

AI Washing

Abstract

This paper investigates AI washing — the exaggeration or misrepresentation of corporate investment in artificial intelligence (AI). Using large language models, we develop novel measures of forward-looking AI investment plans from earnings call transcripts (“AI talk”) and AI-related workforce expertise from employee resumes (“AI walk”) for U.S. public firms from 2016 to 2024. We validate these measures by demonstrating that AI walk, but not AI talk, predicts subsequent AI patent quantity and quality. Within firms, past AI talk does not forecast future AI walk, and AI washing incidents have surged since 2019, particularly among smaller, less capital-intensive manufacturing firms. Market rewards talk in the short run but discounts it in the long run, while walk earns large, persistent valuation gains only over longer horizons. Institutional investors are more discerning than the broad market, allocating capital toward high-walk firms early on. Finally, firms with strong managerial incentives are more likely to raise talk without increasing walk, consistent with strategic hype. Overall, the results reveal a measurable and growing disconnect between AI rhetoric and real investment, reflecting the tension between short-term market incentives and long-term value creation.

Keywords: Artificial Intelligence, Large Language Models (LLMs), Corporate Disclosure, Investment, Technology and Innovation

JEL classifications: D22, E22, G10, G32, O32

1 Introduction

Artificial intelligence (AI) has rapidly emerged as one of the most transformative technologies of the modern era. AI’s capability to analyze big data, automate complex tasks, and improve decision-making suggests that it has the potential to function as a general-purpose technology for business across industries. Despite its potential, however, AI’s economic impact remains an open question. A critical challenge in assessing AI’s impact is that firms do not uniformly disclose AI investments, making it difficult to differentiate between genuine AI-investing firms and those who merely claim to be investing in it.

As investors and stakeholders increasingly recognize the value of AI capabilities, some firms may be tempted to engage in *AI washing* — exaggerating or falsely claiming AI investment plans. Recent anecdotal evidence suggests that the Securities and Exchange Commission (SEC) is scrutinizing firms’ claims about AI investment, highlighting the regulatory importance of this issue.¹ Beyond regulatory concerns, AI washing can lead to capital misallocation, diverting investment toward firms that merely signal AI investment rather than those genuinely engaged in AI development.

In this study, we frame AI washing through the lens of agency conflict between managers and shareholders, drawing on [Stein \(1988\)](#) model of managerial myopia. Managers, pressured by short-term stock price expectations and performance-based compensation, may engage in AI washing by prioritizing *AI talk* (i.e., public claims of AI investment) over *AI walk* (i.e., tangible investments in AI). Given the information asymmetry between managers and investors, firms may use cheap talk as a low-cost signal to exploit investor enthusiasm for AI. This can create a divergence between market perceptions and firm fundamentals: while AI talk may temporarily boost stock prices, only AI walk contributes to long-term innovation and value creation. When firms overstate their AI engagement without follow-through, shareholders ultimately bear the cost of misallocated capital and missed innovation

¹<https://www.bloomberg.com/news/articles/2024-06-20/sec-s-ai-crackdown-signals-trickle-of-cases-will-turn-to-flood?embedded-checkout=true>

opportunities.

We provide the first systematic examination of AI washing among U.S. public firms from 2016Q1 to 2024Q2. Specifically, we quantify the discrepancies between firms’ past AI talk, as measured by forward-looking in-house AI investment claims in quarterly earnings conference calls, and their future AI walk, as measured by AI-related human capital. While regulatory filings (e.g., 10-K and 8-K reports) contain AI-related disclosures, we focus on earnings conference calls, as they represent high-stakes, real-time communications where managers articulate firm-specific information to investors and analysts with relatively greater discretion. Over the sample period, AI talk has risen sharply: the percentage of firms mentioning the word “AI” in conference calls increased from virtually zero in 2016 to about 20% by mid-2024 (Figure 1), raising questions about the credibility and consistency of these claims.

[Insert Figure 1 here]

Given the rapid development of AI technology over the past few years, our analysis accounts for evolving terminology and language change by constructing dynamic word embeddings based on conference call transcripts. We identify keywords that are most closely associated with “AI” for each year, which then allow us to quantify the extent to which firm managers discuss their AI commitments. To refine our measurement of AI talk, we apply large language models (LLMs) to classify only those statements that describe active, in-house AI investments requiring human capital and are framed in a forward-looking manner. This approach yields a firm–quarter measure of AI talk that reflects active investment claims rather than generic references to AI.

We measure AI walk using detailed employee resume data that captures the composition of firms’ AI-related human capital, following the approach of Babina et al. (2024). This dataset allows us to identify active AI-related positions at the firm–quarter level and quantify the share of the workforce engaged in AI activities. By construction, this measure captures

substantive workforce investment in AI and enables a direct comparison between firms’ stated AI ambitions and their realized commitments.

We validate our measures by examining their relationship with subsequent AI innovation, as captured by AI patent quantity, value, and citations. AI walk is a strong and robust predictor: a one-standard-deviation increase is associated with roughly 17% more AI patents, 25% higher patent value, and 33% more citations in the following period. These economically large and statistically significant relationships suggest that the walk measure reflects substantive AI capabilities that translate into meaningful and higher-quality innovation. By contrast, holding AI walk constant, AI talk has no significant association with any innovation outcome, suggesting that rhetorical emphasis on AI, has limited technological impact.

Equipped with the granular firm-quarter level measures, we then examine the dynamic relationship between AI talk and AI walk. Cross-sectionally, firms that talk more about AI also tend to have larger AI-related workforces in subsequent periods. However, this relationship is largely explained by time-invariant firm characteristics such as size, age, and R&D intensity. Once we exploit within-firm variation by including firm fixed effects, the predictive link between past AI talk and future AI walk disappears, indicating that managerial rhetoric around AI investment generally does not translate into subsequent workforce expansion. This disconnect has widened over time: before 2019, firms tended to “walk the talk,” whereas in the post-2019 period, AI talk no longer predicts future AI walk.

This breakdown in the within-firm relationship between AI talk and subsequent AI walk suggests that some firms use AI-related rhetoric without following through on substantive investment. To examine this behavior systematically, we define AI washing incidents as cases where a firm announces forward-looking AI investment plans but undertakes zero AI-related workforce capacity within the subsequent two years. We identify 165 AI washing firms — those with at least one incident over the sample period — and document a substantial rise in such behavior since 2019, concentrated in manufacturing and more prevalent among smaller and less capital-intensive firms. This pattern is consistent with the idea that resource-

constrained firms may use AI talk as a low-cost signaling strategy in place of substantive investment.

We then examine how the stock market responds to such behavior. In the short run, AI talk corresponds to a positive and statistically significant stock price reaction of approximately 1 basis point higher three-day cumulative abnormal returns (CARs) per 3 additional AI-related keywords on average, controlling for the level of AI walk. However, this initial enthusiasm fades over time: by the 360-day horizon, firms with higher AI talk experience statistically significant under-performance. In contrast, firms with higher AI walk show muted short-run reactions but earn sustained and significant positive abnormal returns over longer horizons — a one-standard-deviation higher AI walk is associated with 25 basis points higher BHAR over the 360-day horizon. This divergence suggests that the market may initially reward AI-related rhetoric, only to later recognize its lack of substance and penalize it, while substantive AI investment yields slower but far greater long-term gains in market value.

On the other hand, institutional investors can better distinguish between rhetoric and substance in firms' AI investment disclosures. For AI-focused mutual funds and ETFs, a one-standard-deviation increase in AI walk corresponds to roughly 0.3 additional fund holding the stock, while AI talk shows negative association within three quarters. In the broader institutional fund universe, a one-standard-deviation increase in walk is associated with 12-14 more funds holding the stock across the subsequent four quarters, whereas talk remains insignificant throughout the horizon. These patterns indicate that institutional investors specialized in AI tend to reward substantive AI investment and, within three quarters, reduce exposure to firms with elevated AI talk, while generalist funds largely ignore rhetorical emphasis and respond primarily to tangible investment.

Finally, we investigate managerial incentives for AI washing. Using CEO *delta*, which is a measure of the sensitivity of an executive's wealth to the firm's stock price, as a proxy for pay-for-performance incentives, we compare changes in talk and walk before and after the launch of ChatGPT for high- versus low-delta firms. A difference-in-differences design reveals that

high-delta firms increase AI talk sharply in the post-ChatGPT period, while showing only small and statistically insignificant changes in AI walk. Parallel-trend tests confirm that these differences emerge only after the event. These findings suggest that when executives' personal wealth is more tightly tied to stock price movements, they may have greater incentives to strategically emphasize AI in public communications without committing commensurate resources.

We find similar patterns when we focus on seasoned equity offerings (SEOs), which create strong short-term incentives to boost market sentiment. AI talk spikes sharply in the SEO quarter, while coinciding with no observable increase in AI walk even after four quarters. The market responds disproportionately: during SEO quarters, the return premium to AI talk is more than seven times larger than in other periods, adding roughly 1.6 basis points to three-day CARs for a one-standard-deviation increase in talk. These results are consistent with opportunistic narrative management, where managers strategically amplify AI-related rhetoric when they stand to benefit from a higher stock price, despite the absence of substantive follow-through.

Together, these findings provide novel evidence that AI washing is both increasingly prevalent and strategically motivated, shaped by a tension between short-term market incentives and long-term value creation. By separately measuring AI talk and AI walk, we uncover a clear disconnect between public communication and actual investment, with important implications for investors and regulators.

This paper contributes to the growing literature on the role of AI in firms. Prior research has examined the impact of AI on firm growth and workforce composition (e.g., [Acemoglu and Restrepo \(2018\)](#); [Agrawal et al. \(2019\)](#); [Webb \(2019\)](#); [Acemoglu et al. \(2022\)](#); [Babina et al. \(2023a\)](#); [Babina et al. \(2023b\)](#); [Eisfeldt et al. \(2023\)](#); [Babina et al. \(2024\)](#); [Hampole et al. \(2025\)](#)). However, to the best of our knowledge, we are the first to systematically investigate the gap between firms' AI claims and their actual investments up to 2024. By leveraging granular textual data from conference calls and employee resumes, our study sheds

light on how firms strategically communicate AI engagement and the extent to which such communication aligns with real technological development.

Our study also contributes to the broader literature on corporate disclosure and strategic communication. Prior studies suggest that managers selectively shape firm narratives to influence investor perceptions and stock prices (e.g., [Larcker and Zakolyukina \(2012\)](#); [Gow et al. \(2021\)](#); [Flugum and Souther \(2023\)](#)). Specifically, there is a growing literature on the phenomenon of greenwashing, where firms misreport on their environmental, social, and governance (ESG) practices to appear more sustainable than they truly are (e.g., [Chen \(2022\)](#); [Bauer et al. \(2024\)](#); [He et al. \(2024\)](#)). While greenwashing often relates to long-term reputational positioning and is increasingly subject to regulatory scrutiny, AI washing is distinct in that it centers on a fast-moving technological frontier where disclosure standards are nascent, verification is more difficult, and market enthusiasm can produce rapid valuation effects.

From a methodological perspective, this study contributes to the burgeoning literature that applies natural language processing (NLP) techniques in finance research (e.g., [Jha et al. \(2022\)](#); [Bandyopadhyay et al. \(2023a,b\)](#); [Hirshleifer et al. \(2023\)](#); [Lopez-Lira and Tang \(2023\)](#); [Sautner et al. \(2023\)](#); [Bybee et al. \(2024\)](#); [Jha et al. \(2024\)](#); [Li et al. \(2024\)](#); [van Binsbergen et al. \(2024\)](#); [Hirshleifer et al. \(2025\)](#)). Early methodologies, such as Word2Vec and BERT, exhibit limitations in fully interpreting context, particularly in complex corporate disclosures. Recent developments in LLMs offer improvements in contextual understanding but present challenges in terms of computational cost and scalability when applied to large corpora of text. To address these challenges, this study proposes a novel hybrid approach that combines dynamic word embeddings with LLM-based analysis. Relatedly, [Jha et al. \(2024\)](#) use GPT-based analysis of earnings calls to extract forward-looking corporate policies such as investment and employment. They validate their measures against survey data and show that LLM-inferred expectations predict realized firm behavior. Their findings demonstrate the power of LLMs in extracting credible policy signals from managerial language. Our study

builds on this idea but emphasizes that when applied to AI discourse specifically, rhetoric often outpaces realized workforce investment, underscoring the existence of AI washing.

There are a series of contemporaneous studies on firms' AI-related communication. [Barrios et al. \(2025\)](#) examine U.S. public firms' AI disclosures in earnings calls from 2016–2023, showing that cross-sectionally, the release of ChatGPT in late 2022 sparked a rapid shift toward generative AI topics. [Jia et al. \(2025\)](#) focus on a short window of two quarters after the ChatGPT release, providing descriptive evidence that firms, on average, increase their discussions about generative AI. In a contemporaneous study on firms' AI disclosure and labor investment, [Liu \(2025\)](#) finds that from 2010–2018, AI-exposed firms increased subsequent AI-related hiring. Consistent with this, we find that before 2019, AI talk significantly predicts AI walk. However, extending the sample through 2024 reveals that this link has largely disappeared in recent years, suggesting a growing divergence between firms' AI narratives and their actual investments.

Our study complements these papers while offering four novel insights. First, by exploiting within-firm variation over the recent 2016–2024 period, we provide a granular and timely assessment of the credibility and consequences of corporate AI communication. Second, we examine how the stock market responds differently to AI talk and AI walk across horizons, showing that initial enthusiasm for talk reverses over time, whereas walk delivers sustained value creation. Third, we contrast institutional investors with the broader market, finding that institutions, particularly AI-focused funds, are better able to distinguish between genuine investment and rhetorical signaling, and act on this information more quickly. Finally, we show that managerial incentives can trigger opportunistic spikes in AI talk without corresponding investment, offering direct evidence of managerial myopia in AI disclosure.

The remainder of the paper is organized as follows. Section 2 describes the data sources. Section 3 outlines the methodology and presents key stylized facts. Section 4 validates the AI talk and AI walk measures. Section 5 analyzes the dynamics of these measures and documents descriptive patterns of AI washing firms. Section 6 examines financial market

reactions to AI washing, and Section 7 investigates the role of managerial incentives. Section 8 concludes.

2 Data

2.1 Earnings Call Transcripts

To measure the intensity of firms' AI talk, we analyze quarterly earnings call transcripts sourced from the Capital IQ database. These transcripts provide a rich textual dataset capturing how firms communicate their strategic priorities, technological initiatives, and financial performance to investors and analysts. Our focus is on U.S. public firms from 2016 to 2024.

Earnings calls represent a unique and informative venue for analyzing firms' AI-related discourse. Unlike regulatory filings, which follow standardized reporting requirements, earnings calls offer greater managerial discretion and flexibility in shaping the narrative around firm performance, competitive positioning, and innovation strategies. Importantly, these calls are not subject to direct regulatory constraints, allowing managers to emphasize certain aspects of their operations, including AI investments, based on strategic considerations.

Our analysis specifically focuses on the managerial presentation segment of the earnings call, where executives deliver prepared remarks about their firms. This section is particularly relevant because it reflects the firm's intentional messaging, rather than being driven by external questioning. While the Q&A segment involves interactive discussions with analysts, the presentation segment offers a clearer and more controlled measure of the firm's AI-related disclosures.

2.2 Employee Resume

To measure firms' actual investments in AI, we follow [Babina et al. \(2023a,b, 2024\)](#) and use AI-related human capital as a proxy. We draw on resume data from Revelio Labs, which

compiles structured employment histories from online profiles. Our dataset is a snapshot as of August 2024, in which individuals report past roles, descriptions, and skills. Restricting to U.S. publicly traded firms yields roughly 75 million position records for about 44 million unique employees, including company affiliations, job titles and descriptions, tenure in roles, listed skills, and educational backgrounds. We aggregate these records to the firm–quarter level to characterize the composition of each firm’s active workforce.

