otherwise. Fiscal year 2022 is treated as a transition year and is not coded as part of the post-shock period because ChatGPT was released on November 30, 2022, and firms differ in fiscal-year ends and filing dates. Standard errors are clustered at the firm level to address potential serial correlation within firms over the sample period.

The coefficient of primary interest is the interaction between Above Social Aspiration and Post. For Hypothesis 1, the dependent variable is AI disclosure intensity. For Hypothesis 2, I replace the dependent variable with competitive AI framing share. For Hypothesis 3, I introduce a triple interaction involving industry AI salience to test whether the post-shock effect of performance above social aspirations is stronger in industries where AI is more strategically salient. Importantly, because industry AI salience is measured as a time-invariant, pre-shock industry characteristic, its main effect is absorbed by the firm fixed effects. Consequently, the model relies on the interaction terms to identify how the post-shock effect of performance above social aspirations varies across different levels of industry salience. Because the public release of ChatGPT affected the broader informational environment rather than any one firm in isolation, this design is intended to capture heterogeneous disclosure responses to a common field-level shock. The estimates, therefore, identify whether firms with different performance positions relative to social aspirations responded differently in their disclosures after the shock.

3.4 Measure Validation

Because the empirical design relies on model-assisted text classification, I conduct a manual audit to assess the validity of the AI disclosure measures. The validation focuses on three questions: whether the keyword gate omits substantively AI-related passages, whether the first-stage classifier correctly identifies substantive AI discussion, and whether the second-stage category assignments align with human judgment.

The audit yields three main conclusions. First, the keyword gate appears highly specific. In a random sample of excluded windows, virtually none were coded by human evaluators as containing substantive AI-related discussion, suggesting that the keyword screen does not generate substantial false exclusions. Second, the first-stage classifier appears conservative. Relative to human coding, it achieves very high precision, approximately .97 to .98, but somewhat lower recall, approximately .81 to .85. This indicates that the principal source of measurement error is missed positive cases rather than false positives, suggesting that any resulting bias is more likely to attenuate than inflate the estimated relationships.

Third, I assess the validity of the second-stage category scores using a multi-label category-presence framework. Because the five categories are not mutually exclusive, a disclosure window can contain more than one AI-related frame. I therefore evaluate whether each category identified by the model is substantively present in the text according to human coders. Under this criterion, the model performs well across all five dimensions. Precision is consistently high, averaging approximately .97 across categories, while recall averages approximately .76. Performance is similarly strong for the competitive AI framing category, which is central to Hypothesis 2, with precision of approximately .97 and recall of approximately .77. These results suggest that the model reliably identifies the presence of competitive AI framing when it is detected, although some relevant passages may be missed. This pattern again points toward conservative measurement error rather than systematic over-detection.

The audit also indicates that some residual noise is concentrated in executive biographies, board descriptions, and other boilerplate passages that mention AI without reflecting substantive firm-level disclosure. To address this concern, I instruct the classifier to prioritize substantive discussion over incidental mentions and, in additional analyses, employ a stricter disclosure measure that excludes biography-like and boilerplate-like windows. The results are substantively similar. Overall, these findings suggest that the measurement approach captures meaningful variation in firm-level AI discourse and is unlikely to generate the main findings through systematic over-detection.

Ⅳ. Results

4.1 Main Results

Table 1 reports descriptive statistics and pairwise correlations for the main variables. The correlations are generally modest and do not indicate serious multicollinearity concerns.

Table 2 presents the main results for AI Disclosure Intensity. Model 1 reports the baseline estimates. Firms performing above social aspirations exhibit significantly higher AI disclosure intensity. By contrast, the coefficients for the historical aspiration variables are small and statistically insignificant.

〈Table 2〉 AI disclosure intensity

DV: AI Disclosure Intensity	Model 1	Model 2
Above Social Aspiration × Post		0.069*** (0.026)
Below Social Aspiration × Post		-0.042*** (0.012)
Above Historical Aspiration × Post		-0.018 (0.014)
Below Historical Aspiration × Post		0.021 (0.019)
Above Social Aspiration	0.058*** (0.019)	0.031* (0.018)
Below Social Aspiration	0.005 (0.005)	0.009** (0.004)
Above Historical Aspiration	0.005 (0.004)	0.005 (0.004)
Below Historical Aspiration	-0.003 (0.006)	0.002 (0.005)
Firm Size	-0.001 (0.010)	0.003 (0.010)
Profitability	-0.010 (0.007)	-0.009 (0.007)
Leverage	-0.010** (0.004)	-0.003 (0.004)
Cash Holdings	0.006 (0.038)	0.018 (0.038)
Intangible Asset Intensity	0.061 (0.063)	0.052 (0.063)
Tobin's Q	0.019* (0.012)	0.017 (0.012)
Filing Length	0.805*** (0.193)	0.797*** (0.194)
Firm fixed effects	Y	Y
Year fixed effects	Y	Y
Number of observations	13,527	13,527
R^2	0.761	0.762

Note: Standard errors clustered at the firm level are reported in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$

This pattern is consistent with the argument that peer-relative standing is particularly relevant in this setting, where firms communicate their position to external audiences in the wake of a field-wide technological shock.

