AI Washing*

Boyuan Li

University of Florida

July 2025

Abstract

This paper investigates AI washing — the practice by which firms exaggerate or falsely claim their investments in artificial intelligence (AI). We leverage large language models to analyze earnings conference call transcripts ("AI talk") and employee resume data ("AI walk") from U.S. public firms between 2016 and 2024. Our analysis reveals that AI talk does not predict future AI walk, even over multi-year horizons. Substantive AI walk, rather than AI talk, positively predicts future AI patenting activity, in terms of both quantity and impact. Firms engaging in empty AI talk without corresponding AI walk generate fewer and lower-quality AI patents. Importantly, institutional investors appear to recognize this disconnect, allocating more capital to firms with higher AI walk. While AI talk is associated with short-term stock return gains, potentially motivating inflated disclosures, only AI walk is correlated with superior long-run stock performance. In addition, firms with high managerial incentives are significantly more likely to increase AI talk without a corresponding rise in walk, suggesting strategic hype. Overall, our findings highlight a critical disconnect between firms' AI rhetoric and their substantive AI investments, revealing a misalignment 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

*I am grateful for the invaluable guidance of Joel Houston (advisor), Sehoon Kim, Alejandro Lopez-Lira, and Baolian Wang. I am also thankful for the helpful comments from Mark Flannery, Zhiguo He, Nitish Kumar, Mahendrarajah Nimalendran, Jay Ritter, Michael Ryngaert, Yuehua Tang, Yanbin Wu, along with seminar participants at the University of Florida and the 2024 FMA Doctoral Student Consortium. Email address: boyuan.li@warrington.ufl.edu.

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 in their business operations. 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-related workforce or innovation). Given the information asymmetry between managers and investors in emerging technologies, 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

[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

the cost of misallocated capital and missed innovation opportunities.

We provide the first empirical examination of AI washing among U.S. public firms over the period of 2016-2024. Specifically, we quantify the discrepancies between firms' AI talk — measured by their discussions in quarterly earnings conference calls — and their actual AI walk — measured by their AI-related workforce composition. 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.

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 filter the discussions along multiple dimensions — including whether a firm is actively investing in AI, whether it is outsourcing AI development, and whether the statements are forward-looking.

To measure AI walk, we leverage detailed employee resume data that captures firms' AI-related human capital, following the approach of Babina et al. (2024). This allows us to assess substantive workforce investment in AI across firms and compare it to their public statements. Our data show a sharp rise in AI talk over the sample period: the percentage of firms mentioning "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]

Our empirical analysis yields four key findings. First, we find that AI talk does not predict future AI workforce investment, even over a two-year horizon. This lack of alignment between rhetoric and action is consistent with agency-driven AI washing, where managers use AI talk strategically without follow-through, and is becoming more salient in recent years.

Second, to validate our measures, we show that AI walk is strongly associated with future AI innovation, as measured by patent quantity, value, and citations. In contrast, AI talk is not a significant predictor of innovation outcomes, suggesting that only substantive investment — rather than rhetorical signaling — drives meaningful technological progress.

Third, we examine how institutional investors respond to AI engagement. Controlling for AI talk, we find that firms with greater AI walk are significantly more likely to be held by AI-focused mutual funds and ETFs, and to be held more heavily. Moreover, institutional investors discount AI talk when it is not supported by substantive action, suggesting that they can detect and respond to AI washing behavior.

Finally, we explore potential incentives for AI washing by examining stock price reactions. We find that AI talk significantly boosts short-term cumulative abnormal returns (CARs) around earnings announcements, whereas AI walk does not, suggesting that firms may engage in AI talk to exploit short-term investor enthusiasm. However, when we turn to longer-run stock performance, measured by buy-and-hold abnormal returns (BHAR) over the subsequent 180 trading days from earnings announcements, the institutional investors eventually penalizes firms whose talk is not backed by tangible AI investments. The delayed recognition likely reflects the time it takes for investors to update beliefs about the credibility of AI communication. This divergence between short- and long-term market responses reinforces the strategic temptation for firms to talk without walking. Additionally, firms with high managerial incentives — proxied by equity compensation sensitivity (Delta) — are significantly more likely to increase AI talk without a corresponding rise in walk, suggesting strategic hype.

Together, these findings provide novel evidence that AI washing is both prevalent and strategically motivated, shaped by a tension between short-term market incentives and long-term value creation. By separating AI talk from 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); Babina et al. (2024)). However, to the best of our knowledge, we are the first to systematically investigate the gap between firms' AI claims and their actual investments. 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)). We extend this literature by documenting how AI-related disclosures serve as a strategic tool for managerial signaling, despite limited follow-through in actual AI investment.

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); Bybee et al. (2024); Hirshleifer et al. (2023); Jha et al. (2023); Lopez-Lira and Tang (2023); Sautner et al. (2023); 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 keyword-based filtering with LLM-based analysis.

From a policy perspective, our results underscore the need for enhanced regulatory scrutiny of AI-related corporate disclosures. The SEC's recent enforcement actions indicate growing concerns about misleading AI claims, and our findings provide empirical support for such regulatory interventions. For investors, our study highlights the importance of distinguishing between firms that genuinely invest in AI and those that merely engage in AI-related rhetoric.

The remaining of the paper is organized as follows. Section 2 describes the data sources. Section 3 outlines the methodology and presents some stylized facts. Section 4 provides validation of the measures. Section 5 illustrates firm characteristics of AI talkers and AI walkers. Section 6 examines the dynamics of AI talk and AI walk and provides suggestive evidence of AI washing. Section 7 examines whether institutional investors can distinguish AI washing behavior. Section 8 investigates potential motivations of AI washing, followed by the conclusion in Section 9.

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 companies 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 or deemphasize 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 companies. 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 (i.e., AI walk), we follow Babina et al. (2023a,b, 2024) and use AI-related human capital as a proxy. We obtain employee resume data from the Revelio Labs, which provides structured employment histories for professionals working in U.S. public companies, offering granular insights into firms' workforce composition and skill distribution. The dataset includes information such as employee names, company affiliations, job descriptions, tenure at specific positions, listed skills, and educational backgrounds.

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 2016 Q1 and 2024 Q2 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 provide the entire section to LLM and ask it to classify the 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.

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.

Our sample includes all AI patents, as defined by the USPTO, that were issued through the end of 2023. To better reflect the timing of innovation, we aggregate patenting activity at the firm-quarter level using the filing date rather than the grant date, given the typical lag between the two. We believe this approach provides a more accurate, near real-time measure 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

[2]The fund is classified 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 is classified as an AI-using fund if it applies AI in its own investment strategies (e.g., AI-driven algorithms for portfolio selection).

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 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 a broader swath of the 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 per firm in the dataset. 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 2016 Q1 to 2024 Q2. 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 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.[3] 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 results 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]

[3] For example, "*artificial*" and "*intelligence*" are two separate words, but are frequently used together in the bigram, "*artificial_intelligence*".

3.2 ChatGPT Analysis

To refine our measurement of AI talk, we apply LLMs to filter the discussions. In particular, we provide ChatGPT with AI-related text from conference call transcripts, as identified by the AI-related keywords, and ask it to answer the following prompt with a series of questions —

"Analyze the following manager presentation excerpt from an earnings conference call. Based on this source only, answer concisely in JSON format:

Input: row['Text']

Questions (be brief and specific):

1. Is the company actively investing in AI? (Yes/No/N/A)

2. Is the company outsourcing or collaborating with other companies on AI? (Yes/No/N/A)

3. What is the primary focus of the firm's AI investment? (If applicable, answer with a list of specific keywords)

4. Does this AI investment require AI-related skills from employees? (Yes/No/N/A)

5. Is the firm discussing past, current, or planned AI initiatives?

6. Does the company provide specific evidence about their AI investment? (Yes/No/N/A)

7. Do the firm's AI-related claims align with measurable outcomes? (Yes/No/N/A)

8. Has the company outlined a clear long-term AI vision or road map? (Yes/No/N/A)"

3.3 AI Talk Measure

Based on the output from ChatGPT, we further filter the conference call text by requiring *Q1: investment* = *"Yes", Q2: outsourcing* $\neq$ *"Yes", Q4: employee* = *"Yes", and Q5: timeline* = *"current."* This ensures that the discussions are truly related to current in-house AI investment that the firms are actively engaging in, which allows for comparison with real-time firm employee profiles. 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.

3.4 AI Walk Measure

To identify AI-related position, we require at least one AI keyword from the corresponding yearly AI dictionary to be present in the position description text. Then we construct the AI walk measure as following —

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

where i denotes firm and t denotes quarter.