Relative to job postings, resume data provide a better measure for realized AI investment (i.e., stocks of AI-skilled employees) rather than intended demand (i.e., vacancy flows that may be duplicated, unfilled, or noisy). Furthermore, resume data capture internal moves, promotions, and hires that would not appear in postings, and they offer richer insights into job descriptions for assessing actual AI development.

2.3 Institutional Holdings

To examine institutional investors’ exposure to AI-related firms, we utilize fund holdings data from the CRSP Mutual Fund Database. This dataset provides comprehensive coverage of U.S. mutual funds and ETFs, including their portfolio compositions, investment strategies, and quarterly holdings.

Given the absence of a formal definition of AI-focused funds in the existing literature, we develop a novel classification approach based on fund prospectus using LLM. Out of all the fund prospectuses that are filed between 2016Q1 and 2024Q2 from the SEC website, we first require the presence of “AI,” “artificial intelligence,” or “technology” in the *strategy narrative* section of the prospectus. Then, we feed the entire section to LLMs to classify each fund to be AI-investing, AI-using, or neither.² This procedure yields 98 unique AI-investing fund portfolios, for which we obtain quarterly portfolio holdings from the CRSP Mutual Fund Database.

²The fund is classified as an AI-Investing fund if it explicitly states that it invests in firms that develop, advance, or heavily utilize AI in their business models (e.g., AI research, AI infrastructure, machine learning applications). The fund is classified as an AI-using fund if it applies AI in its own investment strategies (e.g., AI-driven algorithms for portfolio selection). The complete prompt is shown in Section A2 of the Appendix.

2.4 AI Patents

We utilize the Artificial Intelligence Patent Dataset developed by the United States Patent and Trademark Office (USPTO), which classifies AI patents into several categories, including machine learning, evolutionary computation, natural language processing, vision, speech, knowledge processing, planning and control, and AI hardware (Pairolero et al. (2025)). The dataset includes firm identifiers, enabling us to track firms’ technological innovation in AI over time. To assess both the *quantity* and *quality* of AI innovation, we supplement this dataset with measures of economic value and forward citations from Kogan et al. (2017), which serve as proxies for the patents’ economic and scientific significance.

The USPTO mandates that all patent applications filed on or after November 29, 2000, be made public 18 months after the initial filing date.³ To account for truncation bias, we limit our analysis on innovation outcomes to the end of 2022. To avoid truncation bias from unpublished applications, we restrict our analysis of innovation outcomes to patents issued by U.S. public firms through the end of 2022. Patenting activity is aggregated at the firm-quarter level using the *filing date*, rather than the grant date, to better reflect the timing of innovation given the typical lag between filing and grant. This approach yields a firm-quarter measure of newly generated AI patents and provides a more accurate, near-real-time indicator of firms’ AI-related innovation output.

2.5 AI Inventors

We construct an alternative measure of firms’ substantive engagement in AI based on the hiring of AI inventors. Using data from the USPTO, we extract information on the filing date, inventor name, and inventor location for each AI-related patent. We then match inventors to employee profiles from Revelio Labs by performing exact name and location matching, and requiring that the matched individual is employed at the firm at the time of patent filing. This matching procedure enables us to track firms’ recruitment of AI inventors

³<https://www.uspto.gov/web/offices/pac/mpep/s1120.html>

and aggregate this information to the firm-quarter level.

This alternative AI walk measure offers a narrow but highly selective view of firms' AI investment, focusing specifically on those employees who have contributed to patented innovations. In contrast, our baseline walk measure captures the broader AI workforce, including engineers, data scientists, and other AI-related professionals who may not be directly involved in patenting. Later in the analysis, we show that this inventor-based measure yields consistent evidence of AI washing, providing robustness to our main findings and reinforcing the distinction between symbolic AI talk and substantive AI walk.

2.6 Sample Summary Statistics

To obtain our final sample, we impose two important filters. First, we exclude firms in the information and technology industries, following [Babina et al. \(2024\)](#). This is because we are interested in studying firms that invest in AI to benefit their business operations rather than those who merely sell AI products. Second, we impose a minimum threshold of 100 active employees in the resume data for a firm to be included in the sample. This criterion helps mitigate concerns related to firms with limited workforce representation, which could otherwise distort our AI investment measure. We focus on firms that have discussed AI-related topics at least once in their conference calls from 2016Q1 to 2024Q2. Our final sample includes 20,135 firm-quarter observations covering 721 unique U.S. public firms. The summary statistics of the final sample are presented in [Table 1](#).

[Insert [Table 1](#) here]

3 Methodology and Descriptive Evidence

3.1 Identifying AI-related Keywords

Given the rapid development of AI technology, it is important to account for the evolving terminology used by firm managers. To this end, we use the skip-gram implementation of the *Word2Vec* algorithm, a widely used natural language processing (NLP) technique that generates vector representations of words based on the prediction of surrounding context words. Unlike traditional bag-of-words models, *Word2Vec* captures semantic relationships and contextual similarity. We follow the approach proposed by [Houston et al. \(2024\)](#) and train separate *Word2Vec* models for each year in our sample using conference call transcripts, thereby constructing a dynamic word embedding that reflects the evolution of AI-related language over time.

When training our models, we identify bi-grams and tri-grams in the corpus to preserve multi-word phrases.⁴ We use a vector dimension of 300 and a context window of 10 words, and drop infrequent words that appear fewer than five times. To improve model stability, we run the skip-gram algorithm 20 times (rather than the default of 5). The resulting embeddings are robust to alternative window sizes of 5 and 15.

From each yearly *Word2Vec* model, we identify the top 100 terms most closely associated with “AI” based on cosine similarity scores. As shown in [Figure 2](#), there has been substantial change in the composition of the top-100 most similar words to “AI” over time. For instance, generative AI terms begin to appear prominently in 2023, shortly after the launch of ChatGPT at the end of 2022.

[Insert [Figure 2](#) here]

⁴For example, “*artificial*” and “*intelligence*” are two separate words, but are frequently used together in the bigram, “*artificial_intelligence*”.

3.2 AI Talk Measure

To construct our measurement of AI talk, we first apply LLMs to further analyze firms' discussions containing any AI-related keywords. In particular, we classify the AI-related text from conference call transcripts using the GPT-4o-mini model with the following prompt:⁵

Instruction: Analyze the following manager presentation excerpt from a company's conference call transcript. Based solely on this source, determine whether the company is making a clear commitment to AI-related activities that require human capital—such as hiring AI talent, expanding AI R&D teams, or building internal AI development capacity.

Classification Criteria:

- *AI Investment Talk: The firm describes in-house AI activities involving human capital, including developing AI technology, expanding R&D staff, hiring AI specialists, or creating dedicated AI teams.*
- *Not AI Investment Talk: No mention of in-house AI activities requiring human capital. General remarks on AI's importance, reliance on external AI providers, or AI activities limited to procurement of tools or infrastructure without internal staffing should be classified here.*

Time Frame:

- *Backward-Looking: AI activities involving human capital that have already occurred (e.g., “We hired...,” “Last year, we built an AI team...”).*
- *Forward-Looking: Planned AI activities involving human capital (e.g., “We will expand our AI research team in 2025...”).*
- *If both backward- and forward-looking activities are present, report past and future time frames separately.*

Time Frame Detail:

- *Identify any specific years, quarters, or relative time references mentioned.*

⁵As an alternative, we also use the gemini-2.0-flash model and the classification outputs are largely identical.

Output Format:

1. *Commitment Status:* (AI Investment / Not AI Investment)
2. *Time Frame:* (Backward-Looking / Forward-Looking / NA)
3. *Time Frame Detail:* (Exact year/quarter or relative reference)
4. *Supporting Evidence:* (Quote key phrases from the transcript that justify the classification)

Based on the output, we further filter the conference call text by requiring *Commitment Status* = “AI Investment” and *Timing* = “Forward-Looking.” This ensures that any remaining discussions are truly related to on-going or future in-house AI investment that the firms are actively engaging in, which allows for comparison with real-time firm employee profiles.⁶ After this step, we construct the AI talk measure as following:

$$AI\ Talk_{i,t} = \frac{\sum_{k=1}^K Similarity\ Score_{i,t,k} \times AI\text{-related}\ Keyword\ Occurrence_{i,t,k}}{Total\ Number\ of\ Words_{i,t}} \times 100, \quad (1)$$

where i denotes firm, t denotes quarter, and k denotes keyword.

The term *AI-related Keyword Occurrence* represents the frequency of each AI-related term in firm i ’s earnings call transcript during quarter t . The term *Total Number of Words* serves as a normalization factor to ensure that the AI talk measure reflects the relative prominence of AI discussion rather than variations in transcript length. By incorporating word similarity scores, our measure assigns greater weight to terms that are more closely related to AI concepts in a given year, allowing for a more accurate representation of a firm’s AI discourse.

Of all forward-looking AI investment discussions, roughly 65% specify a clear time frame. Among these, over 90% indicate that the investment is expected to occur within two years of the statement. This pattern is important, as it informs the time horizon we adopt in our subsequent analyses.

⁶For examples of such AI Investment discussions, see Section A1 of the Appendix.

3.3 AI Walk Measure

We use AI-related human capital as a proxy for firms’ actual walk in AI. For job positions that span multiple quarters, we treat them as *active* in each quarter. To identify AI-related positions, we require at least one AI keyword from the corresponding yearly AI dictionary to be present in the position description text. To ensure that a position containing AI keywords is genuinely tied to the development of AI technology, we further filter the remaining descriptions using an LLM-based classification prompt. This additional step helps exclude roles where AI mentions refer only to peripheral activities, such as adopting pre-built AI tools in non-technical workflows or referencing AI tangentially without contributing to in-house AI development or infrastructure.⁷ We use the following prompt with GPT-4o-mini:

Instruction: Analyze the following job description. Based solely on this description, determine whether the position is AI-related.

Classification Criteria:

- *AI-Related: The role contributes to the firm’s in-house AI development or infrastructure, indicating real AI investment. This can include building or deploying machine learning models, researching AI algorithms, or developing tools and systems to support AI work.*
- *Not AI-Related: The role does not directly involve AI model development or infrastructure, such as generic IT or computer science roles.*

Output Format:

1. *AI Classification: (AI-Related / Not AI-Related)*
2. *Supporting Evidence: (Quote key phrases from the description that justify the classification)*

With the filtered AI-related job positions, we construct the firm-quarter level AI walk measure as following:

⁷For examples of AI-related job positions, see Section A1 of the Appendix.

$$AI\ Walk_{i,t} = \frac{\text{Number of Active AI-related Positions}_{i,t}}{\text{Number of All Active Positions with Descriptions}_{i,t-1}} \times 100, \quad (2)$$

where i denotes firm and t denotes quarter.

We acknowledge two limitations of using job description text. First, position descriptions may not always align with the exact timing of work activities, since employees typically record them retrospectively. Job titles, in contrast, are more contemporaneous and thus may better reflect real-time AI engagement. To assess robustness, we construct an alternative walk measure based solely on job titles. Specifically, we use the AI-relatedness scores from [Babina et al. \(2024\)](#) to identify the ten most AI-intensive job titles and then apply an LLM-based classification to categorize all job titles in our data as AI-related or not.⁸ This job-title-based measure correlates strongly with our description-based measure (correlation coefficient = 0.64), and we will show regression results based on this alternative classification as robustness analysis.

Second, even though description text is the richest source of information for identifying AI-related employees, it is not available for every job record. To mitigate this issue, we scale the number of AI-related positions by the lagged number of positions with descriptions available. This normalization should address missing-text concerns as long as there is no systematic selection into providing descriptions based on AI engagement. To test this assumption directly, we calculate the percentage of missing descriptions separately for AI-related and non-AI-related job titles. As shown in [Figure A1](#), while both series exhibit an upward trend over the sample period (with newer records less likely to contain descriptions), the trends are largely parallel between the two groups. This pattern suggests that missing descriptions are not systematically related to AI engagement and are unlikely to bias our results.

[Figure 3](#) plots the time-series trends of the AI talk and AI walk measures. Both measures have exhibited a steady upward trajectory since early 2016, reflecting the growing emphasis

⁸The complete prompt is shown in Section A3 of the Appendix.

on AI-related discussions and investments across firms.

[Insert Figure 3 here]

We also find substantial variation in both AI talk and AI walk across industries. Figure 4 plots the average talk and walk by industry, defined using firms’ two-digit NAICS codes. Both measures are highest in the “Information” industry, consistent with the findings of Babina et al. (2024) over the sample period of 2010-2018. Several other industries also exhibit elevated levels of AI-related discourse and investment, suggesting that AI is not confined to a single sector.

[Insert Figure 4 here]

For easier interpretation, we standardize AI talk and walk measures such that each measure has mean zero and standard deviation of one. We use the standardized measures in remaining analyses.

3.4 Focus of AI Talk

To better understand the focus of firms’ AI talk, we examine the prevalence of key topics extracted from earnings conference call transcripts. Specifically, we apply LLMs to identify the focus of AI investment within firm-level AI talk, isolating themes mentioned in the context of AI-related workforce commitments. Figure 5 presents the time-series frequency of four selected themes — automation, generative AI, cloud, and a generic AI category. Automation emerges as the dominant and most persistent focus, rising steadily from 2016 to a peak of nearly 12% of firms in the sample discussing it as the primary focus of AI investment by 2022, before declining thereafter. Generative AI appears abruptly in late 2022, coinciding with the public release of ChatGPT, and quickly reaches almost 8% of firms in 2023, underscoring the speed at which breakthrough technologies reshape investment narratives. Cloud remains relatively stable at 2–3% across the sample period, reflecting its

role as enabling infrastructure rather than a shifting focal point. In contrast to these specific and well-defined areas of AI investment, mentions of the generic AI category follow a similar post-2022 surge, suggesting that firms increasingly reference AI in broad, non-specific terms — possibly to signal engagement with AI without disclosing concrete investment goals or focus areas. These trends highlight both the persistence of established AI investment themes and the rapid adoption of new, less specific narratives in response to major technological events, reinforcing that our AI talk measure effectively captures the evolving content of firms’ AI investment discussions.

[Insert Figure 5 here]

4 Validation of Measures: Innovation Outcomes

To validate our AI Talk and AI Walk measures, we examine whether they are predictive of subsequent innovation outcomes. If our measures are valid, we expect AI Walk to be positively associated with innovation outcomes, as actual AI capability investments should facilitate the development of novel technologies. Conversely, controlling for AI Walk, AI Talk should have weaker association with innovation as it reflects symbolic signaling rather than substantive action.

Specifically, we use three different measures of AI innovation at the firm-quarter level: (1) the number of AI patents granted, (2) the total economic value of those patents, and (3) the number of forward citations to those patents. The following regression specification is estimated:

$$Y_{i,t} = \beta_1 Talk_{i,t-1} + \beta_2 Walk_{i,t-1} + \gamma \mathbf{X}_{i,t-1} + \alpha_{firm} + \lambda_{industry-quarter} + \varepsilon_{i,t}, \quad (3)$$

where $Y_{i,t}$ takes on one of the three dependent variables, and the key independent variables are $Talk_{i,t-1}$ and $Walk_{i,t-1}$, which measure the standardized levels of AI talk and AI walk

of firm i , respectively, lagged by one period.⁹ The vector $\mathbf{X}_{i,t-1}$ includes a set of lagged control variables capturing firm characteristics and financial indicators, including firm size, cash/assets, R&D intensity, age, and CAPEX. The specification also includes firm fixed effects, α_{firm} , and industry-quarter fixed effects, $\lambda_{industry-quarter}$, to account for time-invariant firm characteristics and time-specific industry-level shocks.