Model 2 provides the main test of Hypothesis 1. Consistent with the hypothesis, the interaction between Above Social Aspiration and Post is positive and statistically significant ($\beta = 0.069$, $p < 0.01$). Following the emergence of generative AI, firms whose prior performance exceeded social aspiration levels increased AI-related disclosure more than otherwise similar firms. The interaction between Below Social Aspiration and Post is negative and statistically significant ($\beta = -0.042$, $p < 0.01$), indicating that firms performing below social aspirations were less likely to expand AI-related disclosure after the shock. These results suggest that peer-relative performance shaped firms' disclosure responses to the emergence of generative AI.

Figure 1 provides event-study evidence for AI Disclosure Intensity. I use fiscal year 2021 as the reference year because it is the last full fiscal year before the public release of ChatGPT on November 30, 2022. Fiscal year 2022 is treated as a transition year,

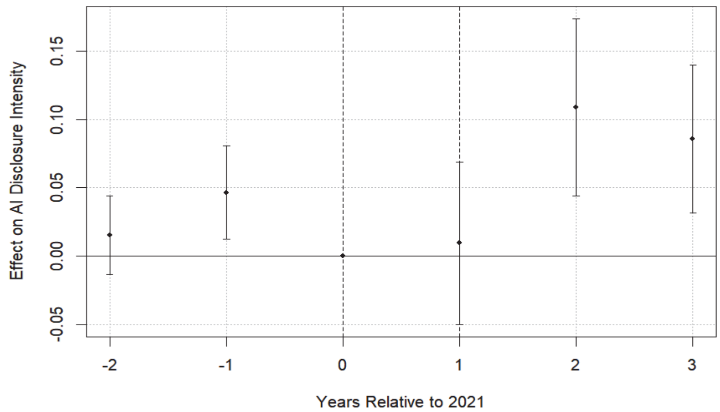

〈Figure 1〉 Event study: AI disclosure intensity with 95% confidence intervals

and fiscal years 2023 and 2024 represent the post-shock period. Thus, in the event-study figure, t = 0 refers to fiscal year 2021. The estimates indicate that the relationship between above-social-aspiration performance and AI disclosure becomes substantially larger after the generative AI shock. While some pre-shock estimates are positive, the post-shock coefficients are larger and more consistently positive, especially in 2023 and 2024. This pattern suggests that above-aspiration firms expanded AI-related disclosure more strongly after the shock, rather than that the estimates merely reflect a stable cross-sectional tendency to discuss AI.

Table 3 examines whether, among firms that discussed AI, above-aspiration firms also changed how they framed it. The dependent variable is Competitive AI Framing Share. Because this measure captures the relative emphasis on competitive framing within AI-related disclosure, it is defined only for firm-years with positive AI disclosure. Accordingly, the analysis is restricted to AI-disclosing firm-years, and the results should be interpreted as evidence about the framing of AI disclosure among firms that discuss AI rather than as evidence about the initial decision to disclose AI.

Consistent with Hypothesis 2, the interaction between Above Social Aspiration and Post is positive and statistically significant ($\beta = 0.040$, $p < 0.10$). Among AI-disclosing firm-years, firms exceeding social aspiration levels devote a larger share of their AI-related disclosure to competitive framing after the generative AI shock. This result suggests that, conditional on discussing AI, firms performing above social aspirations placed greater relative emphasis on the competitive implications of AI.

Figure 2 presents analogous event-study

〈Table 3〉 Competitive AI framing share

DV: Competitive AI Framing Share	Model 1
Above Social Aspiration × Post	0.040* (0.021)
Below Social Aspiration × Post	0.003 (0.007)
Above Historical Aspiration × Post	-0.001 (0.006)
Below Historical Aspiration × Post	0.012 (0.010)
Above Social Aspiration	0.014 (0.020)
Below Social Aspiration	0.002 (0.005)
Above Historical Aspiration	0.003 (0.002)
Below Historical Aspiration	-0.014** (0.007)
Firm Size	-0.010 (0.008)
Profitability	-0.001 (0.004)
Leverage	-0.004 (0.003)
Cash Holdings	-0.033 (0.025)
Intangible Asset Intensity	0.024 (0.039)
Tobin's Q	0.002 (0.008)
Filing Length	-0.075 (0.102)
Firm fixed effects	Y
Year fixed effects	Y
Number of observations	3,287
R^2	0.709

Note: Standard errors clustered at the firm level are reported in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$