For job positions that span multiple quarters, we treat them as *active* in each quarter they remain open. To account for differences in firm employment size, we scale the number of AI-related positions by the lagged total number of active positions with available job descriptions from the previous quarter. This standardization also helps mitigate bias due to missing position texts — particularly in later periods where more jobs may lack descriptions.

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 on AI-related discussions and investments among 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.

3.5 Focus of AI Talk

To gain further insights into the focus of firms' AI talk, we compare the frequency of commonly used terms over time. Figure 5 presents the time-series frequency of select AI-related focus appearing in earnings conference calls. Notably, the frequency of terms such as "machine learning" and "big data" began increasing earlier in the sample period. In contrast, terms like "generative AI" only emerged in 2023, aligning with the recent surge in interest and advancements in generative AI technologies.

These trends provide further validation that our AI talk measure effectively captures the evolution and prevalence of AI discussions by firms over time. The shifting terminology reflects how firms adapt their AI narratives in response to technological progress and market developments.

[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, AI Talk alone may show weaker or even negative relationships with innovation if it reflects symbolic signaling rather than substantive action.

Specifically, we investigate three key dimensions of AI innovation: (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:

$$\mathrm{Y}_{i,t} = \beta_1 \mathrm{Talk}_{i,t-1} + \beta_2 \mathrm{Walk}_{i,t-1} + \gamma \mathbf{X}_i + \alpha_{\text{firm}} + \lambda_{\text{industry-quarter}} + \varepsilon_{i,t}, \tag{3}$$

where $\mathrm{Y}_{i,t}$ takes on one of the three dependent variables, and the key independent variables are $\mathrm{Talk}_{i,t-1}$ and $\mathrm{Walk}_{i,t-1}$, which measure the standardized levels of AI talk and AI walk of firm i, respectively, lagged by one period.[4] The vector $\mathbf{X}_{i,t}$ includes a set of control variables capturing firm characteristics and financial indicators. The main specification also includes firm fixed effects, α_{firm}, and industry-quarter fixed effects, $\lambda_{\text{industry-quarter}}$, to account for time-invariant firm characteristics and time-specific industry-level shocks.

[Insert Table 3 here]

Table 3 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: firms with higher walk measures in the previous quarter exhibit significantly more AI-related patenting activity, with large and statistically significant coefficients on patent counts (0.224 - 0.247), patent value (0.441

[4] In distributed lag analysis, we use multiple lagged periods as independent variables.

- 0.485), and patent citations (0.322 - 0.334). These results support the interpretation of walk as a meaningful proxy for substantive AI engagement.

In contrast, AI Talk ($Talk_{t-1}$) is negatively associated with innovation outcomes, with statistically significant negative coefficients across all three dependent variables. While these estimates are statistically robust, the magnitudes are relatively modest, suggesting that excessive AI-related language alone has limited or even slightly detrimental 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 A1 here]

Given the time lag between 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. As illustrated in Figure A1, the predictive power of AI walk 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 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 2016 Q1. Table 2 presents descriptive statistics that compare firms based on their engagement with artificial intelligence (AI), distinguishing between firms that ever talk about AI (Panel A) and those that ever walk the talk (Panel B).

Across several firm characteristics, Ever Talk firms are systematically larger with more resource. On average, Ever Talk firms have significantly higher sales (6.08 vs. 5.49), more cash holdings (0.19 vs. 0.14), and greater capital expenditures (2.99 vs. 2.46), all with statistically significant differences. These firms are also slightly older (26.8 vs. 25.4) and invest more in R&D (0.012 vs. 0.009). 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, Ever Walk firms — those with concrete AI-related activity — exhibit even stronger distinctions in firm characteristics. These firms are substantially larger, with a mean log sales of 6.40 compared to 4.96 among non-walkers, and demonstrate significantly greater investment capacity (3.32 vs. 1.93). They also tend to be older (29.0 vs. 22.5), 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 (0.03 vs. –0.69). 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.

6 AI Talk-Walk Dynamics

With measures of AI talk and AI walk, we are able to examine the dynamic relationship between them using the following regression specification —

$$\text{AI Walk}_{i,t} = \alpha + \sum_{h=1}^{8} \beta_h \text{AI Talk}_{i,t-h} + \gamma \mathbf{X}_{i,t} + \alpha_{\text{firm}} + \lambda_{\text{industry-quarter}} + \varepsilon_{i,t}, \tag{4}$$

where AI $\text{Walk}_{i,t}$ represents the level of AI investment for firm i at quarter t. The key independent variable, AI $\text{Talk}_{i,t-h}$, captures the extent to which firm i discusses AI in its earnings call at quarter $t-h$. The summation term $\sum_{h=1}^{8} \beta_h$ reflects the estimated effect of AI talk at different time horizons, from one to eight quarters ahead.[5] The vector $\mathbf{X}_{i,t}$ includes control variables such as ROA, leverage, size, sales growth, R&D, and CAPEX, with γ denoting the corresponding coefficients. In the first set of regressions, we incorporate industry-quarter fixed effects $\lambda_{\text{industry-quarter}}$ 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 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, 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 2, 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, 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 AI talk and AI walk is consistent with the practice of AI washing, where firms emphasize AI narratives in public communications without backing them with substantive implementation.

[5] In deciding the number of lags to include, we have to balance the tradeoff between examining a broader scope and including more observations in the analysis. Given the fact that the majority of AI talk is about investment plans within one to two years, we include lagged talk measures up to 8 quarters.

[Insert Table 5 here]

Furthermore, when we split the sample into pre-2019 and post-2019 periods, we find that firms did "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.

7 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:

$$\mathrm{Y}_{i,t} = \beta_1 \mathrm{Talk}_{i,t-1} + \beta_2 \mathrm{Walk}_{i,t-1} + \beta_3 \left(\mathrm{Talk}_{i,t-1} \times \mathbf{1}_{\text{no walk},i,t-1}\right) + \gamma \mathbf{X}_{i,t} + \alpha_{\text{firm}} + \lambda_{\text{industry-quarter}} + \varepsilon_{i,t}, \quad (5)$$

where $\mathrm{Y}_{i,t}$ is one of two dependent variables: (1) the number of funds holding the firm's stock, and (2) the percent of market value held by funds. In separate regressions, we construct these two outcomes variables for AI-focused mutual funds and ETFs and all all institutional funds, respectively. Additionally, the specification includes an interaction term between $\mathrm{Talk}_{i,t-1}$ and an indicator variable that equals 1 if firm i exhibits zero walk in all four subsequent quarters, and 0 otherwise. The vector $\mathbf{X}_{i,t}$ includes a set of control variables capturing firm characteristics and financial indicators. The main 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. The specification does *not* include a stand-alone term for the no-walk indicator, $\mathbf{1}_{\text{no walk},i,t-1}$, because it would be undefined: by construction, the indicator equals 1 only when the firm also engages in talk, meaning there are no observations where the dummy is 1 and $\mathrm{Talk}_{i,t-1} = 0$. As a result, the main effect of the indicator is not separately identified from the interaction term.

[Insert Table 6 here]

Panel A of Table 6 focuses on AI-focused mutual funds and ETFs. In Columns (1) and (2), the number of AI-focused funds holding a firm is positively associated with both lagged AI talk and AI walk. Specifically, a one standard deviation increase in AI walk is associated with approximately 0.28–0.29 more AI funds holding the firm's stock. AI talk also has a positive association (coefficients 0.05), although smaller in magnitude. Importantly, the interaction term is negative and statistically significant, suggesting that firms that talk more about AI but lack corresponding walk are penalized by AI-focused institutional investors. The effect size of the interaction (around -0.06 to -0.07) nearly offsets the positive coefficient on AI talk, implying that talk without walk can deter AI-focused institutional interest.

Columns (3) and (4) show that while AI walk has a consistently positive (though imprecise) association with the percent of market value held by AI funds, AI talk alone has no significant effect. However, the negative and significant interaction again suggests that AI-washing behavior reduces AI fund ownership.

The relationships observed in Panel B generally mirror those in Panel A but with some key differences in magnitude and strength. AI walk remains a strong and statistically significant predictor of institutional ownership, with coefficients around 12 to 14 for the number of funds (Columns 5–6) and 5.6 to 5.7 for percent of market value held (Columns 7–8). This reinforces the idea that actual AI-related investments are rewarded broadly by the market.

AI talk also shows a positive association with both fund count and market value held. However, the interaction term coefficient is smaller in magnitude with weaker statistical significance compared to Panel A. While the coefficients remain negative (e.g., -2.42 to -3.32 for fund count), they are only marginally significant, and the coefficients for percent market value held are not statistically significant.

This suggests that while all institutional investors do respond to empty AI talk, their reaction is less sensitive and more moderate compared to AI-focused investors. In other words, generalist institutional funds may be slower or less attuned to detecting AI washing, whereas AI funds are more discerning in penalizing firms that overstate their AI engagement.