[Insert Table 2 here]

Table 2 presents regression results examining the relationship between firms' AI engagement and their subsequent AI innovation. Across all specifications, AI walk ($Walk_{t-1}$) is a strong and consistent predictor of subsequent innovation: a one-standard-deviation increase in Walk in the previous quarter is associated with approximately 17% more AI patents, 25% higher AI patent value, and 33% more citations in the following period. These economically large and statistically significant relationships indicate that the Walk measure captures substantive AI-related capabilities that translate into meaningful and higher-quality innovation outputs.

In contrast, AI Talk ($Talk_{t-1}$) is not significantly associated with innovation outcomes, suggesting that excessive AI-related language alone has limited and economically trivial implications for actual innovation. Taken together, these results support the validity of the talk-walk distinction: whereas AI Walk is meaningfully associated with tangible innovation outputs, AI Talk may instead reflect hype, aspirational signaling, or strategic communication without real technological follow-through.

[Insert Figure A2 here]

Given the common time lag between technological investment and actual innovation output, we perform distributed lag analysis by regressing patent outcomes on AI walk and AI talk of various lags, along with control variables and fixed effects as in the main specification.

⁹In distributed lag analysis, we lag independent variables for up to eight quarters.

As illustrated in Figure A2, the predictive power of AI talk on future innovation outcomes is robust and strong for 8 quarters ahead, whereas empty AI talk is trivially associated with future innovation, even in the longer term.

5 AI Talk-Walk Dynamics

5.1 Firm Characteristics of AI Talkers and Walkers

To study what type of firms tend to talk and walk on AI, we first have to extend our sample to include all firms, regardless of whether they have mentioned AI over our sample period or not. This step leads to 54,639 firm-quarter observations covering 2,094 unique U.S. public firms. Then we perform univariate analysis to compare several ex-ante firm characteristics as of 2016Q1, which are available for 1,496 firms. Table 3 presents descriptive statistics that compare firms based on their engagement with AI, distinguishing between firms that ever talk about AI (Panel A) and those that ever make AI-related hiring (Panel B).

Across several firm characteristics, firms that talk about AI at least once in our sample period (i.e., *Ever Talk* firms) are systematically larger with more resources. On average, Ever Talk firms have significantly higher sales (6.08 vs. 5.43), more cash holdings (0.19 vs. 0.15), and greater capital expenditures (3.01 vs. 2.43), all with statistically significant differences. These firms are also slightly older (26.9 vs. 25.3) and invest more in R&D (0.018 vs. 0.012). In contrast, profitability (ROA) is nearly identical across both groups, suggesting similar financial performance as of 2016. The magnitude and significance of the mean differences highlight that Ever Talk firms are, on average, more established and better positioned to engage with and discuss emerging technologies like AI.

In comparison, firms that walk at least in one quarter in our sample period (i.e., *Ever Walk* firms) exhibit even stronger distinctions in firm characteristics. These firms are substantially larger, with a mean log sales of 6.37 compared to 4.94 among non-walkers, and demonstrate significantly greater investment capacity (3.31 vs. 1.95). They also tend to be older (29.9

vs. 22.7), indicating a more established corporate foundation. While cash holdings and R&D intensity are moderately higher, profitability (ROA) is notably better among Ever Walk firms (-0.001 vs. -0.010). These patterns suggest that AI investment is more common among mature, better-capitalized firms, reinforcing the notion that translating AI talk into action requires substantial organizational resources and capabilities.

In the remainder of the analysis, we focus on firms that have discussed AI-related topics at least once in their conference calls from 2016Q1 to 2024Q2. Given that the differences in ex-ante firm characteristics between AI-engaging firms and non-engaging firms are relatively small in magnitude, this restriction leaves a sample that remains broadly representative of U.S. public firms while focusing on the subsample most relevant to our study of AI washing.

5.2 Talk-Walk Gap in AI

With measures of AI talk and AI walk, we are now equipped to examine their dynamic relationship with the following regression design:

$$AI\ Walk_{i,t} = \alpha + \sum_{h=1}^8 \beta_h AI\ Talk_{i,t-h} + \gamma \mathbf{X}_{i,t-1} + \alpha_{firm} + \lambda_{industry-quarter} + \varepsilon_{i,t}, \quad (4)$$

where $AI\ Walk_{i,t}$ represents the level of AI investment for firm i in quarter t and the key independent variable, $AI\ Talk_{i,t-h}$, captures the intensity of firm i 's AI investment discussions in its earnings call in quarter $t - h$. The coefficients β_h reflect the estimated predictability of past AI talk at each of the past eight quarters on current AI walk.¹⁰ The vector $\mathbf{X}_{i,t-1}$ includes lagged control variables such as size, cash/assets, R&D, age, and CAPEX. In the first set of regressions, we incorporate industry-quarter fixed effects $\lambda_{industry-quarter}$ to control for time-specific industry-level shocks. In another set of regressions, we further include firm fixed effects λ_{firm} to control for time-invariant firm-specific characteristics that may influence

¹⁰In deciding the number of lags to include, we have to balance the tradeoff between examining a longer time horizon and retaining more observations in the analysis. Given the fact that the vast majority of AI investment talk is about plans within one to two years, we include lagged talk measures up to eight quarters.

both the independent and dependent variables.

[Insert Table 4 here]

The results presented in Table 4 reveal two key findings. First, in the absence of firm fixed effects (columns 1 - 3), past AI talk exhibits strong predictive power for future AI walk, suggesting that cross-sectionally, firms that discuss AI more extensively are also more likely to expand their AI-related workforce. This is also consistent with our findings in Table 3, where more established and better capitalized firms engage in both AI talk and AI walk more. In sharp contrast, however, once we control for spurious cross-sectional variation by incorporating firm fixed effects (columns 4 - 6), this relationship disappears. Specifically, there is no significant association between AI talk over the past eight quarters and current AI walk at the firm level, implying that much of the cross-sectional association may be driven by time-invariant firm characteristics such as size, age, or R&D intensity. These findings underscore the importance of accounting for firm-level heterogeneity when evaluating dynamic relationships between AI rhetoric and implementation.

These findings support the notion that managerial discourse around AI investments often fails to translate into actual organizational commitment, such as hiring AI talent. The absence of a predictive relationship between past AI talk and future AI walk is consistent with the practice of AI washing, where firms emphasize AI narratives in public communications without backing them with substantive implementation later on.

In a series of robustness analyses, we confirm that our results are not sensitive to alternative definitions of AI walk. First, when we measure walk as the flow of new AI employees rather than the stock of AI-related workforce, the findings remain unchanged (as shown in Table A1). Second, relying on job titles rather than descriptions to classify AI-related positions yields similar results, alleviating concerns about missing or retrospective description text (as shown in Table A2). Third, we construct a narrower walk measure based on the hiring of AI inventors, which captures a highly selective but arguably more innovation-focused

dimension of AI investment; the results are again consistent (as shown in Table A3). Finally, when we re-define AI talk and walk using dummy variables that capture the extensive margin of engagement, our conclusions hold (as shown in Table A4).

[Insert Table 5 here]

In a contemporaneous working paper by Liu (2025), the author studies a sample period from 2010 to 2018 and finds that AI-exposed firms subsequently increase their investments in AI-related labor. Consistent with this finding, once we split the sample into pre-2019 and post-2019 periods, we find that firms indeed “walk the talk” prior to 2019, with AI talk significantly predicting subsequent AI walk, as shown in Table 5. However, in the more recent period, this relationship weakens considerably, suggesting a growing disconnect between firms’ AI-related claims and their actual follow-through.

5.3 Firm-Level Determinants of AI Washing

The vast majority (over 90%) of forward-looking AI investment talk specifies that investments will occur within the next two years. Therefore, if a firm makes such statements without subsequently building substantive AI workforce capacity within this time frame, we classify it as engaging in AI washing. Specifically, we define an *AI washing incident* as an instance where firm i discusses AI investment plans in quarter t but records zero AI walk from quarter $t + 1$ through quarter $t + 8$. A firm is labeled as an *AI Washing firm* if it has at least one AI washing incident over the sample period. Based on this definition, we identify 165 firms that have engaged in at least one AI washing incident between 2016Q1 and 2024Q2.

Panel A of Figure 6 shows that AI washing incidents have become increasingly frequent over the sample period. After fluctuating at relatively low levels between 2016 and 2018, the number of quarterly incidents began trending upward from 2019 onward, with a sharp acceleration after 2020. The peak occurs in 2022Q1, when nearly 30 AI washing episodes were identified in a single quarter. Panel B of Figure 6 highlights the industry distribution of

AI washing firms. Manufacturing dominates, accounting for more than half of all identified firms, followed by Finance & Insurance and Real Estate. Other industries with notable representation include Retail Trade, Administrative Support, and Wholesale. Incidents are much less frequent in sectors such as Agriculture, Arts & Entertainment, and Construction, suggesting that AI washing is concentrated in industries where AI narratives are both more marketable and more plausibly integrated into firm operations.

We next investigate which ex-ante firm characteristics predict AI washing. Table 6 reports results from estimating the following specification:

$$AI\ Washing_i = \alpha + \beta Firm\ Variable_{i,2016Q1} + \lambda_{industry} + \varepsilon_i, \quad (5)$$

where $AI\ Washing_i$ represents whether or not firm i is identified as an AI washing firm. All regressions control for industry fixed effects $\lambda_{industry}$. $Firm\ Variable_{i,2016Q1}$ denotes one of the ex-ante firm characteristics as of 2016Q1, including log of sales in column 1, *Cash*) cash/assets in column 2, leverage in column 3, age in column 4, R&D in column 5, log of CAPEX in column 6, Tobin’s Q in column 7, and ROA in column 8. A multivariate specification with all variables are shown in column 10.

[Insert Table 6 here]

The results show that log sales is a strong and negative predictor of AI washing: larger firms are significantly less likely to engage in AI washing. Higher R&D intensity and greater CAPEX are also associated with a lower likelihood of AI washing, consistent with these variables capturing genuine investment capacity and innovation orientation. By contrast, other variables, including cash holdings, leverage, firm age, Tobin’s Q, and ROA, do not consistently predict AI washing behavior.

Taken together, these results suggest that smaller firms and those with lower tangible and intangible investment levels are more prone to AI washing. This aligns with the view that AI washing may be used disproportionately by resource-constrained firms as a low-cost

signaling strategy, while better-capitalized firms are more likely to deliver on their stated AI ambitions.

6 Market Reactions to AI Washing

6.1 Short-run and Long-run Stock Returns

We explore market reactions to AI talk and walk by analyzing both short-run and long-run abnormal returns. Table 7 presents the results from estimating the following specification:

$$Abnormal\ Return_{i,t} = \alpha + \beta_1 AI\ Talk_{i,t} + \beta_2 AI\ Walk_{i,t-1} + \gamma \mathbf{X}_{i,t-1} + \alpha_{firm} + \lambda_{industry-quarter} + \varepsilon_{i,t}, \quad (6)$$

where $Abnormal\ Return_{i,t}$ represents either market-adjusted CARs, measured over the three-day window surrounding earnings calls, or market-adjusted BHARs, measured over the subsequent 180 days, 270 days, or 360 days from the earnings call dates for firm i in quarter t . We control for a comprehensive set of lagged firm variables (size, cash/assets, R&D, age, CAPEX, earnings surprise, ROA, market-to-book ratio), along with firm fixed effects α_{firm} and industry-quarter fixed effects $\lambda_{industry-quarter}$.

[Insert Table 7 here]

In the short run, higher AI talk is associated a positive market response. In Columns (1)–(2), the coefficient on Talk is 0.196 and 0.212, respectively, both statistically significant at the 1% level, implying that a one-standard-deviation increase in AI-related communication is associated with approximately 0.2 basis point higher CARs in the three-day event window, controlling for other firm characteristics and fixed effects. Translating this into raw units, a back-of-the-envelope calculation suggests that mentioning three additional AI-related keywords in an earnings call corresponds to roughly 1 basis point higher CAR, holding all else constant. By contrast, AI walk (lagged one quarter) is not related to CARs in a significant

way, implying that in the short term, market rewards talk rather than walk, consistent with the idea that investor attention is disproportionately captured by rhetorical signaling over substantive capability.

Over longer horizons, this relationship changes significantly. In the 180-day BHAR regressions (Columns (3)-(4)), the coefficients on Talk remain positive (0.165 and 0.308) but are statistically insignificant, indicating that the initial market enthusiasm does not translate into consistently higher medium-term returns. By the 270-day horizon (Columns (5)-(6)), the coefficients on Talk turn negative and remain insignificant, before becoming both economically large and statistically significant in the 360-day BHAR regressions (-2.520 and -1.704 in Columns (7)-(8), significant at the 1% and 5% levels, respectively). This reversal suggests that firms relying on AI talk without corresponding action may face a market correction within a year.

In contrast, while the market reaction to firms' AI walk is muted in the short run, it becomes strongly and consistently positive across all long-horizon specifications. For example, the coefficient on Walk is 14.5 in the 180-day window (Column 4), rising to 24.8 in the 360-day window (Column 8), all significant at the 1% level, suggesting that a one-standard-deviation higher AI walk is associated with 25 bps higher BHAR over the 360-day horizon, all else being equal. These results indicate that substantive AI workforce investments generate sustained positive abnormal returns, in sharp contrast to the eventual under-performance following unsubstantiated AI talk.

Overall, the evidence points to a clear divergence in investor responses: the market reacts quickly and favorably to AI-related communications, but without real follow-through, this initial optimism fades and eventually reverses. By contrast, tangible AI investment yields slower but more durable and significant market rewards.

6.2 Institutional Ownership

Next, we want to understand whether and to what extent institutional investors can distinguish AI washing behavior. To answer this question, we estimate the following regression:

$$Y_{i,t} = \alpha + \beta_1 \text{Talk}_{i,t-k} + \beta_2 \text{Walk}_{i,t-k} + \gamma \mathbf{X}_{i,t-1} + \alpha_{\text{firm}} + \lambda_{\text{industry-quarter}} + \varepsilon_{i,t}, \quad (7)$$

where $Y_{i,t}$ is the number of funds holding a firm’s stock. In separate regressions, we construct the outcome variable for AI-focused mutual funds and ETFs and all all institutional funds, respectively. The vector $\mathbf{X}_{i,t-1}$ includes a set of lagged control variables capturing firm characteristics and financial indicators. The specification also includes firm fixed effects, α_{firm} , and industry-quarter fixed effects, $\lambda_{\text{industry-quarter}}$, to account for unobserved heterogeneity across firms and industry-specific shocks over time. We estimate this equation separately for four different lag structures of AI talk and AI walk, where $k \in \{1, 2, 3, 4\}$ quarters, thereby examining how the relationship between AI engagement and institutional ownership evolves over time. In each regression, only one lag of talk and walk is included, allowing us to isolate the timing of the relationship.

[Insert Table 8 here]

Panel A of Table 8 focuses on AI-focused mutual funds and ETFs. Columns (1)–(4) show a sharp split between walk and talk. Across all four lags, AI walk is strongly positive and significant: a one-standard-deviation increase in walk is associated with roughly 0.24–0.30 additional AI-focused fund holding the stock. By contrast, AI talk shows no association with AI-focused fund holdings at $t-1$ or $t-2$, but turns negative and statistically significant with longer lags. Taken together, AI-focused investors appear to reward substantive AI investment and within 3–4 quarters, actively reduce holdings in firms whose AI communication is not backed by real investment.