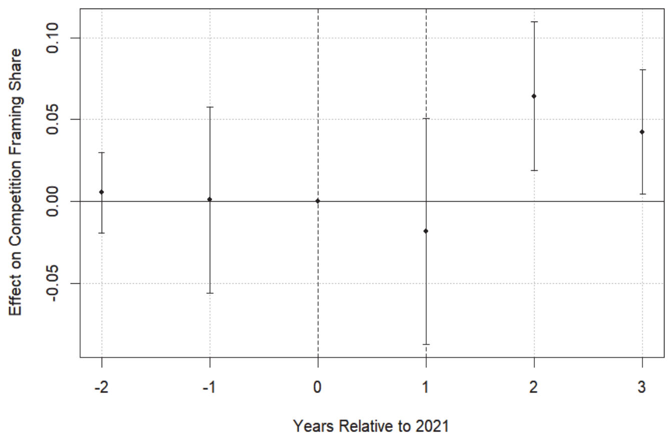

〈Figure 2〉 Event study: Competitive framing share with 95% confidence intervals

estimates for Competitive AI Framing Share. As in Figure 1, fiscal year 2021 serves as the reference year, fiscal year 2022 is treated as a transition year, and fiscal years 2023 and 2024 are interpreted as the post-shock period. The analysis is restricted to firm-years with positive AI disclosure because Competitive AI Framing Share is defined only for AI-disclosing firm-years. The pre-shock estimates are close to zero and imprecisely estimated, whereas the post-shock estimates become positive. This pattern suggests that, after the generative AI shock, firms performing above social aspirations placed greater relative emphasis on the competitive implications of AI among firm-years in which they discussed AI. Because the framing estimates are less precisely estimated than the disclosure-intensity results, I interpret this evidence as supportive rather than definitive.

Table 4 tests Hypothesis 3 by examining whether the disclosure response varies with Industry AI Salience. The triple interaction among Above Social Aspiration, Post, and Industry AI Salience is positive and statistically significant (β = 0.087, $p < 0.05$). This result indicates that the positive post-shock relationship between performance above social aspirations and AI disclosure is stronger in industries where AI was more salient in firms' public disclosures before the gen-

〈Table 4〉 Industry AI salience

DV: AI Disclosure Intensity	Model 1
Above Social Aspiration × Industry AI Salience × Post	0.087** (0.040)
Below Social Aspiration × Industry AI Salience × Post	-0.008 (0.008)
Above Social Aspiration × Industry AI Salience	-0.069*** (0.026)
Below Social Aspiration × Industry AI Salience	-0.003 (0.005)
Post × Industry AI Salience	0.117*** (0.018)
Above Social Aspiration × Post	0.062** (0.028)
Below Social Aspiration × Post	-0.030*** (0.007)
Above Social Aspiration	0.012 (0.020)
Below Social Aspiration	0.010** (0.004)
Above Historical Aspiration	0.005 (0.003)
Below Historical Aspiration	0.007 (0.006)
Firm Size	0.003 (0.010)
Profitability	-0.010 (0.008)
Leverage	-0.004 (0.005)
Cash Holdings	0.034 (0.038)
Intangible Asset Intensity	0.048 (0.061)
Tobin's Q	0.024** (0.012)
Filing Length	0.803*** (0.188)
Firm fixed effects	Y
Year fixed effects	Y
Number of observations	13,527
R^2	0.774

Note: Standard errors clustered at the firm level are reported in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$

erative AI shock. This pattern is consistent with the proposed mechanism: AI-related disclosure has greater signaling value when AI is already more central to the industry's competitive and informational environment.

4.2 Robustness Tests

The main analyses show that firms performing above social aspirations increased AI-related disclosure after the generative AI shock. I conduct several additional analyses to assess the robustness and interpretation of this finding. These analyses address four related concerns. First, I examine whether the result depends on the particular construction of the AI disclosure measure. Second, I assess whether the observed disclosure response reflects actual AI-use reporting or post-shock changes among intangible-intensive firms, rather than broader AI-related communication. Third, I examine whether the evidence is more consistent with a slack-resource explanation or with a peer-relative signaling interpretation. Fourth, I evaluate whether the result is driven by pre-existing AI orientation, observable firm differences, industry-specific trends, or alternative timing assumptions. Tables 5 through 8 report these analyses.

Table 5 first examines whether the main result is robust to alternative measures of AI-related disclosure. The positive interaction between Above Social Aspiration and Post remains statistically significant across multiple disclosure measures, including overall AI Disclosure Intensity (β = 0.069, $p < 0.01$), AI Window Share (β = 0.002, $p < 0.01$), Length-Adjusted AI Intensity (β = 0.001, $p < 0.10$), and Competitive AI Framing

〈Table 5〉 Multiple AI disclosure measures

	Model 1	Model 2	Model 3	Model 4	Model 5
	DV: AI Disclosure Intensity	DV: AI Window Share	DV: Length-Adjusted AI Intensity	DV: Competitive AI Framing Intensity	DV: Average AI Intensity | AI 〉0
Above Social Aspiration × Post	0.069*** (0.026)	0.002*** (0.0007)	0.001* (0.0007)	0.019** (0.010)	-0.005 (0.030)
Below Social Aspiration × Post	-0.042*** (0.012)	-0.001** (0.0005)	-0.001** (0.0005)	-0.005 (0.004)	-0.026** (0.013)
Firm controls	Y	Y	Y	Y	Y
Firm fixed effects	Y	Y	Y	Y	Y
Year fixed effects	Y	Y	Y	Y	Y
Number of observations	13,527	13,527	13,527	13,527	3,308
R^2	0.763	0.771	0.770	0.653	0.746