8 Why Do Firms AI Wash?

8.1 Short-run and Long-run Stock Returns

We explore potential incentives for AI washing by analyzing both short-run and long-run abnormal returns. Table 7 presents the results from event studies, where the dependent variables are market-adjusted CARs, measured over the three-day window surrounding earnings announcements, or BHARs, measured from 180 days to 360 days from the earnings announcements dates. All models control for a comprehensive set of firm fundamentals (e.g., ROA, R&D, SUE, CAPEX), and include both firm and industry-quarter fixed effects to account for unobserved heterogeneity.

[Insert Table 7 here]

Across both short-term and long-term windows, AI talk is associated with significantly positive market reactions, especially in the short run. For example, in Columns (1)–(2), a one standard deviation increase in Talk is approximately associated with 0.25 basis point higher CAR, which is statistically significant at the 1% level, suggesting that markets react favorably to firms that engage in AI-related communication. This effect generally persists over longer horizons as well, with Talk continuing to predict positive BHARs up to one year out. However, the magnitude of the coefficients is moderate, and statistical significance declines somewhat over longer horizons.

Importantly, the interaction term, which captures whether AI talk is decoupled from actual AI walk, is negative across all specifications. While not statistically significant in the CAR or 180-day BHAR windows, it becomes significant at the 5% level for 270-day and 360-day BHARs. This suggests that the market may initially reward AI talk, but eventually penalizes firms whose talk is not backed by tangible AI investments. The delayed recognition likely reflects the time it takes for investors to update beliefs about the credibility of AI communication.

Taken together, these findings underscore a potential intertemporal disconnect in market reactions to AI engagement. In the short run, firms appear able to influence investor sentiment and stock prices through AI-related communications alone, even in the absence of real investment — creating a potential incentive for AI washing. The strong short-term effect of AI talk on CARs, combined with the lack of investor response to AI walk in the short term, suggests that some firms may strategically use AI rhetoric to shape market perceptions. However, in the long run, only firms that back up their AI narratives with real, tangible investments are rewarded with sustained stock price appreciation. This divergence between short- and long-run market responses highlights a misalignment between investor sentiment and firm fundamentals, and underscores the importance of distinguishing between superficial signaling and substantive action in the evaluation of technological innovation.

8.2 Managerial Incentives and ChatGPT Release

To further investigate managerial incentives behind AI washing, we explore whether firms with higher equity-based CEO incentives 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 symbolic communication strategies, such as emphasizing AI capabilities, to shape investor perceptions.

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

$$Y_{i,t} = \beta_1 \text{Post}_t \times \text{HighDelta}_i + \gamma \mathbf{X}_{i,t} + \alpha_{\text{firm}} + \lambda_{\text{industry-quarter}} + \varepsilon_{i,t}, \tag{6}$$

where $Y_{i,t}$ is the outcome variable for firm i in quarter t, and can be one of the following: (1) AI talk, (2) AI walk, and (3) the talk-walk gap. The variable Post_t is an indicator for the post-ChatGPT period, while 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 $\text{Post}_t \times \text{HighDelta}_i$, hence, identifies whether high-delta firms exhibit differential changes in behavior after ChatGPT's release. The model also includes a vector of firm-level control variables $\mathbf{X}_{i,t}$, as well as 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.

[Insert Table 8 here]

The results, presented in Table 8, 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). The most telling evidence comes from columns (5) and (6), where we examine the talk-walk gap directly. Here, the coefficients remain large (0.875 and 0.843) and highly significant, indicating that the increase in AI talk among high-delta firms is not matched by a corresponding increase in AI implementation.

[Insert Figure 6 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-walk gap on yearly indicator variables, interacted with the HighDelta indicator, and plot the coefficients along with their 95% confidence intervals in Figure 6. 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. The results are consistent with a behavioral explanation for AI washing: when reputational or financial gains can be realized through talk alone, and especially when executives stand to benefit directly from stock price appreciation, firms may have stronger motives to engage in symbolic disclosure.

9 Conclusion

This study provides empirical evidence on the phenomenon of AI washing, where firms strategically emphasize AI investment in their public communications without necessarily making corresponding efforts. By leveraging LLMs on earnings call transcripts and employment data, we document a stark disconnect between AI talk and AI walk. While firms increasingly discuss AI in conference calls, such rhetoric does not translate into measurable AI workforce expansion.

Our findings reveal several important insights. First, AI walk is a strong and consistent predictor of firms' future AI innovation, as evidenced by increase in patent volume, quality, and value. In contrast, AI talk exhibits no significant relationship with future AI-related patents, indicating that mere discussion does not equate to tangible technological advancement. Second, we find that institutional investors distinguish between AI talk and AI walk, allocating capital preferentially to firms with substantive AI investments rather than those that merely engage in AI rhetoric. AI-walking firms attract a higher likelihood of being held by AI-focused mutual funds and ETFs, reinforcing the notion that sophisticated investors recognize and reward actual AI engagement.

Lastly, we examine the potential incentives for AI washing and uncover that AI talk is positively associated with short-term abnormal stock returns around earnings announcements, while AI walk has no immediate effect. However, only AI walk is positively associated with long-run stock performance, indicating that the market eventually rewards substance. This suggests that firms may strategically engage in AI talk to influence investor sentiment and boost short-term market performance, even in the absence of substantive AI investments. Firms with high managerial incentives are significantly more likely to increase AI talk without a corresponding rise in walk, suggesting strategic hype.

These findings have important implications for regulators, investors, and policymakers, highlighting the need for enhanced scrutiny of AI-related claims in corporate disclosures. As AI continues to evolve and firms seek to capitalize on its perceived value, distinguishing genuine AI investment from strategic rhetoric will remain a crucial challenge for market participants.

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*.

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*.

Bybee, L., Kelly, B., Manela, A., and Xiu, D. (2024). Business news and business cycles. *The Journal of Finance, forthcoming*.

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.

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.

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. (2023). Chatgpt and corporate policies. chicago booth research paper no. 23-15; fama-miller working paper, becker friedman institute for economics working paper no. 2023-103.

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.

Lopez-Lira, A. and Tang, Y. (2023). Can chatgpt forecast stock price movements? return predictability and large language models. *arXiv preprint arXiv:2304.07619*.

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 2016 Q1 to 2024 Q2.

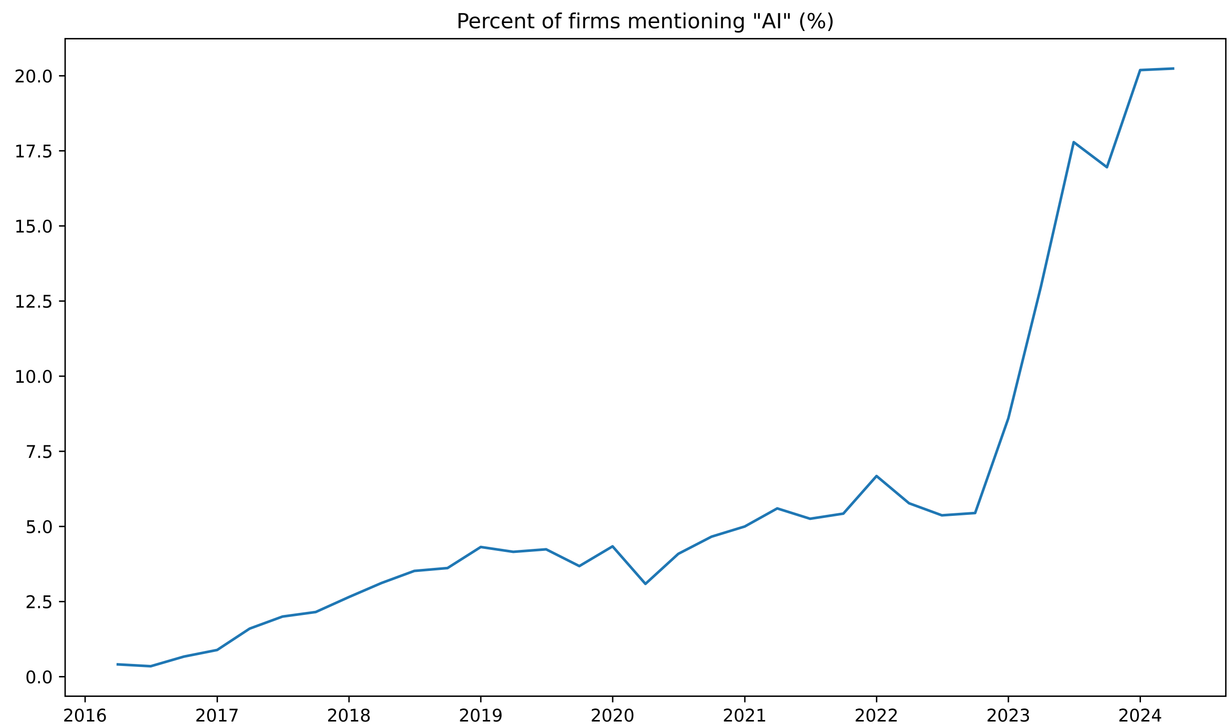

Figure 2. Word Clouds of AI Talk

This figure presents word clouds for select years depicting the top 100 keywords in terms of cosine similarity scores with respect to "AI" in conference calls.