The pattern generalizes to the broader fund universe, but with larger magnitudes and slower adjustment. As shown in Panel B of Table 8, AI walk predicts meaningfully higher

institutional ownership at all lags, with coefficients of 12.6–14.3, all statistically significant at the 5% level. In contrast, AI talk is statistically insignificant at every lag, indicating that generalist funds largely ignore narrative emphasis but instead, focus on firms’ substantive investment.

Compared to the overall market, which reacts strongly to AI talk in the short run before eventually penalizing unsubstantiated claims, institutional investors, especially AI-focused ones, see through AI washing both more quickly and more decisively. They reward walk immediately and begin to punish talk within three quarters, while generalist funds respond more slowly and only to concrete AI investment. This suggests that specialized institutional investors are better equipped to evaluate the credibility of AI-related claims and adjust portfolios accordingly, making them an early and sharper detector of AI washing than the market as a whole.

7 Why Do Firms AI Wash?

7.1 Managerial Incentives and ChatGPT Release

To investigate managerial incentives behind AI washing, we explore whether managers with higher pay-for-performance sensitivity are more likely to inflate AI talk in response to the launch of ChatGPT. We focus on delta, a widely used measure of executive compensation sensitivity, which captures how much the dollar value of a CEO’s wealth changes in response to a 1% increase in the firm’s stock price (Coles et al. (2006), Core and Guay (2002)). High-delta executives have strong incentives to influence short-term stock prices and may thus be more prone to using strategic disclosures, such as emphasizing AI capabilities, to shape investor perceptions.

We implement a difference-in-differences (DID) approach that compares changes in AI talk and AI walk, respectively, before and after the ChatGPT launch between firms with high-delta managers and those with low-delta managers. Specifically, we estimate the following

regression:

$$Y_{i,t} = \alpha + \beta_1 Post_t \times HighDelta_i + \gamma \mathbf{X}_{i,t-1} + \alpha_{firm} + \lambda_{industry-quarter} + \varepsilon_{i,t}, \quad (8)$$

where $Y_{i,t}$ is the outcome variable for firm i in quarter t , and can be either AI talk or AI walk. The variable $Post_t$ is an indicator for the post-ChatGPT period, and $HighDelta_i$ is a dummy variable equal to one for firms in the top quintile of CEO delta and zero for firms in the bottom quintile. The interaction term $Post_t \times HighDelta_i$, hence, identifies whether high-delta firms exhibit differential changes in behavior after ChatGPT’s release. The model also includes a vector of lagged firm-level control variables $\mathbf{X}_{i,t-1}$, as well as firm fixed effects, α_{firm} , and industry-quarter fixed effects, $\lambda_{industry-quarter}$, to account for unobserved heterogeneity across firms and industry-specific shocks over time.

[Insert Table 9 here]

The results, presented in Table 9, reveal a striking pattern. Columns (1) and (2) show that high-delta firms significantly increase their AI talk in the post-ChatGPT period, with coefficients of 0.946 and 0.878, respectively, both statistically significant at the 1% level. In contrast, columns (3) and (4) show only modest or statistically insignificant increases in actual AI walk among high-delta firms (coefficients of 0.070 and 0.035), indicating that the increase in AI talk among high-delta firms is not matched by a corresponding increase in AI implementation.

[Insert Figure 7 here]

The key identifying assumption in our DID model is that the high- and low-delta firms would have exhibited similar behavior in the absence of the ChatGPT launch. To illustrate parallel trends in the pre-event period, we regress talk and walk, respectively, on yearly indicator variables, interacted with the *HighDelta* indicator, and plot the coefficients along with their 95% confidence intervals in Figure 7. The one quarter prior to the ChatGPT

launch is set as the baseline period. Consistent with our main results, we observe positive coefficients immediately after the event, with high statistical significance starting 2 quarters after. Noticeably, none of the pre-event coefficients is significant, suggesting that the control and treated firms do not exhibit any meaningful differences prior to the ChatGPT launch.

Overall, this evidence suggests that managerial incentives play a key role in driving gaps between what firms say and what they do regarding AI investment. When financial gains can be realized through talk alone, and especially when executives stand to benefit directly from stock price increase, firms may have stronger motives to engage in strategic disclosure.

7.2 Managerial Myopia and Seasoned Equity Offerings

Seasoned equity offerings (SEOs) create strong incentives for managers to influence investor perceptions in the short run, as favorable market sentiment can facilitate higher offering prices and more successful capital raising. If AI talk spikes in the SEO quarter without a corresponding increase in AI walk, this likely reflects opportunistic communication rather than substantive investment.

We examine this hypothesis using SEO data from Refinitiv’s SDC Platinum. We retain only the first SEO for each firm, provided it is not followed by a second SEO within two years, and restrict the sample to SEOs occurring within one week of the quarterly earnings conference calls. This ensures that our measure of AI talk reflects information managers disclose in the same quarter as the capital-raising decision. After applying these filters, our sample consists of 92 SEO events.

An event study regression reveals a striking divergence between talk and walk patterns. Figure 8 plots the estimated coefficients from regressions of AI talk and AI walk on relative-quarter indicators, with the SEO quarter normalized to zero. AI Talk exhibits a pronounced spike in the SEO quarter, indicating that managers substantially increase AI-related communication when approaching an equity issuance. In contrast, AI Walk remains essentially flat throughout the event window, with no discernible increase in substantive AI workforce

investment even up to four quarters after the SEO. This divergence between rhetoric and action raises the possibility that managers strategically deploy AI-related communication to influence short-term investor sentiment rather than to signal genuine technological commitment.

We test this directly by estimating CAR regressions with an interaction between AI talk and an SEO-quarter indicator. The results in Table 10 show that AI talk in non-SEO periods is associated with modest but significant positive CARs — about 0.2 basis point per one-standard-deviation increase in talk. SEOs themselves are met with slightly negative market reactions (-2.46 to -2.06 bps), consistent with prior literature that equity issuance can be interpreted as a negative signal (e.g., Loughran and Ritter (1995); Loughran and Ritter (1997); Baker and Wurgler (2000); Baker and Wurgler (2002); Dittmar et al. (2020)).

Noticeably, the interaction term is both economically and statistically significant — in SEO quarters, the market reward to AI talk is amplified by roughly 1.4 bps. This implies that a one-standard-deviation increase in AI talk during an SEO is associated with about 1.6 bps higher CARs, over seven times the effect in non-SEO periods. Coupled with the absence of any contemporaneous rise in AI talk, this finding supports the view that AI talk during SEOs reflects short-term narrative management, allowing managers to temporarily boost market sentiment at the time of equity issuance.

Taken together, the event-study and CAR results are consistent with managerial myopia: managers appear to strategically amplify AI discussions in the short term to influence investor sentiment during capital raising events, even when substantive AI investment does not follow through. This opportunistic timing underscores the importance of distinguishing between rhetorical and substantive AI engagement, particularly when firms face heightened incentives to manage market perceptions.

8 Conclusion

This paper provides systematic evidence on AI washing — the strategic overstatement of AI engagement in public communications without matching real investment. Using LLMs to analyze earnings call transcripts and detailed resume data, we construct novel firm-quarter level measures of AI talk and AI walk. These measures separately capture rhetorical emphasis and substantive investment, allowing us to examine their distinct roles in firms’ AI engagement. We first validate the measures by linking them to future AI innovation. AI walk is a strong predictor of subsequent AI-related patent volume, value, and quality, whereas AI talk shows no such association. This confirms that our measures successfully distinguish between rhetoric and substance.

We then document a pronounced within-firm gap between past talk and future walk: increases in AI talk are not systematically followed by corresponding growth in AI-related employment, consistent with firms overstating AI engagement in their communications. This gap has widened over time, and is concentrated in manufacturing industry and more prevalent among smaller and less capital-intensive firms.

Turning to market reactions, we find that the short-term stock price response to AI talk is positive and significant, but dissipates over longer horizons and reverses within a year. In contrast, AI walk corresponds to muted short-term reactions but sustained and economically meaningful abnormal returns over the long run. Institutional ownership patterns reinforce this divergence. AI-focused mutual funds and ETFs allocate more capital to high-walk firms and begin scaling back exposure to high-talk firms within three quarters. Generalist funds respond more slowly but ultimately exhibit the same preference for substantive investment over rhetoric.

Finally, we show that managerial myopia serves as an incentive for AI washing. Managers with high pay-for-performance sensitivity and those conducting seasoned equity offerings significantly increase AI talk without commensurate increases in walk. During SEO quarters, the short-term return premium to AI talk is more than seven times larger than in other

periods, consistent with opportunistic narrative management aimed at influencing issuance valuations.

Together, these findings provide novel evidence that AI washing is both measurable and strategically motivated. The increasing prevalence of this behavior, particularly among smaller firms and in settings with strong short-term market incentives, highlights the need for investors and regulators to distinguish between rhetorical signaling and genuine technological commitment in corporate disclosures.

References

- Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, 40(S1):S293–S340.
- Acemoglu, D. and Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6):1488–1542.
- Agrawal, A., Gans, J. S., and Goldfarb, A. (2019). Artificial intelligence: the ambiguous labor market impact of automating prediction. *Journal of Economic Perspectives*, 33(2):31–50.
- Babina, T., Fedyk, A., He, A., and Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, 151:103745.
- Babina, T., Fedyk, A., He, A. X., and Hodson, J. (2023a). Artificial intelligence and firms’ systematic risk. *Available at SSRN*.
- Babina, T., Fedyk, A., He, A. X., and Hodson, J. (2023b). Firm investments in artificial intelligence technologies and changes in workforce composition. *NBER Volume on Technology, Productivity, and Economic Growth, forthcoming*.
- Baker, M. and Wurgler, J. (2000). The equity share in new issues and aggregate stock returns. *the Journal of Finance*, 55(5):2219–2257.
- Baker, M. and Wurgler, J. (2002). Market timing and capital structure. *The journal of finance*, 57(1):1–32.
- Bandyopadhyay, A., Mai, D., and Pukthuanthong, K. (2023a). Ai narrative and stock mispricing. *Available at SSRN*.
- Bandyopadhyay, A., Mai, D., and Pukthuanthong, K. (2023b). Diversity narrative and equity in firm leadership. *Available at SSRN*.
- Barrios, J. M., Hochberg, Y. V., and Moshary, S. (2025). Signals or smoke? the boom in corporate ai disclosures. *Available at SSRN 5133107*.
- Bauer, M., Huber, D., Offner, E., Renkel, M., and Wilms, O. (2024). Corporate green pledges. Technical report, IMFS Working Paper Series.
- Bybee, L., Kelly, B., Manela, A., and Xiu, D. (2024). Business news and business cycles. *The Journal of Finance, forthcoming*.
- Chen, S. (2022). Green investors and green transition efforts: talk the talk or walk the walk? *Available at SSRN 4254894*.
- Coles, J. L., Daniel, N. D., and Naveen, L. (2006). Managerial incentives and risk-taking. *Journal of financial Economics*, 79(2):431–468.

- Core, J. and Guay, W. (2002). Estimating the value of employee stock option portfolios and their sensitivities to price and volatility. *Journal of Accounting research*, 40(3):613–630.
- Dittmar, A., Duchin, R., and Zhang, S. (2020). The timing and consequences of seasoned equity offerings: A regression discontinuity approach. *Journal of Financial Economics*, 138(1):254–276.
- Eisfeldt, A. L., Schubert, G., and Zhang, M. B. (2023). Generative ai and firm values. Technical report, National Bureau of Economic Research.
- Flugum, R. and Souther, M. E. (2023). Stakeholder value: A convenient excuse for underperforming managers? *Journal of Financial and Quantitative Analysis*, pages 1–34.
- Gow, I. D., Larcker, D. F., and Zakolyukina, A. A. (2021). Non-answers during conference calls. *Journal of Accounting Research*, 59(4):1349–1384.
- Hampole, M., Papanikolaou, D., Schmidt, L. D., and Seegmiller, B. (2025). Artificial intelligence and the labor market. Technical report, National Bureau of Economic Research.
- He, Q., Marshall, B. R., Nguyen, J. H., Nguyen, N. H., Qiu, B., and Visaltanachoti, N. (2024). Greenwashing: Measurement and implications. *Available at SSRN*.
- Hirshleifer, D., Mai, D., and Pukthuanthong, K. (2023). War discourse and the cross section of expected stock returns. Technical report, National Bureau of Economic Research.
- Hirshleifer, D., Mai, D., and Pukthuanthong, K. (2025). War discourse and disaster premium: 160 years of evidence from the stock market. *The Review of Financial Studies*, 38(2):457–506.
- Houston, J. F., Kim, S., and Li, B. (2024). One hundred and thirty years of corporate responsibility.
- Jha, M., Liu, H., and Manela, A. (2022). Does finance benefit society? a language embedding approach. *A Language Embedding Approach (June 1, 2022)*.
- Jha, M., Qian, J., Weber, M., and Yang, B. (2024). Chatgpt and corporate policies. Technical report, National Bureau of Economic Research.
- Jia, N., Li, N., Ma, G., and Xu, D. (2025). Corporate responses to generative ai: Early evidence from conference calls. *Review of Accounting Studies*, *forthcoming*.
- Kogan, L., Papanikolaou, D., Seru, A., and Stoffman, N. (2017). Technological innovation, resource allocation, and growth. *The Quarterly Journal of Economics*, 132(2):665–712.
- Larcker, D. F. and Zakolyukina, A. A. (2012). Detecting deceptive discussions in conference calls. *Journal of Accounting Research*, 50(2):495–540.
- Li, Q., Shan, H., Tang, Y., and Yao, V. (2024). Corporate climate risk: Measurements and responses. *The Review of Financial Studies*, 37(6):1778–1830.

- Liu, J. (2025). Ai exposure without labor data: Measuring ai’s impact through firms. *Available at SSRN 3482150*.
- Lopez-Lira, A. and Tang, Y. (2023). Can chatgpt forecast stock price movements? return predictability and large language models. *arXiv preprint arXiv:2304.07619*.
- Loughran, T. and Ritter, J. R. (1995). The new issues puzzle. *The Journal of finance*, 50(1):23–51.
- Loughran, T. and Ritter, J. R. (1997). The operating performance of firms conducting seasoned equity offerings. *The journal of finance*, 52(5):1823–1850.
- Pairolero, N. A., Giczy, A. V., Torres, G., Islam Erana, T., Finlayson, M. A., and Toole, A. A. (2025). The artificial intelligence patent dataset (aipd) 2023 update. *The Journal of Technology Transfer*, pages 1–24.
- Sautner, Z., Van Lent, L., Vilkov, G., and Zhang, R. (2023). Firm-level climate change exposure. *The Journal of Finance*, 78(3):1449–1498.
- Stein, J. C. (1988). Takeover threats and managerial myopia. *Journal of Political Economy*, 96(1):61–80.
- van Binsbergen, J. H., Bryzgalova, S., Mukhopadhyay, M., and Sharma, V. (2024). (almost) 200 years of news-based economic sentiment. Technical report, National Bureau of Economic Research.
- Webb, M. (2019). The impact of artificial intelligence on the labor market. *Available at SSRN 3482150*.

Figure 1. Percent of U.S. Public Firms Talking about AI

This figure plots the percent of U.S. public firms that mention “AI” in their quarterly earnings conference calls from 2016Q1 to 2024Q2.