Note: Standard errors clustered at the firm level are reported in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$

Intensity (β = 0.019, $p < 0.05$). Unlike the framing-share measure used in Table 3, Competitive AI Framing Intensity captures the absolute amount of competitive AI framing and is estimated in the full sample. This provides complementary evidence that the competitive-framing pattern is not limited to the sample restricted to AI-disclosing firm-years in Table 3. By contrast, the corresponding coefficient for Average AI Intensity | AI > 0 is small and statistically insignificant. This pattern suggests that firms above social aspirations responded to the generative AI shock primarily by increasing the breadth and prominence of AI-related discussion across the filing, rather than by merely intensifying language within a fixed set of existing AI passages.

A related interpretive concern is that AI-related disclosure may reflect actual AI use or investment rather than external communication. Table 6 addresses this concern by examining whether the main result is concentrated in implementation-related AI disclosure or in firms with greater intangible asset intensity. Model 1 uses AI Use Disclosure as the dependent variable. This measure captures passages referring to AI use, implementation, application, or deployment and therefore most closely approximates disclosure of actual AI-related activity. If the main result were driven primarily by bet-

〈Table 6〉 AI use disclosure and intangible-asset-based alternative explanations

	Model 1	Model 2	Model 3
	DV: AI Use Disclosure	Intangible Asset Intensity	Intangible Asset Intensity × Above Social Aspiration
Above Social Aspiration × Post	0.021 (0.015)	0.053** (0.025)	0.052** (0.026)
Below Social Aspiration × Post	-0.021*** (0.007)	-0.038*** (0.012)	-0.043*** (0.014)
Post × Intangible Asset Intensity		0.094*** (0.011)	0.099*** (0.013)
Above Social Aspiration × Post × Intangible Asset Intensity			-0.008 (0.028)
Below Social Aspiration × Post × Intangible Asset Intensity			-0.011 (0.011)
Firm controls	Y	Y	Y
Firm fixed effects	Y	Y	Y
Year fixed effects	Y	Y	Y
Number of observations	13,527	13,527	13,527
R^2	0.797	0.767	0.767

Note: Standard errors clustered at the firm level are reported in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$

ter-performing firms adopting AI and therefore having more implementation activity to report, the post-shock effect of above-social-aspiration performance should be especially pronounced for this measure. However, the coefficient on Above Social Aspiration × Post is positive but statistically insignificant.

Models 2 and 3 examine whether the result is explained by post-shock disclosure changes among intangible-intensive firms. In Model 2, I add the interaction between Intangible Asset Intensity and Post. This interaction is positive and statistically significant ($\beta = 0.094$, $p < 0.01$), indicating that intangible-intensive firms increased AI disclosure more after the shock. Importantly, however, the coefficient on Above Social Aspiration × Post remains positive and statistically significant ($\beta = 0.053$, $p < 0.05$). Model 3 further adds the triple interaction among Above Social Aspiration, Post, and Intangible Asset Intensity. The triple interaction is statistically insignificant, suggesting that the above-aspiration effect does not occur only among intangible-intensive firms. These results suggest that the main disclosure pattern is unlikely to be explained solely by implementation-related AI reporting or by post-shock disclosure increases among intangible-intensive firms.

I also examine whether the main result is better explained by slack resources. If the disclosure response were primarily driven by resource capacity, firms with greater Cash Holdings should increase AI disclosure more after the shock, and the above-aspiration effect should be stronger among cash-rich firms. Table 7 reports these tests. Model 1 adds the interaction between Cash Holdings and Post. Model 2 adds the triple interaction among Above Social Aspiration, Post, and Cash Holdings. The results do not provide strong support for a slack-based explanation. Cash Holdings × Post is statistically insignificant, and the triple interaction between Above Social Aspiration, Post, and Cash Holdings is also statistically insignificant. The coefficient on Above Social Aspiration × Post remains positive and statistically significant across these specifications. These results are more consistent with the interpretation that favorable peer-relative performance shapes AI-related disclosure through credibility and external communication incentives rather than through slack resources alone. This interpretation is also consistent with the intangible-asset analyses reported in Table 6. Although intangible-intensive firms increased AI disclosure more after the shock, the above-aspiration effect was not concentrated among such firms. Thus, the evidence does not suggest that the main effect is driven solely by firms with greater observable resources or intangible assets.