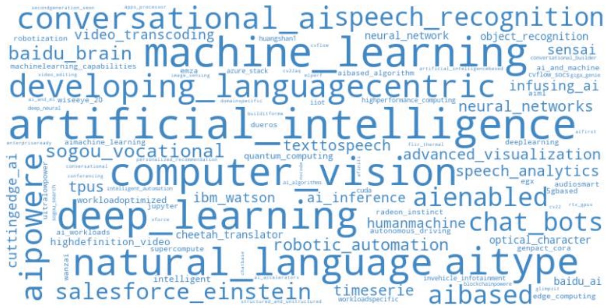

(a) 2019

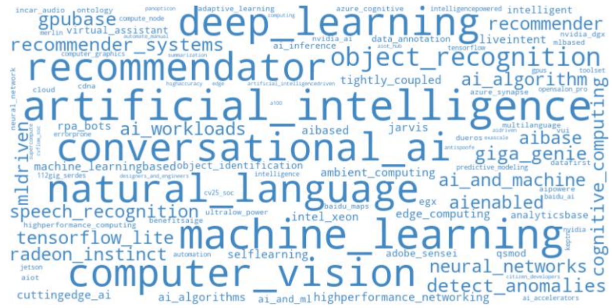

(b) 2020

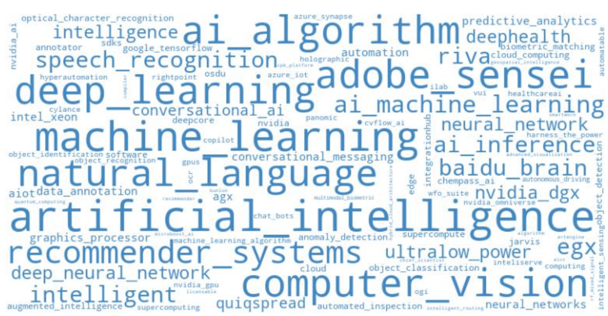

(c) 2021

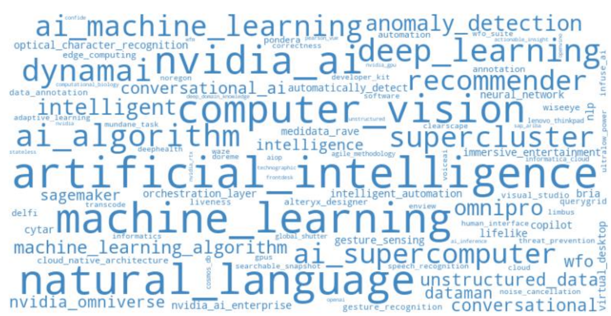

(d) 2022

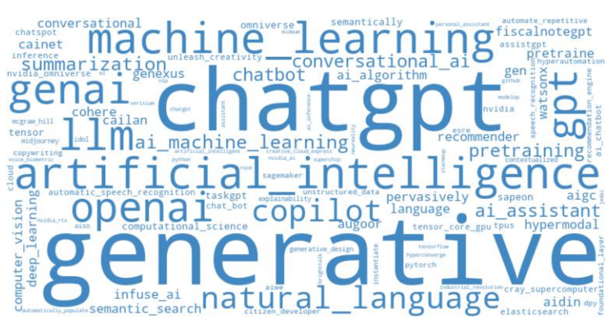

(e) 2023

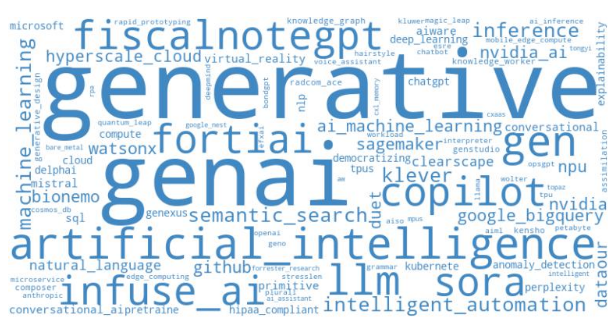

(f) 2024

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) from 2016 Q1 to 2024 Q2.