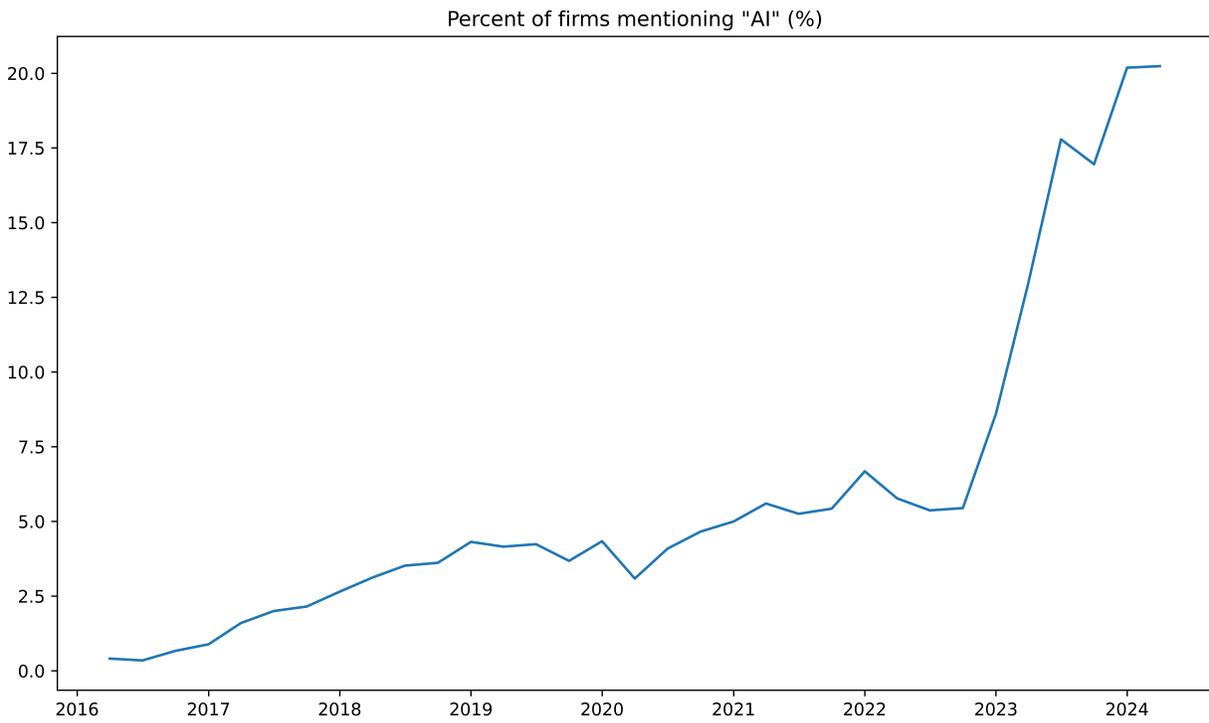


Figure 3. Time Series of AI Talk and AI Walk

This figure plots the quarterly average AI Talk (Panel A) and AI Walk (Panel B) as in percentage terms from 2016Q1 to 2024Q2.

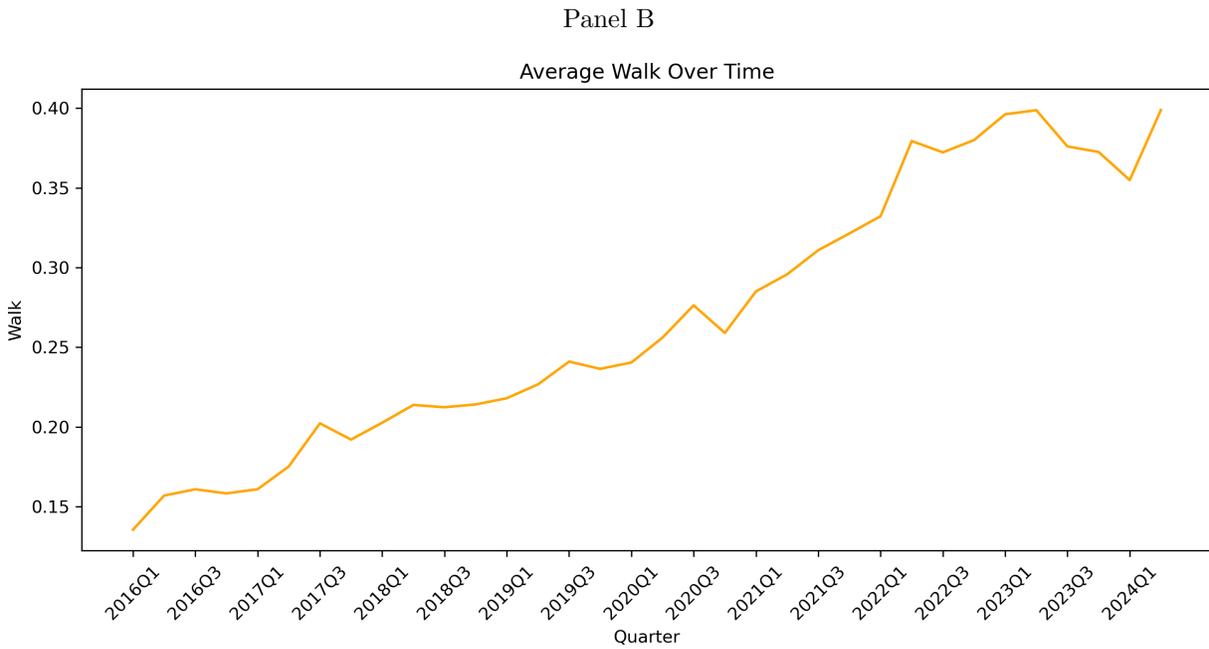
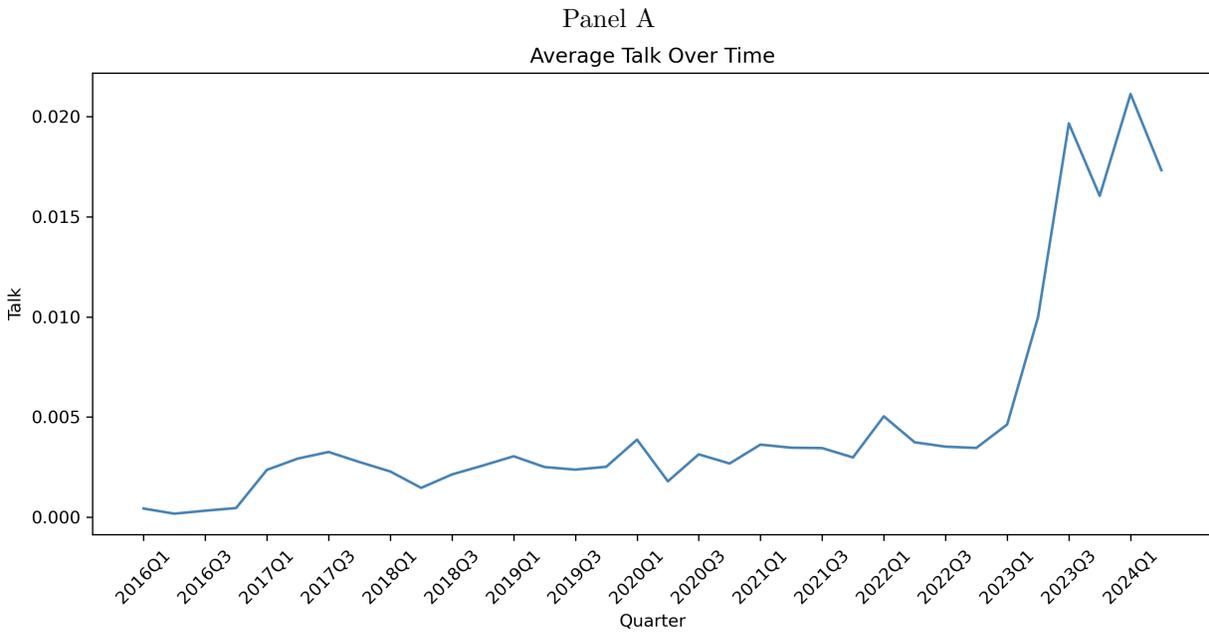
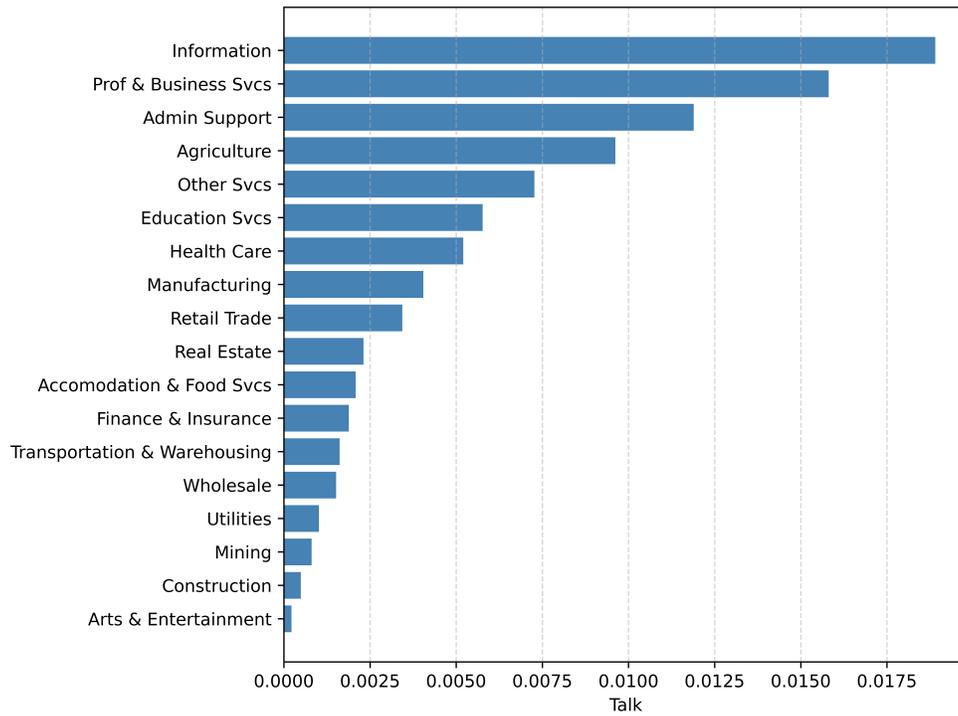


Figure 4. AI Talk and AI Walk by Industry

This figure plots the average AI Talk (Panel A) and AI Walk (Panel B) by industry, defined using firms' two-digit NAICS codes, from 2016Q1 to 2024Q2.

Panel A: AI Talk by Industry



Panel B: AI Walk by Industry

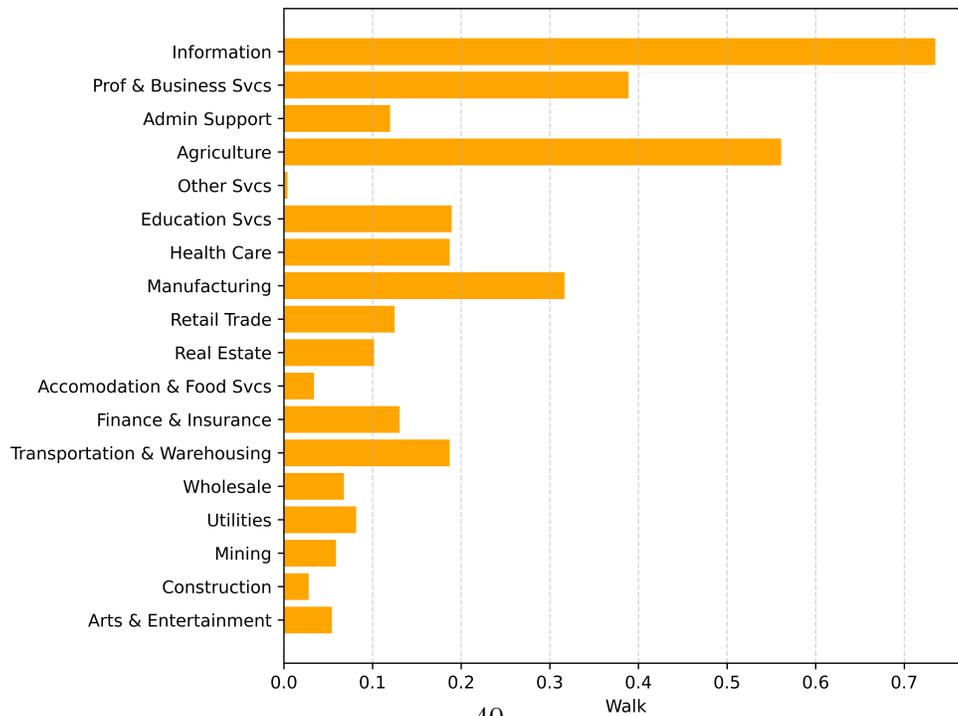


Figure 5. Time Series of Topics in AI Talk

This figure illustrates the time-series evolution of selected AI investment topics extracted from firms' earnings call discussions, 2016Q1 – 2024Q2. The y-axis reports the percentage of conference call transcripts at each quarter in the sample that cover each selected topic – AI, automation, generative AI, and cloud.

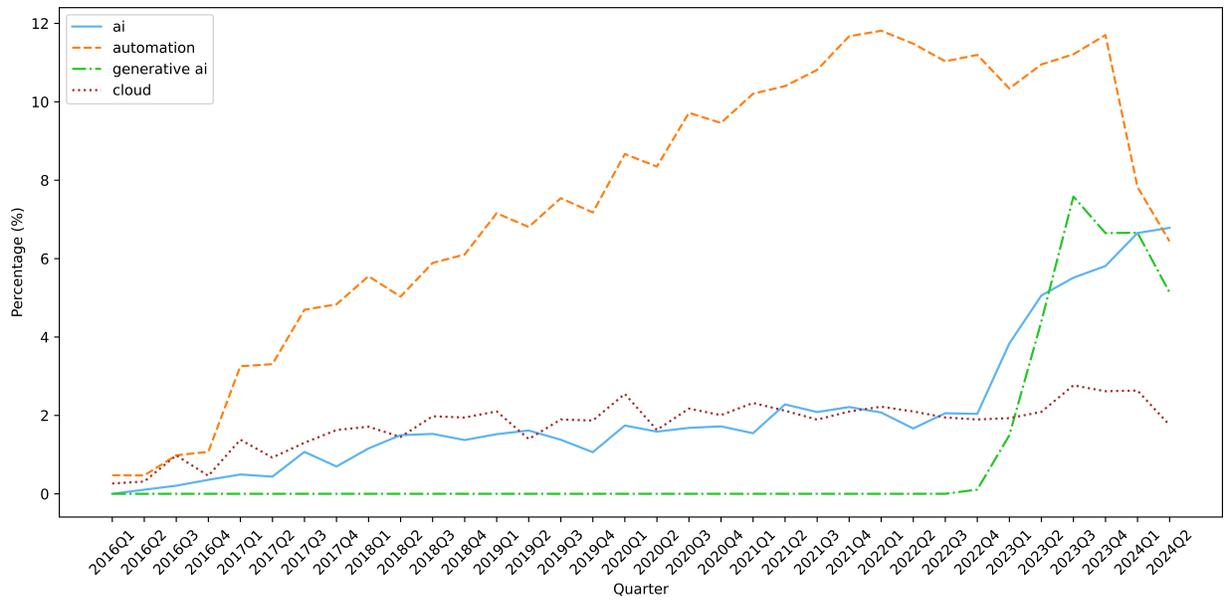
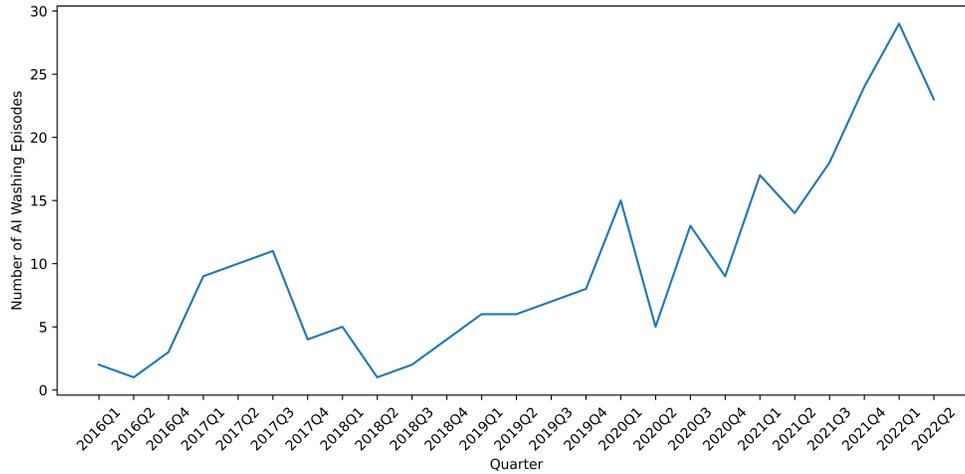


Figure 6. AI Washing

This figure presents descriptive patterns of firms' AI washing behavior, as defined by having AI talk in a given quarter with zero AI walk in the subsequent eight quarters. Panel A plots the time series of number of AI washing incidents from 2016Q1 to 2022Q2. Panel B plots the industry distribution of the number of firms that experience at least one AI washing incident over the sample period.