〈Table 7〉 Slack resource tests

	Model 1	Model 2
DV: AI Disclosure Intensity	Slack Resource	Slack Resource × Above Social Aspiration
Above Social Aspiration × Post	0.072*** (0.026)	0.068*** (0.026)
Below Social Aspiration × Post	-0.042*** (0.012)	-0.045*** (0.013)
Post × Cash Holdings	-0.016 (0.010)	-0.009 (0.012)
Above Social Aspiration × Post × Cash Holdings		-0.023 (0.023)
Below Social Aspiration × Post × Cash Holdings		0.002 (0.004)
Firm controls	Y	Y
Firm fixed effects	Y	Y
Year fixed effects	Y	Y
Number of observations	13,527	13,527
R^2	0.763	0.763

Note: Standard errors clustered at the firm level are reported in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$

Table 8 addresses a broader endogeneity concern: firms performing above social aspirations may possess stronger pre-existing organizational capabilities that influence both performance and AI-related disclosure. Model 1 controls for Pre-shock AI Orientation by interacting Post with firms' average pre-shock AI disclosure intensity. The coefficient on Above Social Aspiration × Post remains positive and statistically significant (β = 0.069, $p < 0.01$), while the interaction between Post and Pre-shock AI Orientation is not statistically significant. Model 2 excludes firms with any pre-shock AI disclosure. The coefficient remains positive and statistically significant (β = 0.046, $p < 0.05$), suggesting that the main result is not driven solely by firms that were already visibly AI-oriented before the generative AI shock.

Model 3 reports an entropy-balanced difference-in-differences specification. In this analysis, treatment is defined as being above the social aspiration level before the post-shock period, and the comparison group is weighted to resemble the treated group on observable pre-treatment characteristics, including firm size, ROA, leverage, cash holdings, intangible asset intensity, Tobin's Q, filing length, pre-shock AI orientation, and pre-shock AI window share. The weighted estimate remains positive and statistically significant (β = 0.053, $p < 0.05$), indicating that the result is not simply due to observable pre-treatment differences.

Models 4 and 5 provide additional checks on industry-level trends and temporal ordering.

〈Table 8〉 Endogeneity and alternative explanations

	Model 1	Model 2	Model 3	Model 4	Model 5
DV: AI Disclosure Intensity	Pre-Shock AI Orientation	Exclude Pre-Shock AI Firms	Entropy-Balanced DiD	Industry-Year FE	Two-Year Lag
Above Social Aspiration × Post	0.069*** (0.026)	0.046** (0.021)	0.053** (0.026)	0.101*** (0.039)	0.069** (0.032)
Below Social Aspiration × Post	-0.042*** (0.012)	-0.027*** (0.010)		-0.031** (0.013)	-0.034*** (0.011)
Post × Pre-shock AI Orientation	0.006 (0.015)				0.005 (0.017)
Firm controls	Y	Y	Y	Y	Y
Firm fixed effects	Y	Y	Y	Y	Y
Year fixed effects	Y	Y	Y	Y	Y
Number of observations	13,527	10,372	13,206	13,527	10,645
R^2	0.763	0.505	0.778	0.787	0.773

Note: Standard errors clustered at the firm level are reported in parentheses. * *p* 〈 .10, ** *p* 〈 .05, *** *p* 〈 .01

Model 4 replaces year fixed effects with industry-year fixed effects, thereby comparing firms within the same three-digit SIC industry and fiscal year. The coefficient on Above Social Aspiration × Post remains positive and statistically significant (β = 0.101, p 〈 0.01). Model 5 uses two-year lagged performance-feedback measures, further separating the measurement of aspiration- based performance from subsequent AI disclosure. The coefficient remains positive and statistically significant (β = 0.069, p 〈 0.05).

I also examine whether the main result is driven by industries in which AI was already highly salient before the generative AI shock. In the sample restricted to industries below the median of pre-shock Industry AI Salience, the coefficient on Above Social Aspiration × Post remains positive and statistically significant (β = 0.059, p 〈 0.05). When I further restrict the sample to the bottom tercile of Industry AI Salience, the coefficient remains positive but is less precisely estimated (β = 0.038, n.s.). This pattern suggests that the main result is not confined to high-AI-salience industries, although the disclosure response is weaker where AI was least salient before the shock. Overall, the results in Table 8 and the low-salience subsample analyses suggest that the main disclosure finding is unlikely to be explained solely by pre-existing AI orientation, observable firm differences, industry-specific disclosure trends, or short- term simultaneity between performance and disclosure.

Ⅴ. Discussion and Conclusion

5.1 Theoretical and Managerial Implications

This study examines how firms respond to the emergence of generative AI through corporate disclosure. Drawing on the behavioral theory of the firm, I find that firms performing above social aspiration levels increased AI-related disclosure following the generative AI shock. Among firms with AI disclosure, they also placed greater relative emphasis on competitive AI framing. In addition, the disclosure response was stronger in industries where AI was more salient before the shock.

These findings deepen our understanding of the relationship among performance feedback, voluntary disclosure, and the communication of technological change. The behavioral theory of the firm has traditionally emphasized how aspiration levels shape internal responses such as problemistic search and organizational adjustment (Greve, 2008; Kuusela et al., 2017). The evidence here suggests that performance feedback also has an externally oriented dimension. In the context of a field-wide technological shock, favorable peer-relative performance appears to provide firms with a stronger basis for communicating preparedness for technological change to external audiences (Baum et al., 2005; Lounsbury and Glynn, 2001). This does not imply that disclosure is purely symbolic or detached from firms' underlying capabilities. Rather, it suggests that performance relative to peers shapes how firms publicly interpret and communicate a highly uncertain technological shift.