Panel A

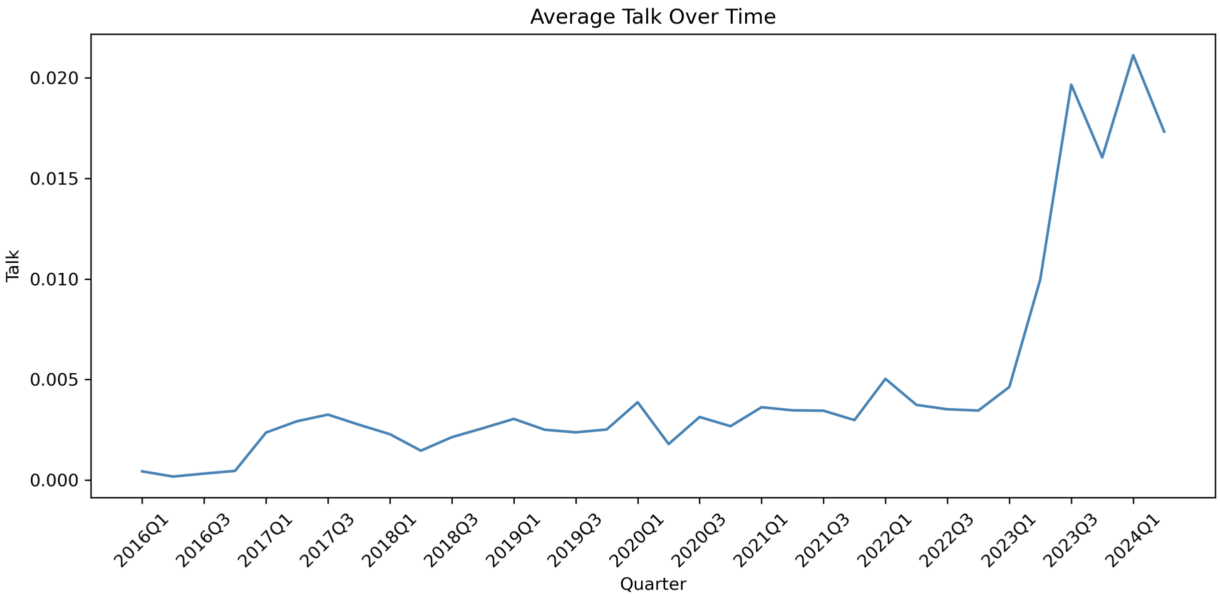

Panel B

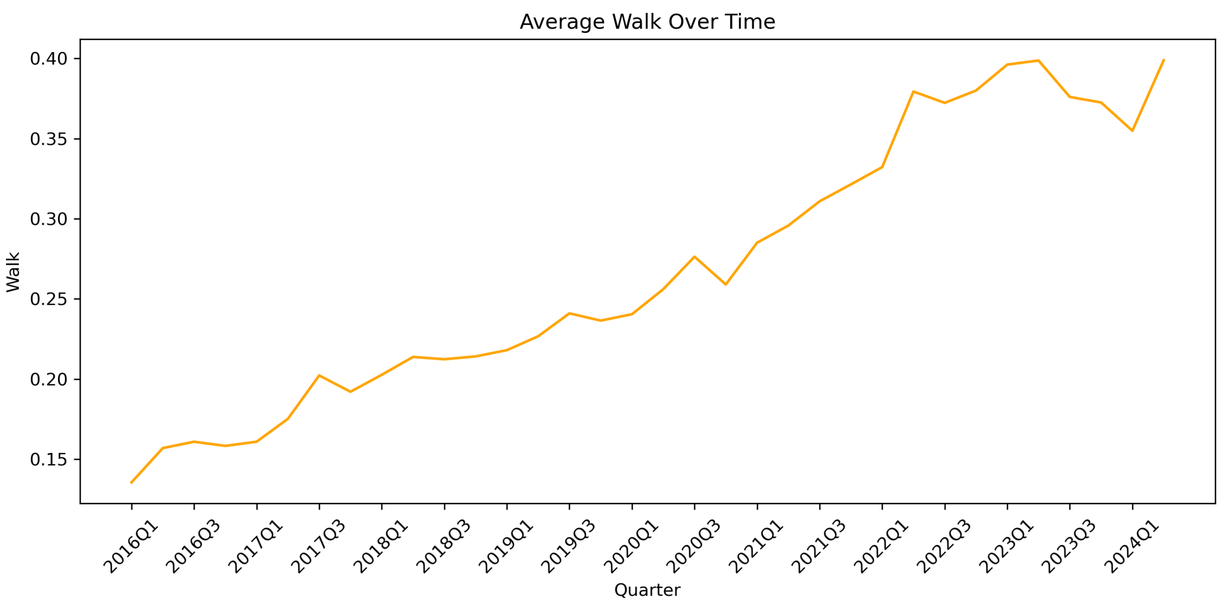

Figure 4. AI Talk and AI Walk by Industry

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

Panel A

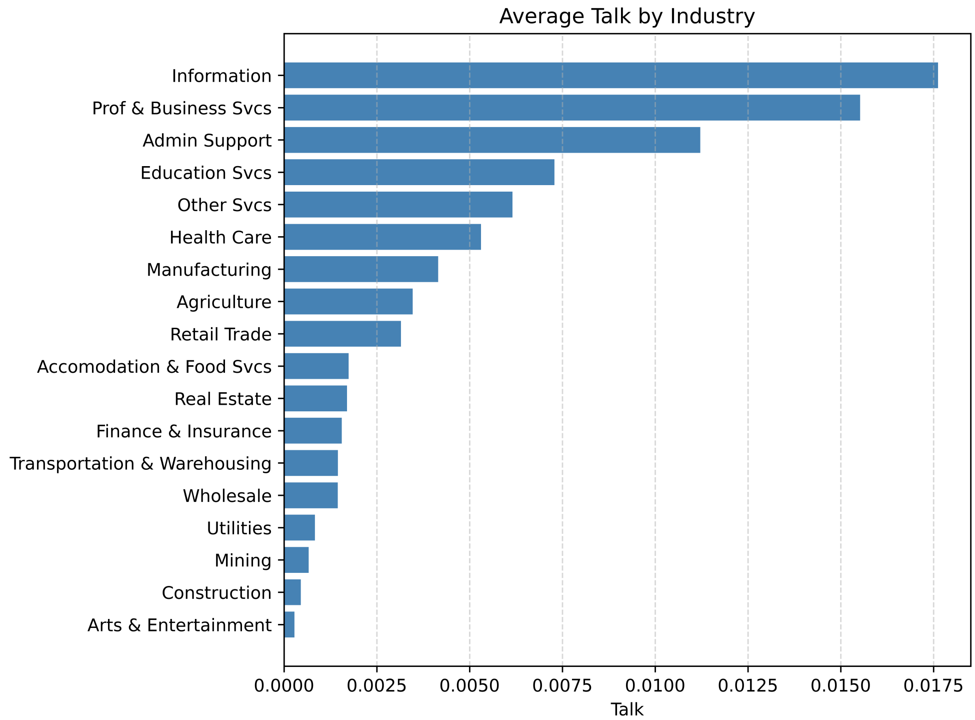

Panel B

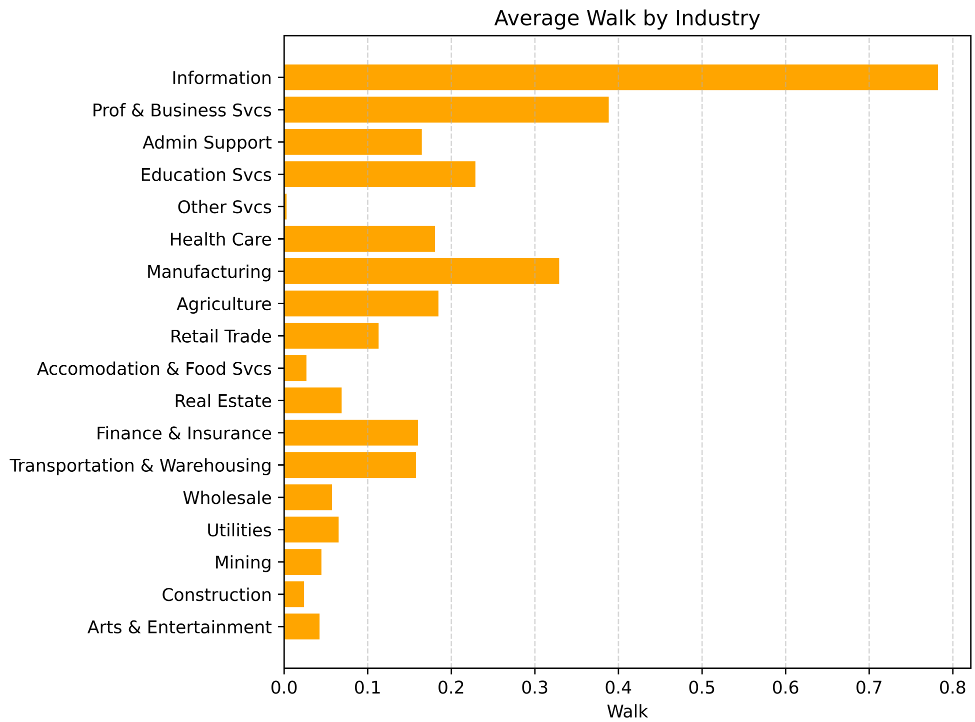

Figure 5. Time Series of Topics in AI Talk

This figure plots the frequency of select AI-related phrases in quarterly earnings conference calls from 2016 Q1 to 2024 Q2.