Panel A: AI Washing Time Series



Panel B: AI Washer by Industry

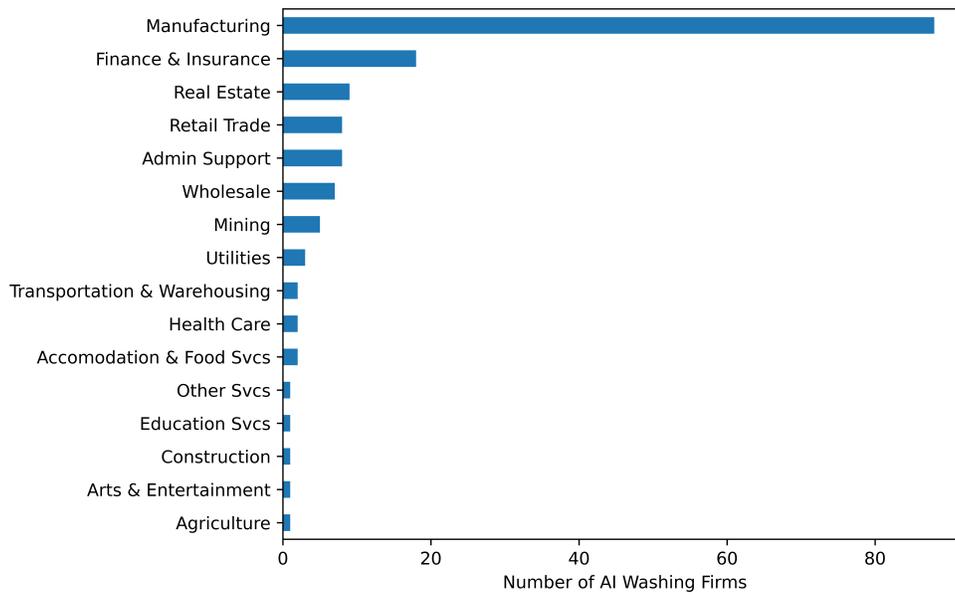


Figure 7. Pre- and Post-Event Differences Between High- and Low-Delta Firms

This figure presents results from the event-study regression using the release of ChatGPT in 2022Q4 as an exogenous shock. The estimated coefficients from past 4 quarters to future 4 quarters are plotted, along with their 95% confidence intervals. The quarter prior to ChatGPT release is omitted. Standard errors are clustered at the firm level.

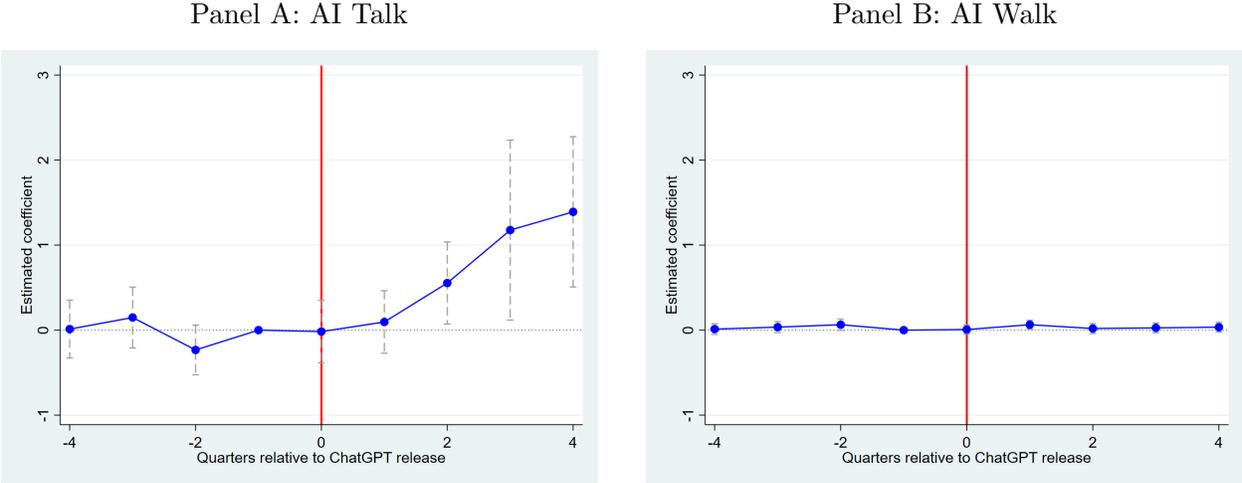
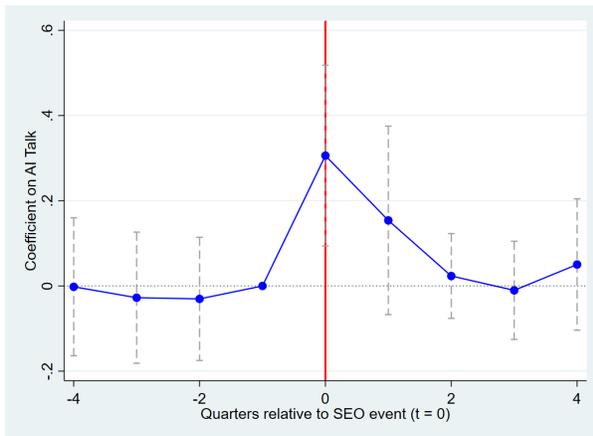


Figure 8. Event Study Around Seasoned Equity Offerings

This figure presents results from the event-study regression using firms' seasoned equity offering as events. The estimated coefficients from past 4 quarters to future 4 quarters are plotted, along with their 95% confidence intervals. The quarter prior to each SEO is omitted. Standard errors are clustered at the firm level.

Panel A: AI Talk



Panel B: AI Walk

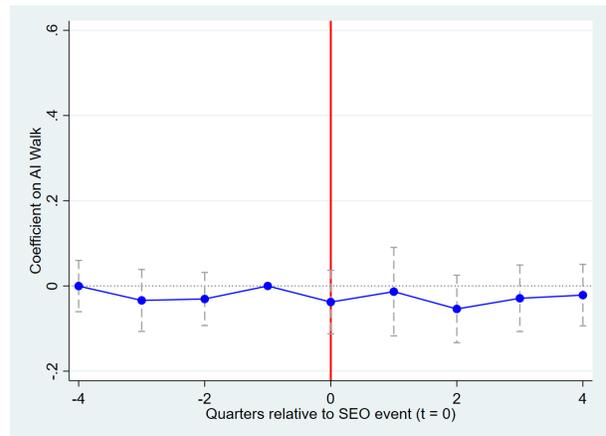


Table 1. Summary Statistics

This table presents summary statistics for the variables used in the analysis. It includes the mean, standard deviation, selected percentiles (5th, 25th, 50th, 75th, and 95th), and the number of observations (N).

Variable	Mean	Std. Dev.	p5	p25	p50	p75	p95	N
AI Talk	0.0089	0.0511	0.0000	0.0000	0.0000	0.0000	0.0483	20,135
AI Walk	0.3942	1.2281	0.0000	0.0000	0.0419	0.3072	1.7879	20,135
Number of AI Patents	2.0379	13.9895	0.0000	0.0000	0.0000	0.0000	6.0000	20,135
Econ. Value of Patents (\$ mil)	82.0553	853.5465	0.0000	0.0000	0.0000	0.0000	157.1905	20,135
AI Patent Citations	3.6650	53.9750	0.0000	0.0000	0.0000	0.0000	3.0000	20,135
Number of AI Funds	0.8306	1.8481	0.0000	0.0000	0.0000	1.0000	5.0000	20,135
ROA	-0.0009	0.0518	-0.0924	-0.0028	0.0083	0.0207	0.0500	20,127
Leverage	0.2370	0.2088	0.0004	0.0775	0.1823	0.3459	0.6715	18,917
log Sales	6.2579	2.0416	2.5551	5.0142	6.4123	7.6348	9.6229	20,127
Sales Growth	0.0432	0.2633	-0.2535	-0.0435	0.0195	0.0907	0.3709	19,912
Return	0.0340	0.2533	-0.3461	-0.1013	0.0215	0.1427	0.4303	20,046
MTB	1.8194	2.0269	0.1541	0.5624	1.1313	2.2160	6.1115	20,128
R&D	0.0116	0.0240	0.0000	0.0000	0.0000	0.0149	0.0548	20,128
Earnings Surprise (SUE)	0.2153	2.6805	-2.2379	-0.2108	0.0548	0.4283	3.2017	17,837
log CAPEX	3.6061	2.1037	0.2086	2.0380	3.5272	5.0689	7.3447	20,135

Table 2. AI Talk, Walk, and Innovation Outcomes

This table presents results from firm-quarter panel regressions analyzing the relationship between AI talk (lagged), walk (lagged), and firm innovation outcomes. The dependent variable is log of number of AI patents, log of AI patent value, and log of AI patent citations, respectively. Control variables include size, cash/assets, R&D, age, and capital expenditures, and are lagged by one period. We also control for firm fixed effects and industry-quarter fixed effects. Standard errors are clustered at the firm level and are shown in parentheses. (* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

Variable	Log AI Patent Count _t		Log AI Patent Value _t		Log AI Patent Citations _t	
	(1)	(2)	(3)	(4)	(5)	(6)
Talk _{t-1}	0.005 (0.005)	0.004 (0.005)	0.006 (0.012)	0.005 (0.012)	-0.009 (0.012)	-0.009 (0.012)
Walk _{t-1}	0.154** (0.065)	0.171*** (0.065)	0.208** (0.084)	0.246*** (0.082)	0.324*** (0.088)	0.334*** (0.088)
Controls	N	Y	N	Y	N	Y
Firm FE	Y	Y	Y	Y	Y	Y
Industry-Quarter FE	Y	Y	Y	Y	Y	Y
Adj R ²	0.858	0.859	0.789	0.789	0.630	0.631
Observations	16,080	16,031	16,080	16,031	16,080	16,031

Table 3. Firm Characteristics by Talk and Walk

This table presents univariate analysis results along several firm characteristics as of 2016. The means, medians, standard deviations, and number of observations of these variables are shown separately for firms that ever (never) engage in AI talk (Panel A) and firms that ever (never) engage in AI walk (Panel B) over our sample period from 2016Q1 to 2024Q2. The differences in means between the two groups of firms and the corresponding t-statistics are also reported.

Panel A	Ever Talk				Never Talk				Mean Diff	t-stat
	Variable	Mean	Median	SD	N	Mean	Median	SD		
log Sales ₂₀₁₆	6.077	6.303	2.055	539	5.427	5.540	1.981	957	0.650	6.015
Cash/Assets ₂₀₁₆	0.190	0.119	0.201	539	0.145	0.062	0.211	957	0.045	4.059
Age ₂₀₁₆	26.916	21.000	21.611	537	25.251	21.000	20.441	956	1.665	1.480
R&D ₂₀₁₆	0.018	0.000	0.038	539	0.012	0.000	0.039	957	0.005	2.426
log CAPEX ₂₀₁₆	3.013	2.936	1.957	539	2.429	2.270	1.926	957	0.584	5.599
ROA ₂₀₁₆	-0.004	0.008	0.057	539	-0.005	0.005	0.053	957	0.001	0.357

Panel B	Ever Walk				Never Walk				Mean Diff	t-stat
	Variable	Mean	Median	SD	N	Mean	Median	SD		
log Sales ₂₀₁₆	6.366	6.596	2.034	759	4.936	5.140	1.755	737	1.430	14.542
Cash/Assets ₂₀₁₆	0.175	0.098	0.202	759	0.146	0.063	0.214	737	0.028	2.650
Age ₂₀₁₆	28.898	22.000	23.709	758	22.706	21.000	16.934	735	6.192	5.792
R&D ₂₀₁₆	0.019	0.000	0.044	759	0.010	0.000	0.033	737	0.008	4.224
log CAPEX ₂₀₁₆	3.307	3.295	1.980	759	1.952	1.651	1.675	737	1.355	14.275
ROA ₂₀₁₆	-0.001	0.008	0.049	759	-0.010	0.004	0.058	737	0.009	3.214

Table 4. AI Talk and AI Walk

This table presents results from distributed-lag regressions analyzing the relationship between AI walk and AI talk lagged by various horizons. Control variables include size, cash/assets, R&D, age, and capital expenditures, and are lagged by one period. We control for industry fixed effects in column (1), industry fixed effects and quarter fixed effects in column (2), industry-quarter fixed effects in column (3), firm fixed effects in column (4), firm fixed effects and quarter fixed effects in column (5), and industry-quarter fixed effects in column (6). Standard errors are clustered at the firm level (***) $p < 0.01$, ** $p < 0.05$, * $p < 0.1$).

Variable	Dependent Variable: AI Walk _t					
	(1)	(2)	(3)	(4)	(5)	(6)
Talk _{t-1}	0.0310 (0.0189)	0.0327 (0.0198)	0.0340 (0.0210)	0.0040 (0.0032)	0.0045 (0.0032)	0.0047 (0.0033)
Talk _{t-2}	0.0454** (0.0189)	0.0475** (0.0197)	0.0487** (0.0201)	0.0021 (0.0025)	0.0023 (0.0025)	0.0023 (0.0026)
Talk _{t-3}	0.0541** (0.0244)	0.0552** (0.0251)	0.0572** (0.0263)	-0.0024 (0.0042)	-0.0020 (0.0041)	-0.0019 (0.0043)
Talk _{t-4}	0.1113** (0.0467)	0.1115** (0.0467)	0.1127** (0.0479)	-0.0020 (0.0072)	-0.0020 (0.0072)	-0.0022 (0.0076)
Talk _{t-5}	0.1059** (0.0450)	0.1045** (0.0447)	0.1039** (0.0455)	-0.0032 (0.0063)	-0.0034 (0.0063)	-0.0040 (0.0065)
Talk _{t-6}	0.1034** (0.0428)	0.1029** (0.0429)	0.1031** (0.0438)	-0.0013 (0.0061)	-0.0017 (0.0061)	-0.0024 (0.0064)
Talk _{t-7}	0.0991** (0.0421)	0.0990** (0.0426)	0.0992** (0.0435)	0.0035 (0.0058)	0.0035 (0.0058)	0.0031 (0.0059)
Talk _{t-8}	0.1145** (0.0503)	0.1156** (0.0511)	0.1159** (0.0525)	0.0031 (0.0050)	0.0029 (0.0049)	0.0025 (0.0051)
Controls	Y	Y	Y	Y	Y	Y
Firm FE	N	N	N	Y	Y	Y
Industry FE	Y	Y	N	N	N	N
Quarter FE	N	Y	N	N	Y	N
Industry-Quarter FE	N	N	Y	N	N	Y
Adj R ²	0.292	0.292	0.269	0.933	0.933	0.931
Observations	13,602	13,602	13,598	13,595	13,595	13,593

Table 5. AI Talk and AI Walk Pre- and Post-2019

This table presents results from distributed-lag regressions analyzing the relationship between AI walk and AI talk lagged by various horizons for pre-2019 and post-2019 subsamples, respectively. Control variables include size, cash/assets, R&D, age, and capital expenditures, and are lagged by one period. We also control for firm fixed effects and industry-quarter fixed effects. Standard errors are clustered at the firm level (***) $p < 0.01$, ** $p < 0.05$, * $p < 0.1$).