The study also contributes to voluntary disclosure research by showing that disclosure following a technological shock is shaped not only by information asymmetries or capital-market incentives, but also by firms' behavioral reference points. Under conditions of interpretive uncertainty, where many of the capabilities relevant to a new technology are difficult for outsiders to observe directly (Brynjolfsson et al., 2021), disclosure becomes an important means through which firms communicate their readiness and strategic position (Healy and Palepu, 2001). In this sense, the findings suggest that firms performing above social aspirations are better positioned to connect current performance to forward-looking claims about technological readiness and future competitiveness.

In addition, by moving beyond disclosure volume to examine the framing of AI-related communication, the study contributes to research on how organizations interpret and present technological discontinuities (Kaplan, 2008; Suddaby and Greenwood, 2005). Firms respond to technological change not only through the decision to disclose but also through how they define the technology's strategic significance for external audiences. Framing AI in relation to rivalry, market positioning, industry shifts, and competitive advantage allows firms to connect current performance to a forward-looking account of compet- itiveness.

These findings also speak to emerging research on the organizational consequences of generative AI. Much of this literature has focused on productivity, labor substitution, investment, or financial-market responses (Brynjolfsson et al., 2025; Eisfeldt et al., 2023; Felten et al., 2021). The evidence here suggests that an additional dimension of organizational response lies in corporate discourse. In the early stages of a technological shift, firms appear to compete not only through operational changes and investments but also through how they publicly interpret, frame, and position the new technology for external audiences.

The findings also carry managerial implications. When a field-wide technological shock increases uncertainty about the future basis of competition, disclosure becomes a means by which firms manage external evaluations. Managers of firms performing strongly relative to their peers may be better able to connect current success to future relevance. At the same time, the results suggest that the value of such communication depends on industry context. In industries where AI is more central to value creation and competitive differentiation, stakeholders may pay closer attention not only to whether firms discuss AI, but also to how they frame its significance.

5.2 Limitations and Future Research

This study has several limitations that help define its scope and point to directions for future research. First, the analysis identifies disclosure behavior rather than the internal decision processes or AI-related activities that produce it. I argue that favorable peer-relative performance provides firms with a more credible basis for communicating preparedness for technological change, but archival disclosure data do not allow me to observe managers' cognitive interpretations, internal deliberations, or communication strategies directly (Greve, 1998; Iyer and Miller, 2008; Kim et al., 2015). Future research could complement this design with interviews, surveys, or experiments to examine more directly how managers interpret performance feedback and how those interpretations shape communication following technological shocks.

A related limitation concerns the interpretation of AI-related disclosure. Although I separately examine AI-use disclosure and account for post-shock disclosure changes among intangible-intensive firms, 10-K filings do not allow me to directly observe firms' internal AI investments, adoption decisions, or experimentation processes. Some of the observed disclosure responses may therefore reflect actual changes in AI-related activity rather than external communication alone. Similarly, the observational design cannot fully eliminate the possibility that unobserved time-varying organizational capabilities influence both firm performance and AI-related disclosure. The additional analyses using pre-shock AI orientation, entropy-balanced comparisons, industry-year fixed effects, two-year lagged performance-feedback measures, and low-AI-salience industry subsamples reduce these concerns, but they do not exploit an instrumental variable or another quasi-random source of variation in performance feedback. The results should therefore be interpreted as evidence of a robust association between peer-relative performance and post-shock AI disclosure, rather than as definitive causal evidence. Future research could combine disclosure data with direct measures of AI adoption, AI investment, AI hiring, AI patents, cloud-computing expenditures, or internal technology adoption to more sharply distinguish performance-feedback effects from capability-based explanations.

Second, the study focuses on 10-K filings, which represent only one component of firms' broader communication strategies. Form 10-K filings are useful because they provide standardized, regulated, and comparable disclosures, but they are also less flexible and less immediate than other communication channels. Firms also communicate about emerging technologies through earnings calls, investor presentations, press releases, websites, and social media, each of which differs in audience, timing, and rhetorical flexibility (Dyer et al., 2017). Future research could examine whether similar performance-feedback dynamics appear across multiple communication channels and whether firms tailor AI-related narratives to different external audiences.

Third, the measures rely on LLM-assisted text classification. Although the validation results suggest that the approach is conservative and performs well in identifying substantive AI disclosure and category-level framing, measurement error remains possible. In particular, some AI-related passages may be missed, and some category boundaries—such as the distinction between AI opportunity and competitive AI framing—may remain difficult to classify perfectly. Future work could refine these measures through alternative models, larger context windows, human-in-the-loop annotation procedures, or more granular coding schemes that distinguish, for example, proactive competitive positioning from reactive competitive pressure (Li et al., 2021).

Fourth, the empirical setting centers on generative AI following the public release of ChatGPT. This setting is analytically useful because it provides a visible and field-wide technological shock, but it is also distinctive in the speed with which the technology entered public discourse and became salient to managers and investors. As a result, caution is warranted in generalizing the magnitude or timing of these effects to other technological discontinuities. Future research could examine whether similar disclosure dynamics emerge around technologies that diffuse more slowly, are less visible to investors, or vary more sharply in their relevance across firms and industries.