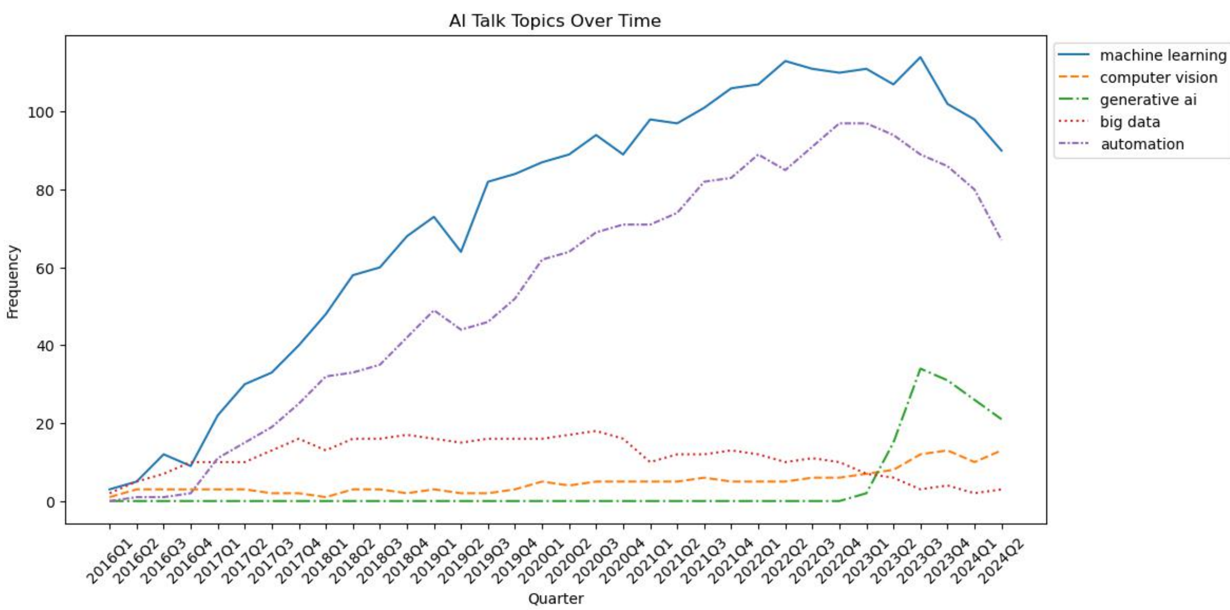

Figure 6. Pre- and Post-Event Differences Between High- and Low-Delta Firms

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

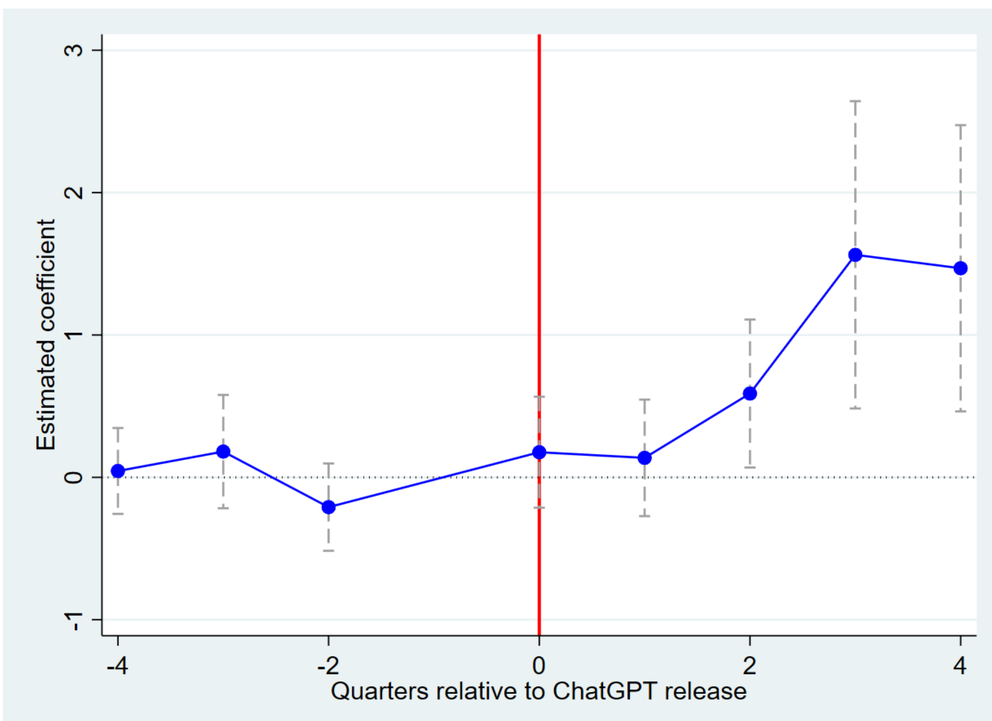

Table 1. Summary Statistics

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
Held by A Funds	0.2893	0.4534	0.0000	0.0000	0.0000	1.0000	1.0000	20,135
Number of AI Funds	0.8306	1.8481	0.0000	0.0000	0.0000	1.0000	5.0000	20,135
% Market Value Held	0.1699	0.6382	0.0000	0.0000	0.0000	0.0062	0.8763	20,128
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. 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 2016 Q1 to 2024 Q2. The differences in means between the two groups of firms and the corresponding t-statistics are also reported.

Panel A	Ever Talk				Never Talk					
Variable	Mean	Median	SD	N	Mean	Median	SD	N	Mean Diff	*t*-stat
$\log(\text{Sales})_{2016}$	6.0798	6.3027	2.0439	17929	5.4890	5.5725	1.9498	30278	0.5907	31.57
$\text{Cash/Assets}_{2016}$	0.1875	0.1176	0.1972	17929	0.1448	0.0625	0.2081	30278	0.0427	22.20
Age_{2016}	26.8264	21.0000	17.7326	17792	25.3611	21.0000	17.3211	30169	1.4653	7.63
$\text{R\&D}_{2016}$	0.0119	0.0000	0.0262	17929	0.0093	0.0000	0.0299	30278	0.0026	9.66
ROA_{2016}	-0.0031	0.0083	0.0514	17929	-0.0032	0.0056	0.0476	30278	0.00003	0.06
$\log(\text{CAPEX})_{2016}$	2.9907	2.9134	1.9612	17929	2.4557	2.2996	1.9330	30278	0.5350	29.21

Panel B	Ever Walk				Never Walk					
Variable	Mean	Median	SD	N	Mean	Median	SD	N	Mean Diff	*t*-stat
$\log(\text{Sales})_{2016}$	6.3970	6.6179	2.0023	25087	4.9619	5.1579	1.7223	23120	1.4351	84.03
$\text{Cash/Assets}_{2016}$	0.1739	0.0969	0.1997	25087	0.1463	0.0634	0.2100	23120	0.0276	14.76
Age_{2016}	29.0448	22.0000	23.8026	24977	22.4923	21.0000	16.9735	22984	6.5526	34.45
$\text{R\&D}_{2016}$	0.0117	0.0000	0.0282	25087	0.0087	0.0000	0.0290	23120	0.0030	11.38
ROA_{2016}	0.0003	0.0086	0.0480	25087	-0.0069	0.0042	0.0521	23120	0.0073	15.93
$\log(\text{CAPEX})_{2016}$	3.3185	3.2958	1.9732	25087	1.9343	1.6449	1.6714	23120	1.3842	82.75

Table 3. 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, including AI patent counts, AI patent value, and AI patent citations. The dependent variable is the log of the number of AI patents. Control variables include size, cash/assets, R&D, age, and capital expenditures. 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$_{t-1}$	-0.029***	-0.029***	-0.061***	-0.062***	-0.025**	-0.026***
	(0.011)	(0.011)	(0.021)	(0.022)	(0.010)	(0.010)
Walk$_{t-1}$	0.224***	0.247***	0.441***	0.485***	0.322***	0.334***
	(0.072)	(0.073)	(0.115)	(0.117)	(0.090)	(0.091)
Size$_{t-1}$		0.047***		0.074**		0.044***
		(0.015)		(0.032)		(0.017)
Cash/Assets$_{t-1}$		0.082		0.101		0.216**
		(0.072)		(0.174)		(0.104)
R&D$_{t-1}$		0.585		0.990		0.244
		(0.590)		(1.212)		(0.687)
Age$_{t-1}$		0.076		0.304**		0.049
		(0.047)		(0.149)		(0.042)
CAPEX$_{t-1}$		0.018**		0.027		0.015
		(0.007)		(0.017)		(0.009)
Firm FE	Y	Y	Y	Y	Y	Y
Industry-Quarter FE	Y	Y	Y	Y	Y	Y
Adj R^2	0.732	0.734	0.668	0.671	0.526	0.527
Observations	19,973	19,910	19,973	19,910	19,973	19,910

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 lagged sales, cash, R&D, age, and capital expenditures, as well as firm and industry-quarter fixed effects. Standard errors are clustered at the firm level (*** p<0.01, ** p<0.05, * p<0.1).

	Dependent Variable: AI Walk$_t$					
Variable	(1)	(2)	(3)	(4)	(5)	(6)
Talk$_{t-1}$	0.0310	0.0327	0.0340	0.0040	0.0045	0.0047
	(0.0189)	(0.0198)	(0.0210)	(0.0032)	(0.0032)	(0.0033)
Talk$_{t-2}$	0.0454**	0.0475**	0.0487**	0.0021	0.0023	0.0023
	(0.0189)	(0.0197)	(0.0201)	(0.0025)	(0.0025)	(0.0026)
Talk$_{t-3}$	0.0541**	0.0552**	0.0572**	-0.0024	-0.0020	-0.0019
	(0.0244)	(0.0251)	(0.0263)	(0.0042)	(0.0041)	(0.0043)
Talk$_{t-4}$	0.1113**	0.1115**	0.1127**	-0.0020	-0.0020	-0.0022
	(0.0467)	(0.0467)	(0.0479)	(0.0072)	(0.0072)	(0.0076)
Talk$_{t-5}$	0.1059**	0.1045**	0.1039**	-0.0032	-0.0034	-0.0040
	(0.0450)	(0.0447)	(0.0455)	(0.0063)	(0.0063)	(0.0065)
Talk$_{t-6}$	0.1034**	0.1029**	0.1031**	-0.0013	-0.0017	-0.0024
	(0.0428)	(0.0429)	(0.0438)	(0.0061)	(0.0061)	(0.0064)
Talk$_{t-7}$	0.0991**	0.0990**	0.0992**	0.0035	0.0035	0.0031
	(0.0421)	(0.0426)	(0.0435)	(0.0058)	(0.0058)	(0.0059)
Talk$_{t-8}$	0.1145**	0.1156**	0.1159**	0.0031	0.0029	0.0025
	(0.0503)	(0.0511)	(0.0525)	(0.0050)	(0.0049)	(0.0051)
Controls	Y	Y	Y	Y	Y	Y
Firm FE	N	N	N	Y	Y	Y
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 lagged ROA, leverage, sales, sales growth, R&D, and capital expenditures, as well as firm, industry, quarter, and industry-quarter fixed effects. Standard errors are clustered at the firm level (*** $p<0.01$, ** $p<0.05$, * $p<0.1$).