Variable	Dependent Variable: AI Walk _t	
	(1) Pre-2019	(2) Post-2019
Talk _{t-1}	0.0380* (0.0201)	0.0124 (0.0087)
Talk _{t-2}	0.0443** (0.0215)	0.0134 (0.0088)
Talk _{t-3}	0.0433** (0.0205)	-0.0018 (0.0081)
Talk _{t-4}	0.0238 (0.0173)	0.0059 (0.0086)
Talk _{t-5}	0.0177 (0.0141)	0.0072 (0.0099)
Talk _{t-6}	0.0617*** (0.0193)	0.0044 (0.0099)
Talk _{t-7}	0.0584*** (0.0184)	0.0089 (0.0093)
Talk _{t-8}	0.0417** (0.0165)	0.0021 (0.0094)
Controls	Y	Y
Firm FE	Y	Y
Industry-Quarter FE	Y	Y
Adj R ²	0.882	0.821
Observations	3,551	8,832

Table 6. Firm-level Determinants of AI Washing

This table presents results from regressions examining the firm-level determinants of AI washing. The dependent variable is an indicator of any AI washing incidents from 2016Q1 to 2024Q2. The independent variables are ex-ante firm characteristics as of 2016Q1, including log sales, cash/assets, leverage, age, R&D intensity, log capital expenditures, Tobin's Q, and return on assets (ROA). All specifications include industry fixed effects. Standard errors clustered at the industry level are reported in parentheses. (* p<0.1, ** p<0.05, *** p<0.01)

Variable	Ever AI Washing								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
log Sales ₂₀₁₆	-0.0529*** (0.0117)								-0.0365** (0.0169)
Cash/Assets ₂₀₁₆		-0.0342 (0.0438)							-0.1206 (0.0961)
Leverage ₂₀₁₆			0.0657 (0.0669)						0.0874 (0.1210)
Age ₂₀₁₆				-0.0013** (0.0006)					0.0007 (0.0008)
R&D ₂₀₁₆					-0.7692*** (0.0721)				-1.6342*** (0.4282)
log CAPEX ₂₀₁₆						-0.0623*** (0.0077)			-0.0513*** (0.0146)
Tobin's Q ₂₀₁₆							-0.0243* (0.0125)		-0.0236 (0.0195)
ROA ₂₀₁₆								-0.4099 (0.3380)	-0.0392 (0.4171)
Industry FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
Adj. R ²	0.0471	-0.00851	-0.0101	-0.00288	-0.00481	0.0585	-0.00139	-0.00596	0.0846
Observations	539	539	506	537	539	539	539	539	499

Table 7. AI Washing and Stock Market Reactions

This table presents results from regressions examining the relationship between AI talk, walk, and stock market reactions. The dependent variables include cumulative abnormal returns (CAR) and buy-and-hold abnormal returns (BHAR) over various horizons, calculated using the market-adjusted model. Control variables include ROA, sales, stock return, MTB, R&D, SUE, and capital expenditures. All specifications include firm fixed effects and industry-quarter fixed effects. Standard errors clustered at the firm level are reported in parentheses. (* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$)

Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	CAR(-1,1) _t	CAR(-1,1) _t	BHAR(0,180) _t	BHAR(0,180) _t	BHAR(0,270) _t	BHAR(0,270) _t	BHAR(0,360) _t	BHAR(0,360) _t
Talk _t	0.196*** (0.071)	0.212*** (0.077)	0.165 (0.368)	0.308 (0.382)	-0.586 (0.565)	-0.064 (0.602)	-2.520*** (0.746)	-1.704** (0.699)
Walk _{t-1}	-0.447 (0.386)	-0.314 (0.343)	14.783*** (3.245)	14.471*** (3.617)	25.014*** (6.225)	23.748*** (5.710)	28.740*** (9.849)	24.775*** (7.726)
Controls	N	Y	N	Y	N	Y	N	Y
Firm FE	Y	Y	Y	Y	Y	Y	Y	Y
Industry-Quarter FE	Y	Y	Y	Y	Y	Y	Y	Y
Adj. R ²	0.015	0.036	0.172	0.220	0.221	0.285	0.252	0.332
Observations	18,184	17,659	17,628	17,127	16,953	16,464	15,854	15,397

Table 8. AI Washing and Institutional Ownership

This table presents results from firm-quarter panel regressions examining the relationship between AI talk, walk, and institutional ownership. Dependent variable is the number of holdings funds by AI-focused (Panel A) and all institutional funds (Panel B). Control variables include size, cash/assets, R&D, age, and capital expenditures, and are lagged by one period. All specifications include firm fixed effects and industry-quarter fixed effects. Standard errors are clustered at the firm level and are shown in parentheses. (* p<0.1, ** p<0.05, *** p<0.01)

Panel A: AI Funds				
Variable	(1)	(2)	(3)	(4)
Number of Holding Funds _t				
Talk _{t-1}	-0.012 (0.011)			
Walk _{t-1}	0.302*** (0.092)			
Talk _{t-2}		-0.011 (0.014)		
Walk _{t-2}		0.265*** (0.083)		
Talk _{t-3}			-0.025** (0.013)	
Walk _{t-3}			0.266*** (0.082)	
Talk _{t-4}				-0.039*** (0.013)
Walk _{t-4}				0.243*** (0.080)
Controls	Y	Y	Y	Y
Firm FE	Y	Y	Y	Y
Industry-Quarter FE	Y	Y	Y	Y
Adj R ²	0.781	0.797	0.796	0.809
Observations	18,959	18,317	17,644	17,070
Panel B: All Funds				
Variable	(5)	(6)	(7)	(8)
Number of Holding Funds _t				
Talk _{t-1}	1.309 (0.912)			
Walk _{t-1}	14.304** (5.803)			
Talk _{t-2}		1.510 (1.141)		
Walk _{t-2}		12.610** (5.576)		
Talk _{t-3}			-0.080 (1.161)	
Walk _{t-3}			12.671** (5.876)	
Talk _{t-4}				-0.698 (1.314)
Walk _{t-4}				13.391** (6.046)
Controls	Y	Y	Y	Y
Firm FE	Y	Y	Y	Y
Industry-Quarter FE	Y	Y	Y	Y
Adj R ²	0.944	0.946	0.943	0.944
Observations	18,959	18,317	17,644	17,070

Table 9. Difference-in-Differences: Post-ChatGPT AI Talk-Walk Gap

This table presents difference-in-differences regressions estimating whether firms with high equity-based CEO incentives (high delta) responded differently to the ChatGPT launch in terms of AI talk, AI walk, and the resulting talk-walk gap. The main regressor of interest is $Post \times High\ Delta$, which captures the differential effect for high-delta firms in the post-ChatGPT period relative to low-delta firms. Control variables include size, cash/assets, R&D, age, and capital expenditures, and are lagged by one period. All specifications include firm fixed effects and industry-quarter fixed effects. Standard errors are clustered at the firm level. (* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

Variable	AI Talk _t		AI Walk _t	
	(1)	(2)	(3)	(4)
Post _t × High Delta	0.946*** (0.269)	0.878*** (0.263)	0.070** (0.031)	0.035 (0.028)
Controls	N	Y	N	Y
Firm FE	Y	Y	Y	Y
Industry-Quarter FE	Y	Y	Y	Y
Adj R ²	0.208	0.219	0.944	0.949
Observations	7,038	6,599	7,038	6,599

Table 10. Stock Market Reactions Around SEOs

This table presents results from regressions examining the relationship between AI talk, walk, and stock market reactions. The dependent variables include cumulative abnormal returns (CAR) and buy-and-hold abnormal returns (BHAR) over various horizons, calculated using the market-adjusted model. Control variables include ROA, sales, stock return, MTB, R&D, SUE, and capital expenditures. All specifications include firm fixed effects and industry-quarter fixed effects. Standard errors clustered at the firm level are reported in parentheses. (* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$)

Variable	CAR(-1,1) _t	
Talk _t	0.208*** (0.072)	0.223*** (0.078)
SEO Dummy _t	-2.455** (1.195)	-2.059* (1.238)
SEO Dummy × Talk _t	1.117 (0.735)	1.403* (0.758)
Walk _{t-1}	-0.400 (0.371)	-0.295 (0.334)
Controls	N	Y
Firm FE	Y	Y
Industry-Quarter FE	Y	Y
Adj. R^2	0.0151	0.0366
Observations	18,184	17,659

Appendix

A1. Examples of AI Talk and AI Walk

Below are a few examples of discussions about AI-related investment in earnings conference calls (“AI Talk”), followed by examples of AI-related job descriptions from employee resumes (“AI Walk”).

AI Talk Examples

1. “We will be developing features and product concepts that could extend MicroVision’s PicoP scanning technology into emerging market opportunities in both 3D sensing and augmented reality.”
2. “We are also investing in the development of two very important technologies for our future: the next-generation infotainment platform and the new autonomous driving platform, which will create new growth opportunities for the company.”
3. “The strong revenue performance also enabled us to invest more significantly in our product and technology, planting seeds in the areas of AI and machine learning that will provide the foundation of our future... Our technology innovation efforts are focused on emerging computing platform shifts such as AI, virtual and augmented reality, distributed commerce, and the Internet of Everywhere.”

AI Walk Examples

1. “Conversational Engineering using proprietary scripting language and software to build out the natural language processing model for Synchrony’s AI chatbot/digital assistant, Sydney. Analysis on large sets of user behaviour and linguistic data to identify issues and improvements for the NLP system.”
2. “Senior ML engineer leading AR effect recommendation across multiple apps. Skilled at model optimization and experimentation design to achieve significant metrics wins.”
3. “We are the science team building next-gen computer vision algorithms and services for AWS customers. My teams are engaged in cutting-edge research in computer vision in the areas of recognition, segmentation, video understanding, document OCR, and layout/semantic understanding. We translate our work into AWS services and platforms, such as Amazon Rekognition and Amazon Textract.”

A2. Prompt for AI-Fund Classification

Instruction: Analyze the following fund prospectus excerpt. Based on this text only, determine whether the fund is an AI-Investing Fund or an AI-Using Fund.

Classification Criteria:

- The fund qualifies as an AI-Investing Fund if it explicitly states that it invests in companies that develop, advance, or heavily utilize AI in their business models (e.g., AI research, AI infrastructure, machine learning applications).
- The fund does not qualify if it simply mentions using AI in its own investment strategies (e.g., “We use AI-driven algorithms for portfolio selection”).
- If the fund mentions both aspects, classify based on the primary focus described in the text.

Output Format:

1. Fund Classification: (AI-Investing Fund / AI-Using Fund / Unclear)
2. Primary Investment Focus: (Brief summary of the fund’s AI-related investment thesis)
3. Supporting Evidence: (Quote key phrases from the prospectus that justify the classification)

A3. Prompt for Job Title Classification

Instruction: Classify each of the following job titles as *AI-Related* or *Not AI-Related*.

Classification Criteria:

- **AI-Related:** The role contributes to the firm's in-house AI development or infrastructure, indicating real AI investment. This includes building or deploying machine learning models, researching AI algorithms, or developing tools and systems to support AI work.

Examples: Artificial Intelligence Engineer; Senior Data Scientist – Machine Learning Engineer; Lead Machine Learning Scientist – Enterprise Products; AI Consultant; AI Senior Analyst; Machine Learning Engineer; Technician Architecture Delivery Senior Analyst AI; Artificial Intelligence Analyst; Software Engineer; Machine Learning; Artificial Intelligence Architect

- **Not AI-Related:** The role does not directly involve AI model development or infrastructure, such as generic IT or computer science roles.

Output Format:

1. Title: (Job title to be classified)
2. Classification: (AI-Related / Not AI-Related)

A4. Additional Results

Figure A1. Percent of Job Positions without Description Text

This figure plots the percent of job positions without description text for AI-related titles and non-AI-related titles, respectively, from 2016Q1 to 2024Q2.

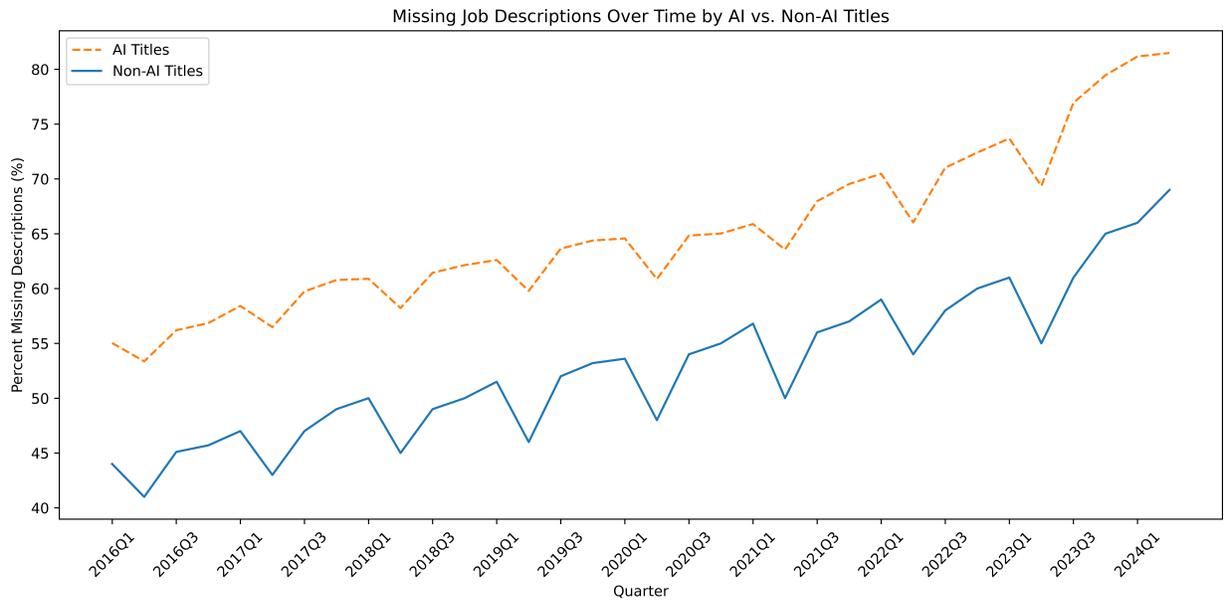
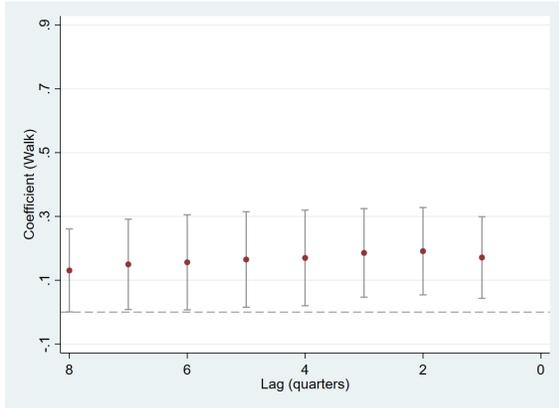
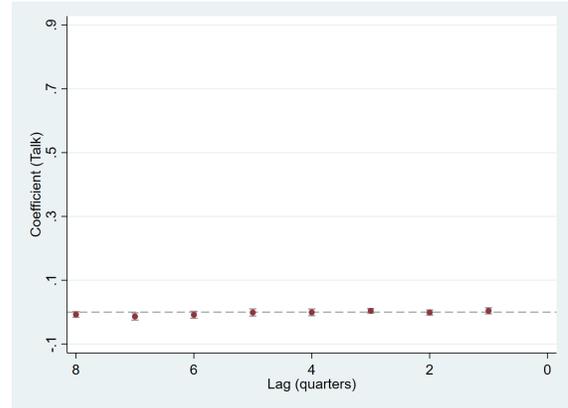


Figure A2. Distributed Lag Analysis on Innovation Outcomes

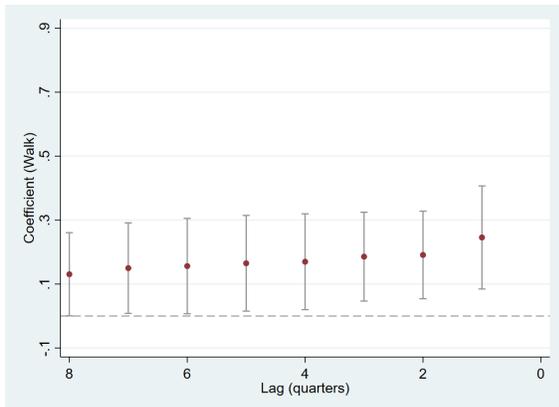
This figure presents results of distributed lag analysis of AI Walk and AI Talk on different innovation outcome variables. The coefficients are represented by dots and 95% confidence intervals are depicted by bars.