Finally, the weaker findings for historical aspirations should not be interpreted as evidence that historical benchmarks are generally unimportant (Audia and Greve, 2006). In the behavioral theory of the firm, firms may draw on multiple aspiration sources, and the behavioral salience of those sources depends on the setting. My argument is narrower: when firms communicate publicly in response to a highly visible, field-wide technological shock, peer-relative standing should be especially salient because both managers and external audiences are likely to interpret performance comparatively. Historical aspirations remain relevant as self-referential benchmarks tied to prior experience, but they are less directly aligned with the externally comparative nature of corporate disclosure in this context. Future research could examine the conditions under which social versus historical aspirations become more influential in shaping organizational communication, including differences in industry structure, technological salience, and the observability of firm capabilities.

References

윤민범, 김재희, 김희웅(2025), “경영학계 AI 학술 연구의 공백 탐색: 산업 트렌드 기반 비교 분석,” **경영학연구**, 54(6), pp.1459-1483.

(Yoon, M., Kim, J., and Kim, H.-W.(2025), “Exploring Academia Research Gaps in Management Studies on Artificial Intelligence: An Industry-Based Comparative Analysis,” *Korean Management Review*, 54(6), pp.1459-1483.)

Acemoglu, D. and Restrepo, P.(2019), “Automation and New Tasks: How Technology Displaces and Reinstates Labor,” *Journal of Economic Perspectives*, 33(2), pp.3-30.

Acemoglu, D. and Restrepo, P.(2022), “Tasks, Automation, and the Rise in U.S. Wage Inequality,” *Econometrica*, 90(5), pp.1973-2016.

Ahn, Y. Y., Chang, J. H., and Lee, J. H.(2007), “Voluntary Disclosure, Analyst Following, Forecast Accuracy and Dispersion,” *Korean Management Review*, 36(3), pp.765-790.

Audia, P. G. and Greve, H. R.(2006), “Less Likely to Fail? Low Performance, Firm Size, and Factory Expansion in the Shipbuilding Industry," *Management Science*, 52(1), pp.83-94.

Autor, D. H., Levy, F., and Murnane, R. J.(2003), "The Skill Content of Recent Technological Change: An Empirical Exploration," *Quarterly Journal of Economics*, 118(4), pp.1279-1333.

Babina, T., Fedyk, A., He, A., and Hodson, J.(2024), "Artificial Intelligence, Firm Growth, and Product Innovation," *Journal of Financial Economics*, 151, pp.103745.

Baum, J. A. C., Rowley, T. J., Shipilov, A. V., and Chuang, Y. T.(2005), "Dancing with Strangers: Aspiration Performance and the Search for Underwriting Syndicate Partners," *Administrative Science Quarterly*, 50(4), pp.536-575.

Benner, M. J. and Tripsas, M.(2012), "The Influence of Prior Industry Affiliation on Framing in Nascent Industries," *Academy of Management Journal*, 55(2), pp.277-299.

Brynjolfsson, E., Rock, D., and Syverson, C.(2021), "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies," *American Economic Journal: Macroeconomics*, 13(1), pp.333-372.

Brynjolfsson, E., Li, D., and Raymond, L. R.(2025), "Generative AI at Work," *The Quarterly Journal of Economics*, 140(2), pp.889-942.

Carlson, N. A. and Burbano, V.(2025), "The Use of LLMs to Annotate Data in Management Research: Foundational Guidelines and Warnings," *Strategic Management Journal*.

Cyert, R. M. and March, J. G.(1963), *A Behavioral Theory of the Firm, Prentice-Hall, Englewood Cliffs, NJ.*

Dyer, T., Lang, M., and Stice-Lawrence, L.(2017), "The Evolution of 10-K Textual Disclosure: Evidence from Latent Dirichlet Allocation," *Journal of Accounting and Economics*, 64(2-3), pp.221-245.

Eisfeldt, A. L., Schubert, G., and Zhang, M. B.(2023), "Generative AI and Firm Values," *NBER Working Paper No. 31222*, National Bureau of Economic Research.

Felten, E., Raj, M., and Seamans, R.(2021), "Occupational, Industry, and Geographic Exposure to Artificial Intelligence: A Novel Dataset and Its Potential Uses," *Strategic Management Journal*, 42(12), pp.2195-2217.

Fiol, C. M.(1989), "A Semiotic Analysis of Corporate Language: Organizational Boundaries and Joint Venturing," *Administrative Science Quarterly*, 34(2), pp.277-303.

Frey, C. B. and Osborne, M. A.(2017), "The Future of Employment: How Susceptible Are Jobs to Computerisation?," *Technological Forecasting and Social Change*, 114, pp.254-280.

Granqvist, N., Grodal, S., and Woolley, J. L.(2013), "Hedging Your Bets: Explaining Executives' Market Labeling Strategies in Nanotechnology," *Organization Science*, 24(2), pp.395-413.