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

Table 6. AI Washing and Institutional Ownership

This table presents results from firm-quarter panel regressions examining the relationship between AI talk, walk, and fund investment behavior. Dependent variables include the number of funds (CRSP portfolios) and the percent of market value held by AI-focused (Panel A) and all institutional funds (Panel B). Control variables include firm size, cash , R&D intensity, age, and capital expenditures. All regressions 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	(1)	(2)	(3)	(4)
Variable	Number of Funds		% Market Value Held by Funds	
$Talk_{t-1}$	0.055***	0.051***	0.004	0.004
	(0.012)	(0.013)	(0.004)	(0.004)
$Talk \times Washing_{t-1}$	-0.066***	-0.060***	-0.011**	-0.010**
	(0.014)	(0.014)	(0.005)	(0.005)
$Walk_{t-1}$	0.288***	0.283***	0.086	0.083
	(0.099)	(0.093)	(0.087)	(0.081)
$Size_{t-1}$		0.303***		0.037
		(0.062)		(0.026)
$Cash/Assets_{t-1}$		-0.527*		-0.045
		(0.308)		(0.099)
$R\&D_{t-1}$		-3.412		-2.126*
		(2.086)		(1.217)
Age_{t-1}		-0.099		0.011**
		(0.077)		(0.005)
$CAPEX_{t-1}$		0.037		-0.009
		(0.029)		(0.009)
Firm FE	Y	Y	Y	Y
Industry-Quarter FE	Y	Y	Y	Y
Adj R^2	0.742	0.747	0.681	0.683
Observations	19,973	19,910	19,966	19,909

Panel B: All Funds	(5)	(6)	(7)	(8)
Variable	Number of Funds		% Market Value Held by Funds	
$Talk_{t-1}$	3.342**	2.737*	1.495*	1.485*
	(1.658)	(1.439)	(0.822)	(0.816)
$Talk \times Washing_{t-1}$	-3.320*	-2.418*	-1.320	-1.300
	(1.726)	(1.464)	(0.863)	(0.854)
$Walk_{t-1}$	14.175**	11.941**	5.744**	5.603*
	(6.639)	(5.522)	(2.905)	(2.905)
$Size_{t-1}$		52.735***		2.509*
		(7.834)		(1.455)
$Cash/Assets_{t-1}$		75.742***		12.468**
		(24.533)		(4.976)
$R\&D_{t-1}$		-332.320**		29.063
		(143.542)		(53.934)
Age_{t-1}		-0.051		-2.886
		(13.080)		(2.617)
$CAPEX_{t-1}$		11.504***		1.126*
		(3.065)		(0.656)
Firm FE	Y	Y	Y	Y
Industry-Quarter FE	Y	Y	Y	Y
Adj R^2	0.932	0.937	0.749	0.749
Observations	19,973	19,910	19,966	19,909

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 models include firm 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)

VARIABLES	(1) CAR(-1, 1)$_t$	(2) CAR(-1, 1)$_t$	(3) BHAR(0, 180)$_t$	(4) BHAR(0, 180)$_t$	(5) BHAR(0, 270)$_t$	(6) BHAR(0, 270)$_t$	(7) BHAR(0, 360)$_t$	(8) BHAR(0, 360)$_t$
Talk$_{t-1}$	0.251***	0.254***	0.644*	0.869**	0.959*	1.473**	0.356	1.393*
	(0.086)	(0.086)	(0.354)	(0.412)	(0.550)	(0.637)	(0.855)	(0.772)
Talk×Washing$_{t-1}$	-0.157	-0.126	-0.510	-0.882	-1.117	-2.042**	-0.750	-2.557**
	(0.159)	(0.159)	(0.559)	(0.597)	(0.863)	(1.007)	(1.039)	(1.202)
Walk$_{t-1}$	-0.455	-0.320	1.685	5.413*	25.379***	24.182***	28.995***	25.199***
	(0.388)	(0.343)	(3.214)	(2.773)	(6.201)	(5.686)	(9.751)	(7.639)
Size$_{t-1}$		-0.156		-11.265***		-16.906***		-22.389***
		(0.367)		(2.242)		(3.363)		(4.376)
Cash/Assets$_{t-1}$		2.257*		-6.737		-21.247*		-20.934
		(1.362)		(8.456)		(11.815)		(14.657)
R&D$_{t-1}$		25.172**		279.780***		413.782***		532.429***
		(12.595)		(72.974)		(100.585)		(120.719)
Age$_{t-1}$		-0.043		-2.145**		-3.068**		-4.831**
		(0.242)		(0.952)		(1.495)		(2.177)
CAPEX$_{t-1}$		-0.569***		-4.086***		-6.501***		-7.940***
		(0.170)		(0.830)		(1.267)		(1.534)
SUE$_{t-1}$		0.389***		-0.058		0.081		-0.105
		(0.054)		(0.246)		(0.347)		(0.411)
ROA$_{t-1}$		24.849***		-1.439		-16.077		-16.258
		(6.197)		(24.620)		(29.936)		(33.043)
MTB$_{t-1}$		-0.895***		-6.342***		-9.501***		-12.073***
		(0.099)		(0.623)		(0.799)		(0.921)
Firm FE	Y	Y	Y	Y	Y	Y	Y	Y
Industry-Quarter FE	Y	Y	Y	Y	Y	Y	Y	Y
Adj R^2	0.0150	0.0363	0.172	0.220	0.221	0.286	0.251	0.332
Observations	18,184	17,659	16,922	16,461	16,953	16,464	15,854	15,397

Table 8. 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* × *High Delta*, which captures the differential effect for high-delta firms in the post-ChatGPT period relative to low-delta firms. All specifications include firm 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 $Walk_t$		Talk-Walk Gap_t	
	(1)	(2)	(3)	(4)	(5)	(6)
$Post_t$ × High Delta	0.946***	0.878***	0.070**	0.035	0.875***	0.843***
	(0.269)	(0.263)	(0.031)	(0.028)	(0.256)	(0.257)
$Size_{t-1}$		0.249*		0.122***		0.127
		(0.136)		(0.042)		(0.124)
$Cash/Assets_{t-1}$		-1.005*		-0.100		-0.904*
		(0.576)		(0.151)		(0.494)
$R\&D_{t-1}$		-9.391		-0.619		-8.772
		(8.782)		(1.482)		(8.277)
Age_{t-1}		0.090		0.002		0.088
		(0.061)		(0.008)		(0.062)
$CAPEX_{t-1}$		-0.015		-0.006		-0.010
		(0.048)		(0.011)		(0.047)
Firm FE	Y	Y	Y	Y	Y	Y
Quarter FE	N	N	N	N	N	N
Industry-Quarter FE	Y	Y	Y	Y	Y	Y
Adj R^2	0.208	0.219	0.944	0.949	0.207	0.217
Observations	7,038	6,599	7,038	6,599	7,038	6,599

Appendix

Table A1. AI Talk and AI Walk: Robustness Test 1

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 variabls as walk and talk measures. Control variables include lagged, sales, cash, R&D, age, and capital expenditures, as well as firm, industry, quarter, and industry-quarter fixed effects. Standard errors are clustered at the firm level (*** p<0.01, ** p<0.05, * p<0.1).

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

Table A2. AI Talk and AI Walk: Robustness Test 2

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 lagged ROA, leverage, sales, sales growth, R&D, and capital expenditures, as well as firm, industry, quarter, and industry-quarter fixed effects. Standard errors are clustered at the firm level (*** p<0.01, ** p<0.05, * p<0.1).

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

Figure A1. 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.