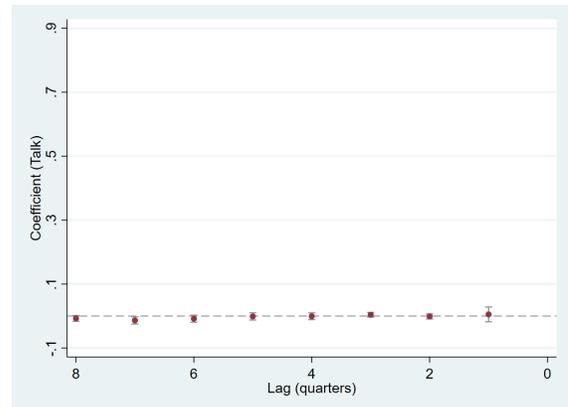
(a) Lagged AI Walk on Patent Counts



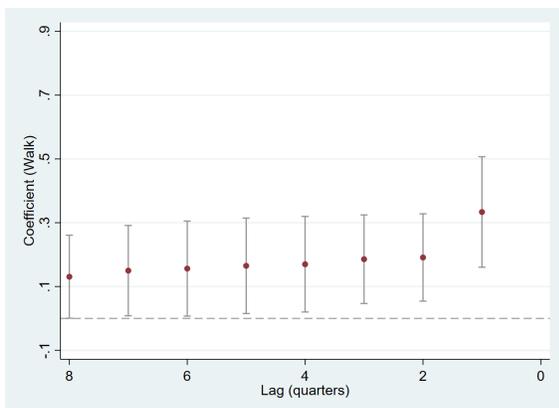
(b) Lagged AI Talk on Patent Counts



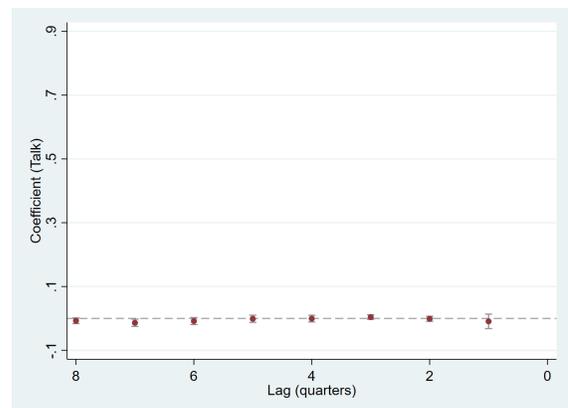
(c) Lagged AI Walk on Patent Value



(d) Lagged AI Talk on Patent Value



(e) Lagged AI Walk on Patent Citations



(f) Lagged AI Talk on Patent Citations

Table A1. AI Talk and AI Walk: New AI Employees

This table presents robustness test results from distributed-lag regressions analyzing the relationship between AI walk and AI talk lagged by various horizons, where the walk measure is constructed based on new AI employees. Control variables include size, cash/assets, R&D, age, and capital expenditures, and are lagged by one period. We control for industry fixed effects in column (1), industry fixed effects and quarter fixed effects in column (2), industry-quarter fixed effects in column (3), firm fixed effects in column (4), firm fixed effects and quarter fixed effects in column (5), and industry-quarter fixed effects in column (6). Standard errors are clustered at the firm level (***) $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

VARIABLES	Dependent Variable: AI Walk _t					
	(1)	(2)	(3)	(4)	(5)	(6)
Talk _{t-1}	0.0053 (0.0050)	0.0085 (0.0055)	0.0090 (0.0057)	-0.0073* (0.0042)	-0.0049 (0.0043)	-0.0051 (0.0044)
Talk _{t-2}	0.0087 (0.0068)	0.0120* (0.0070)	0.0128* (0.0071)	-0.0075* (0.0039)	-0.0059 (0.0040)	-0.0054 (0.0040)
Talk _{t-3}	0.0074 (0.0073)	0.0094 (0.0076)	0.0108 (0.0078)	-0.0098* (0.0052)	-0.0083 (0.0052)	-0.0075 (0.0052)
Talk _{t-4}	0.0322*** (0.0118)	0.0322*** (0.0117)	0.0324*** (0.0119)	0.0034 (0.0062)	0.0033 (0.0063)	0.0032 (0.0063)
Talk _{t-5}	0.0343*** (0.0121)	0.0316*** (0.0119)	0.0300** (0.0122)	0.0071 (0.0062)	0.0052 (0.0061)	0.0037 (0.0061)
Talk _{t-6}	0.0259** (0.0115)	0.0258** (0.0114)	0.0255** (0.0116)	0.0019 (0.0090)	0.0015 (0.0090)	0.0011 (0.0091)
Talk _{t-7}	0.0130 (0.0120)	0.0140 (0.0121)	0.0143 (0.0125)	-0.0088* (0.0051)	-0.0083 (0.0051)	-0.0080 (0.0051)
Talk _{t-8}	0.0226** (0.0109)	0.0245** (0.0109)	0.0258** (0.0111)	0.0018 (0.0079)	0.0017 (0.0079)	0.0028 (0.0080)
Controls	Y	Y	Y	Y	Y	Y
Firm FE	N	N	N	Y	Y	Y
Industry FE	Y	Y	N	N	N	N
Quarter FE	N	Y	N	N	Y	N
Industry-Quarter FE	N	N	Y	N	N	Y
Adj R ²	0.0822	0.0940	0.0788	0.290	0.297	0.286
Observations	13,602	13,602	13,598	13,595	13,595	13,593

Table A2. AI Talk and AI Walk: Job Titles

This table presents robustness test results from distributed-lag regressions analyzing the relationship between AI walk and AI talk lagged by various horizons, where the walk measure is constructed based on job titles. Control variables include size, cash/assets, R&D, age, and capital expenditures, and are lagged by one period. We control for industry fixed effects in column (1), industry fixed effects and quarter fixed effects in column (2), industry-quarter fixed effects in column (3), firm fixed effects in column (4), firm fixed effects and quarter fixed effects in column (5), and industry-quarter fixed effects in column (6). Standard errors are clustered at the firm level (***) $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

VARIABLES	Dependent Variable: AI Walk _t					
	(1)	(2)	(3)	(4)	(5)	(6)
Talk _{t-1}	0.0299 (0.0208)	0.0289 (0.0215)	0.0307 (0.0227)	0.0039 (0.0047)	0.0054 (0.0048)	0.0048 (0.0051)
Talk _{t-2}	0.0331** (0.0148)	0.0326** (0.0155)	0.0336** (0.0160)	0.0006 (0.0040)	0.0017 (0.0041)	0.0016 (0.0043)
Talk _{t-3}	0.0827** (0.0420)	0.0831** (0.0422)	0.0842* (0.0434)	-0.0070 (0.0058)	-0.0065 (0.0057)	-0.0062 (0.0060)
Talk _{t-4}	0.0608*** (0.0136)	0.0599*** (0.0142)	0.0610*** (0.0149)	0.0080 (0.0072)	0.0081 (0.0071)	0.0084 (0.0073)
Talk _{t-5}	0.0710* (0.0416)	0.0691* (0.0416)	0.0692 (0.0425)	0.0015 (0.0053)	0.0010 (0.0053)	0.0012 (0.0055)
Talk _{t-6}	0.0672* (0.0361)	0.0649* (0.0363)	0.0657* (0.0373)	0.0082 (0.0067)	0.0072 (0.0066)	0.0075 (0.0065)
Talk _{t-7}	0.0671* (0.0376)	0.0644* (0.0381)	0.0651* (0.0391)	0.0147 (0.0091)	0.0140 (0.0090)	0.0144 (0.0090)
Talk _{t-8}	0.0833** (0.0364)	0.0807** (0.0371)	0.0820** (0.0382)	0.0164 (0.0110)	0.0161 (0.0110)	0.0165 (0.0110)
Controls	Y	Y	Y	Y	Y	Y
Firm FE	N	N	N	Y	Y	Y
Industry FE	Y	Y	N	N	N	N
Quarter FE	N	Y	N	N	Y	N
Industry-Quarter FE	N	N	Y	N	N	Y
Adj R ²	0.162	0.164	0.140	0.843	0.844	0.842
Observations	13,602	13,602	13,598	13,595	13,595	13,593

Table A3. AI Talk and AI Walk: Inventor Hiring

This table presents robustness test results from distributed-lag regressions analyzing the relationship between AI walk and AI talk lagged by various horizons, using inventor hiring as the walk measure. Control variables include size, cash/assets, R&D, age, and capital expenditures, and are lagged by one period. We control for industry fixed effects in column (1), industry fixed effects and quarter fixed effects in column (2), industry-quarter fixed effects in column (3), firm fixed effects in column (4), firm fixed effects and quarter fixed effects in column (5), and industry-quarter fixed effects in column (6). Standard errors are clustered at the firm level (***) $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

VARIABLES	Dependent Variable: AI Walk _t					
	(1)	(2)	(3)	(4)	(5)	(6)
Talk _{t-1}	0.0079 (0.0082)	0.0105 (0.0087)	0.0113 (0.0091)	-0.0033** (0.0014)	-0.0040*** (0.0014)	-0.0040*** (0.0015)
Talk _{t-2}	0.0181** (0.0078)	0.0206** (0.0081)	0.0216*** (0.0082)	-0.0033** (0.0016)	-0.0040** (0.0016)	-0.0040** (0.0016)
Talk _{t-3}	0.0248*** (0.0094)	0.0262*** (0.0096)	0.0274*** (0.0100)	-0.0031 (0.0019)	-0.0031 (0.0020)	-0.0029 (0.0020)
Talk _{t-4}	0.0408*** (0.0147)	0.0408*** (0.0145)	0.0416*** (0.0149)	-0.0056* (0.0030)	-0.0059** (0.0030)	-0.0056* (0.0031)
Talk _{t-5}	0.0359** (0.0151)	0.0358** (0.0149)	0.0356** (0.0151)	-0.0085*** (0.0028)	-0.0080*** (0.0028)	-0.0076*** (0.0028)
Talk _{t-6}	0.0340** (0.0155)	0.0352** (0.0155)	0.0355** (0.0157)	-0.0068** (0.0028)	-0.0064** (0.0028)	-0.0060** (0.0029)
Talk _{t-7}	0.0290* (0.0157)	0.0307* (0.0158)	0.0313* (0.0161)	-0.0065** (0.0026)	-0.0059** (0.0026)	-0.0053** (0.0026)
Talk _{t-8}	0.0264 (0.0179)	0.0293 (0.0181)	0.0297 (0.0184)	-0.0064** (0.0029)	-0.0070** (0.0029)	-0.0062** (0.0029)
Controls	Y	Y	Y	Y	Y	Y
Firm FE	N	N	N	Y	Y	Y
Industry FE	Y	Y	N	N	N	N
Quarter FE	N	Y	N	N	Y	N
Industry-Quarter FE	N	N	Y	N	N	Y
Adj R ²	0.262	0.265	0.241	0.969	0.969	0.969
Observations	13,602	13,602	13,598	13,595	13,595	13,593

Table A4. AI Talk and AI Walk: Dummy Variables

This table presents robustness test results from distributed-lag regressions analyzing the relationship between AI walk and AI talk lagged by various horizons, using dummy variables as walk and talk measures. Control variables include size, cash/assets, R&D, age, and capital expenditures, and are lagged by one period. We control for industry fixed effects in column (1), industry fixed effects and quarter fixed effects in column (2), industry-quarter fixed effects in column (3), firm fixed effects in column (4), firm fixed effects and quarter fixed effects in column (5), and industry-quarter fixed effects in column (6). Standard errors are clustered at the firm level (***) $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

VARIABLES	Dependent Variable: AI Walk _t					
	(1)	(2)	(3)	(4)	(5)	(6)
Talk _{t-1}	0.0349** (0.0139)	0.0337** (0.0141)	0.0366** (0.0148)	0.0074 (0.0081)	0.0074 (0.0081)	0.0107 (0.0081)
Talk _{t-2}	0.0243* (0.0130)	0.0223* (0.0132)	0.0254* (0.0138)	0.0103 (0.0081)	0.0090 (0.0080)	0.0137 (0.0092)
Talk _{t-3}	0.0191 (0.0118)	0.0179 (0.0118)	0.0223* (0.0125)	-0.0047 (0.0074)	-0.0052 (0.0074)	0.0006 (0.0077)
Talk _{t-4}	0.0174 (0.0119)	0.0213* (0.0119)	0.0226* (0.0123)	0.0022 (0.0076)	0.0019 (0.0076)	0.0050 (0.0076)
Talk _{t-5}	0.0338*** (0.0130)	0.0307** (0.0131)	0.0289** (0.0135)	0.0051 (0.0093)	0.0044 (0.0093)	0.0059 (0.0093)
Talk _{t-6}	0.0361*** (0.0136)	0.0331** (0.0139)	0.0324** (0.0144)	0.0092 (0.0094)	0.0073 (0.0094)	0.0087 (0.0094)
Talk _{t-7}	0.0486*** (0.0142)	0.0467*** (0.0147)	0.0444*** (0.0152)	0.0155* (0.0091)	0.0145 (0.0092)	0.0151 (0.0093)
Talk _{t-8}	0.0334** (0.0159)	0.0383** (0.0163)	0.0363** (0.0170)	0.0110 (0.0097)	0.0103 (0.0098)	0.0105 (0.0100)
Controls	Y	Y	Y	Y	Y	Y
Firm FE	N	N	N	Y	Y	Y
Industry FE	Y	Y	N	N	N	N
Quarter FE	N	Y	N	N	Y	N
Industry-Quarter FE	N	N	Y	N	N	Y
Adj R ²	0.276	0.278	0.263	0.784	0.784	0.785
Observations	13,602	13,602	13,598	13,595	13,595	13,593