Greve, H. R.(1998), "Performance, Aspirations, and Risky Organizational Change," *Administrative Science Quarterly*, 43(1), pp.58-86.

Greve, H. R.(2003), "A Behavioral Theory of R&D Expenditures and Innovations: Evidence from Shipbuilding," *Academy of Management Journal*, 46(6), pp.685-702.

Greve, H. R.(2008), "A Behavioral Theory of Firm

Growth: Sequential Attention to Size and Performance Goals," *Academy of Management Journal*, 51(3), pp.476-494.

Han, S. Y. and Park, K. M.(2014), "Determinants of Technological Breadth in the U.S. Semiconductor Industry: Performance Feedback, Recency of Technological Inputs, and Technology Maturity," *Korean Management Review*, 43(1), pp.41-66.

Harrison, J. S., Josefy, M. A., Kalm, M., and Krause, R.(2023), "Using Supervised Machine Learning to Scale Human-Coded Data: A Method and Dataset in the Board Leadership Context," *Strategic Management Journal*, 44(7), pp.1780-1802.

Healy, P. M. and Palepu, K. G.(2001), "Information Asymmetry, Corporate Disclosure, and Capital Markets: A Review of the Empirical Disclosure Literature," *Journal of Accounting and Economics*, 31(1-3), pp.405-440.

Iyer, D. and Miller, K. D.(2008), "Performance Feedback, Slack, and the Timing of Acquisitions," *Academy of Management Journal*, 51(4), pp.808-822.

Jegadeesh, N. and Wu, D.(2013), "Word Power: A New Approach for Content Analysis," *Journal of Financial Economics*, 110(3), pp.712-729.

Kaplan, S.(2008), "Framing Contests: Strategy Making under Uncertainty," *Organization Science*, 19(5), pp.729-752.

Kaplan, S. and Tripsas, M.(2008), "Thinking about Technology: Applying a Cognitive Lens to Technical Change," *Research Policy*, 37(5), pp.790-805.

Kim, J. Y., Finkelstein, S., and Haleblian, J. (2015), "All Aspirations Are Not Created Equal: The Differential Effects of Historical and Social Aspirations on Acquisition Behavior," *Academy of Management Journal*, 58(5), pp.1361-1388.

Kuusela, P., Keil, T., and Maula, M.(2017), "Driven by Aspirations, but in What Direction? Performance Shortfalls, Slack Resources, and Resource-Consuming vs. Resource-Freeing Organizational Change," *Strategic Management Journal*, 38(5), pp.1101-1120.

Leuz, C. and Verrecchia, R. E.(2000), "The Economic Consequences of Increased Disclosure," *Journal of Accounting Research*, 38 (Supplement), pp.91-124.

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), pp.1049-1102.

Li, K., Mai, F., Shen, R., and Yan, X.(2021), "Measuring Corporate Culture Using Machine Learning," *The Review of Financial Studies*, 34(7), pp.3265-3315.

Loughran, T. and McDonald, B.(2011), "When Is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-Ks," *Journal of Finance*, 66(1), pp.35-65.

Lounsbury, M. and Glynn, M. A.(2001), "Cultural Entrepreneurship: Stories, Legitimacy, and the Acquisition of Resources," *Strategic Management Journal*, 22(6-7), pp.545-564.

Miric, M., Jia, N., and Huang, K. G.(2023), "Using Supervised Machine Learning for Large-Scale Classification in Management Research: The Case for Identifying Artificial Intelligence Patents," *Strategic Management Journal*, 44(2), pp.491-519.

Singh, J. V.(1986), "Performance, Slack, and Risk Taking in Organizational Decision Making," *Academy of Management Journal*, 29(3), pp.562-585.

Stocken, P. C.(2000), "Credibility of Voluntary Disclosure," *RAND Journal of Economics*, 31(2), pp.359-374.

Suddaby, R. and Greenwood, R.(2005), "Rhetorical Strategies of Legitimacy," *Administrative Science Quarterly*, 50(1), pp.35-67.

Tripsas, M.(1997), "Unraveling the Process of Creative Destruction: Complementary Assets and Incumbent Survival in the Typesetter Industry," *Strategic Management Journal*, 18(S1), pp.119-142.

Tushman, M. L. and Anderson, P.(1986), "Technological Discontinuities and Organizational Environments," *Administrative Science Quarterly*, 31(3), pp.439-465.

Verrecchia, R. E.(2001), "Essays on Disclosure," *Journal of Accounting and Economics*, 32 (1-3), pp.97-180.

• Injae Jeon is currently a research fellow at the Institute for Business Research and Education, Korea University. He received his bachelor's degree in Business Administration from Hanyang University, a master's degree in Public Policy from the KDI School of Public Policy and Management, and a Ph.D. in Business Administration from Korea University. While working at the KDI Center for International Development, he gained firsthand experience with changes in the non-market environment and the roles of governments and firms in responding to them. Building on this experience, his research focuses on firms' strategic responses to major environmental shifts, including geopolitical rivalry, climate change, and the AI transformation.