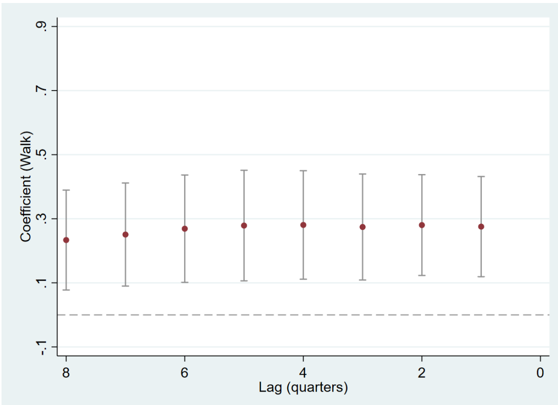

(a) Lagged AI Walk on Patent Counts

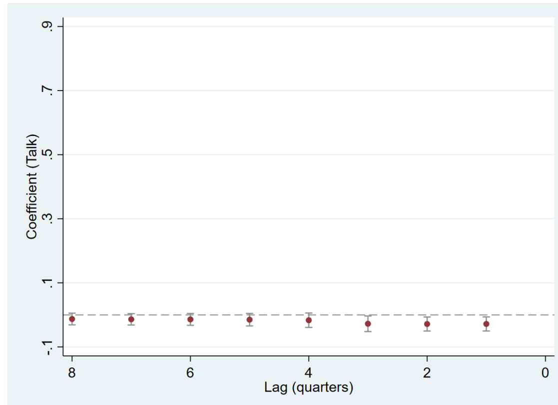

(b) Lagged AI Talk on Patent Counts

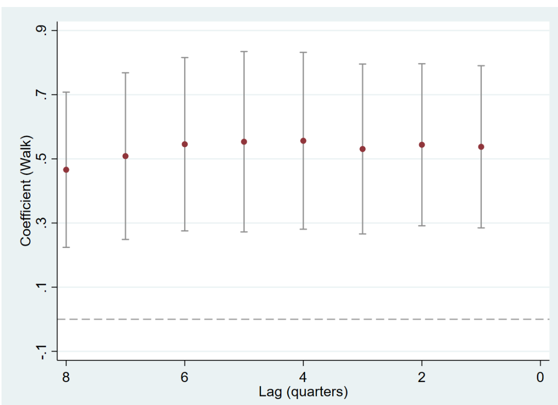

(c) Lagged AI Walk on Patent Value

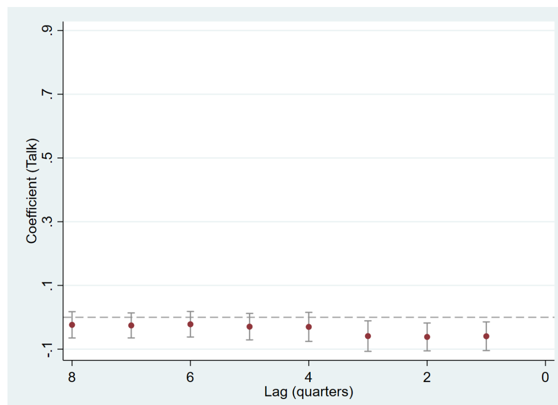

(d) Lagged AI Talk on Patent Value

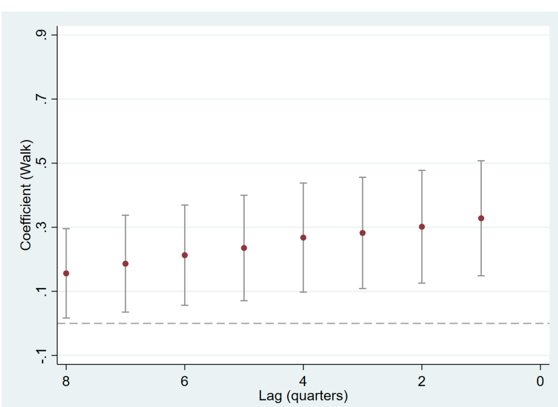

(e) Lagged AI Walk on Patent Citations

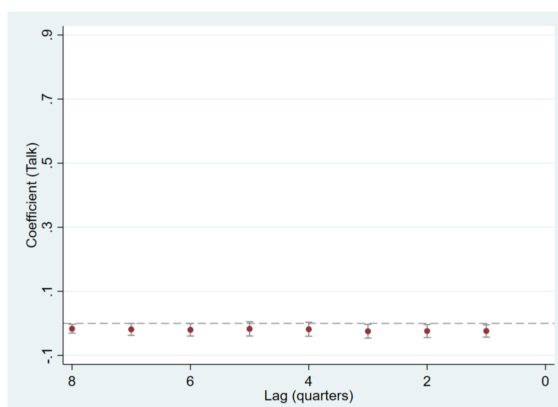

(f) Lagged AI Talk on Patent Citations

Figure A2. Distributed Lag Analysis on Fund Holdings

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

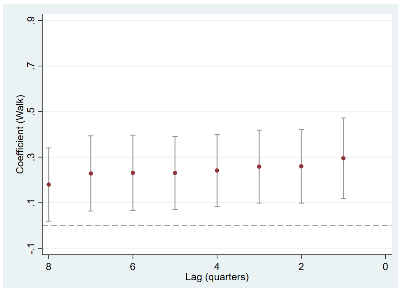

(a) Lagged AI Walk on Number of AI Funds

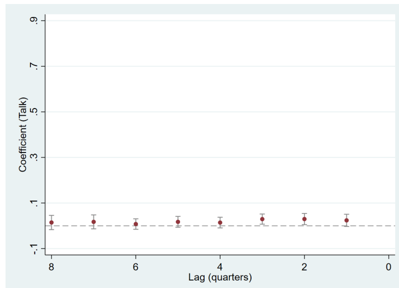

(b) Lagged AI Talk on Number of AI Funds

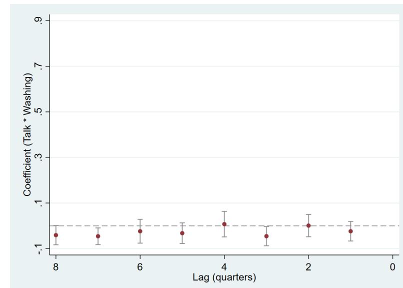

(c) Lagged AI Talk × Washing on Number of AI Funds

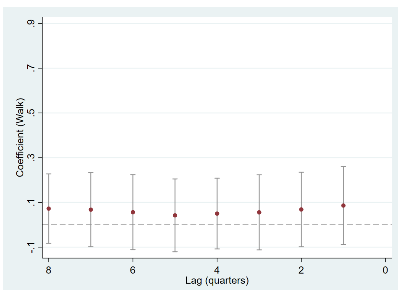

(d) Lagged AI Walk on % Market Value Held by AI Funds

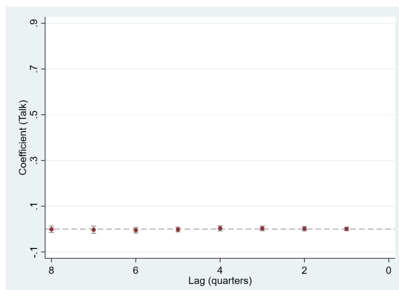

(e) Lagged AI Talk on % Market Value Held by AI Funds

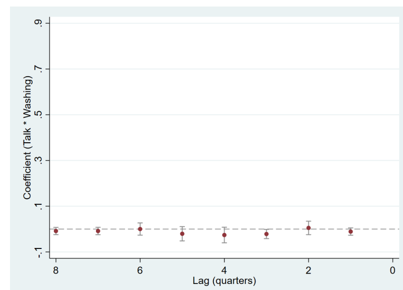

(f) Lagged AI Talk × Washing on % Market Value Held by AI Funds

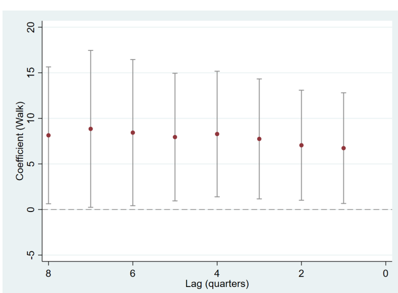

(g) Lagged AI Walk on % Market Value Held by All Funds

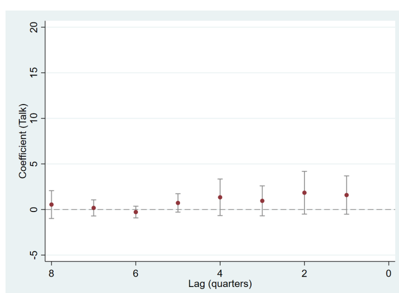

(h) Lagged AI Talk on % Market Value Held by All Funds

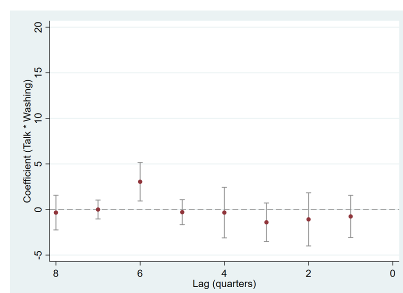

(i) Lagged AI Talk × Washing on % Market Value Held by All Funds

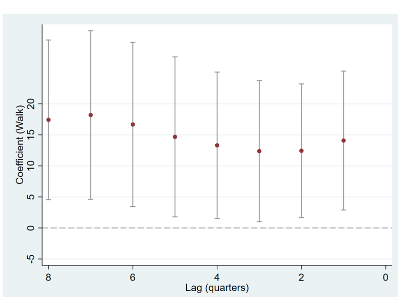

(j) Lagged AI Walk on Number of All Funds

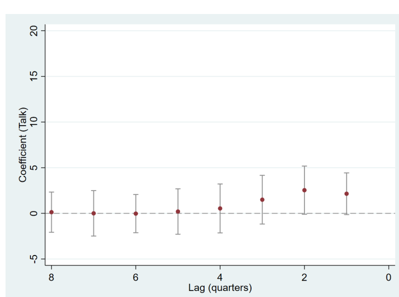

(k) Lagged AI Talk on Number of All Funds

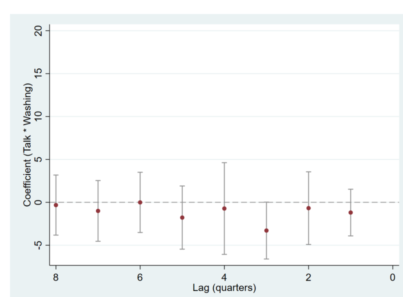

(l) Lagged AI Talk × Washing on Number of All Funds